

The pfSense Documentation

© 2025 Electric Sheep Fencing LLC & Rubicon Communications LLC

Netgate

Dec 23, 2025

CONTENTS

1 Preface 2

2 Introduction 4

3 Releases 20

4 Product Manuals 356

5 Networking Concepts 357

6 Hardware 373

7 Installing and Upgrading 399

8 Configuration 459

9 Netgate® Nexus 528

10 Backup and Recovery 556

11 Interface Types and Configuration 596

12 User Management and Authentication 627

13 Certificate Management 646

14 Firewall 667

15 Network Address Translation 734

16 Routing 769

17 Bridging 784

18 Virtual LANs (VLANs) 796

19 Multiple WAN Connections 806

20 Virtual Private Networks 826

21 L2TP VPN 939

22 Services 944

23 DHCP 1006
24 DNS 1012
25 Traffic Shaper 1019
26 Captive Portal 1037
27 High Availability 1060
28 System Monitoring 1076
29 Diagnostics 1172
30 Packages 1201
31 Virtualization 1383
32 Wireless 1385
33 Cellular Wireless 1406
34 Troubleshooting 1415
35 pfSense® software Configuration Recipes 1544
36 Menu Guide 1920
37 Glossary of Terms 1929
38 Development 1930
39 References 1974
40 Licensing 1996
41 Configuration Recipes 2476
42 Additional Commercial Resources 2477
Index 2478

Thoroughly detailed information and continually updated instructions on how to best operate pfSense® software.

CHAPTER

ONE

PREFACE

1.1 Acknowledgements

This documentation, and the pfSense® project itself, would not be possible without a great team of developers, contributors, customers, and a wonderful community. The project has received code contributions from several hundred individuals. Thousands more have done their part supporting the project by helping others on the forum, social media, and other platforms. And even more have contributed by purchasing hardware, support, and services. Our thanks go out to everyone who has done their part to make the project the great success it has become.

1.1.1 pfSense Developers

The current active pfSense software development team includes the following members (in alphabetical order by surname):

- Glen Barber
- Renato Botelho do Couto
- Leon Dang
- Brad Davis
- Peter Grehan
- Mateusz Guzik
- Reid Linnemann
- Christian McDonald
- Kris Molinari
- Jim Pingle
- Kristof Provost
- Luiz Otavio O Souza
- Steve Wheeler

We also want to give thanks to former project members, significant community contributors, and all FreeBSD developers who have assisted considerably with pfSense project development. Their time and effort throughout the last 15 years is meaningful, and we appreciate their contributions.

1.2 Feedback

Netgate encourages feedback for this documentation and the pfSense® software distribution. Please send suggestions, criticism and/or praise using the feedback forms at the bottom of each page.

For general feedback related to the pfSense project, please post to the forum. Links to these resources can be found at https://www.netgate.com/support/contact-support.

Welcome to *The pfSense Documentation*, written by the pfSense® project team and including contributions from community members.

This set of documents covers topics ranging from the installation process and basic configuration to advanced networking and firewall implementations using this popular open source firewall and router software distribution.

This is designed to be a friendly guide to common networking and security tasks along with a thorough reference for the capabilities of pfSense software. These documents cover the following topics (and more!):

- An introduction to pfSense software and its features.
- Firewall design and hardware planning.
- Installing and upgrading pfSense software.
- Using the web-based configuration interface.
- Backing up and restoring the firewall configuration.
- Firewall fundamentals including defining and troubleshooting rules.
- Port forwarding and Network Address Translation (NAT).
- General networking and routing configuration.
- Virtual LANs (VLANs), Multi-WAN, and Bridging.
- Virtual Private Networks using IPsec and OpenVPN.
- Traffic shaping using ALTQ or Limiters.
- Wireless networking configuration.
- Captive Portal setup.
- High Availability using redundant firewalls.
- Various network-related services.
- Firewall monitoring, logging, traffic analysis, sniffing, packet capturing, and troubleshooting.
- Software package and third-party software installations.

There is also a *Menu Guide* with all standard menu choices available in the pfSense software GUI.

CHAPTER

TWO

INTRODUCTION

2.1 What does pfSense stand for/mean?

The early tag line for the pfSense open source project was "making sense of PF", referring to the packet filter technology at the core of the project.

PF in FreeBSD can perform many of the basic packet filtering and QoS firewall tasks that pfSense software provides, however, pfSense software makes it easier to manage, monitor, and maintain. It accomplishes this by layering an easy-to-use GUI and customized services on top of the operating system and relevant packages, resulting in a complete firewall/router/VPN solution that is capable of much more than the sum of the underlying components.

2.2 What is pfSense® Plus Software?

Netgate announced the creation of pfSense Plus software, and the renaming of the open-source project to pfSense Community Edition (CE), in January 2021. The rationale was simple: The existence of pfSense Plus software would allow Netgate to add advanced features required by business customers. In the time since that announcement, a number of premium capabilities have been added to pfSense Plus software that are not available in pfSense CE software.

2.2.1 Benefits of pfSense Plus Software

More Frequent Software Updates

One of the most significant differences is the release cadence.

Multiple Releases per Year

pfSense Plus software has major updates scheduled three times per year, and additional point releases when required. This allows Netgate to keep pfSense Plus software closely in sync with the many changes and updates that are made 'upstream,' including in FreeBSD.

See also:

Software Release Schedule

Cryptography and VPN Acceleration

pfSense Plus software incorporates a number of capabilities that improve the performance of VPN connectivity.

See also:

VPN Scaling

These exclusive capabilities include:

OpenVPN Data Channel Offload (DCO)

This provides huge performance gains when processing encrypted OpenVPN data by reducing the amount of context switching that happens for each packet.

See also:

OpenVPN Data Channel Offload (DCO)

Intel IPsec Multi-Buffer (IIMB)

This increases VPN performance on Intel, AMD and ARM platforms where extended CPU instructions are present by replacing some cryptographic functions provided by the kernel with accelerated functions that utilize those instructions.

See also:

IPsec-MB

Intel QuickAssist Technology (QAT)

This is an Intel-specific hardware acceleration technology that significantly increases performance, using asynchronous processing, for many cryptographic operations.

See also:

Cryptographic Hardware

SafeXcel cryptographic accelerator

This is an acceleration technology present on some ARM platforms, such as the Netgate 1100 and 2100 appliances.

See also:

Cryptographic Hardware

CESA

This is an acceleration technology present on some ARM platforms such as the Netgate 3100 appliance.

See also:

Cryptographic Hardware

AWS VPC VPN Connection Wizard add-on package

This add-on package automatically creates a VPN tunnel and BGP configuration to communicate with an Amazon AWS VPC.

See also:

AWS VPC Wizard

IPsec Profile Wizard add-on package

This add-on package creates IPsec configuration profiles for Apple devices (iOS and macOS), and IPsec import script bundles for Windows devices.

See also:

IPsec Export Package

OpenVPN Client Import add-on package

This add-on package Imports a unified OpenVPN client configuration file as exported by an OpenVPN server.

See also:

OpenVPN Client Import Package

Additional Features

Additional premium features found in pfSense Plus software include:

ZFS Boot Environment (BE) Management in webConfigurator

This feature makes it easier to take snapshots of key file system areas, resulting in safer upgrades and major changes. If the user encounters problems with an upgrade or configuration change, the firewall can be 'rolled back' to an earlier known good state.

See also:

ZFS Boot Environments (Plus Only)

ZFS dashboard widget (to track status of disks using ZFS)

This feature allows easy monitoring of disks using ZFS.

See also:

ZFS Dashboard Widget

CARP mode (multicast or unicast)

This is an option to choose how CARP (High Availability) operates, either in multicast or unicast mode. Some environments (including virtualization) don't work well, or not at all, with multicast mode. pfSense CE software can only utilize multicast.

See also:

VIP Configuration Options

Ethernet (Layer 2) Filtering Rules

This feature is experimental rule-based pass/block filtering of packets based on Ethernet (Layer 2) header attributes (e.g. MAC addresses). These rules are processed before other (L3) rules in the inbound direction, and after those rules outbound.

See also:

Ethernet (Layer 2) Rules

LDAP Client Certificates

This feature enables pfSense software to use a client certificate when communicating with an LDAP server to identify itself when using an encrypted transport mode.

See also:

LDAP Authentication Servers

GUI Options for WAN 802.1X Authentication Bridging and VLAN 0 PCP Tagging

These options allow directly connecting to certain ISP networks which typically require specific devices at the edge, such as a modem with an authentication certificate.

See also:

WAN Connectivity with 802.1X Authentication Bridging and VLAN 0 PCP Tagging

Native Packet Flow Data Export for NetFlow/IPFIX

Starting with pfSense Plus software version 24.03 the firewall can directly export NetFlow v5 and IPFIX traffic flow data to one or more collectors using the `pflow(4)` feature in PF. The data is collected directly from firewall states and does not require a separate daemon, service, or add-on package.

See also:

Firewall Packet Flow Data

Capabilities For Netgate Hardware

There are also several capabilities in pfSense Plus software that are unique to the appliance hardware found in the Netgate Store.

These include:

- ARM64 devices (for Netgate's ARM-based appliances)
- The Firmware Update add-on package
- MMC Utilities package
- Specialized hardware such as status LEDs, reset buttons, switches, and hardware watchdog devices
- Default optimized configurations for Netgate hardware appliances

2.3 Why FreeBSD?

Numerous factors came under consideration when choosing a base operating system for the project. This section outlines the primary reasons for selecting FreeBSD.

2.3.1 Wireless Interfaces

Wireless interfaces are an important feature for some users of pfSense software. In 2004, wireless interfaces in OpenBSD were very limited compared to FreeBSD. OpenBSD did not include drivers or security protocols and offered no plans for their implementation. To this day, FreeBSD surpasses the wireless capabilities of OpenBSD.

2.3.2 Network Performance

Network performance in FreeBSD is significantly better than that of OpenBSD. For small to mid-sized deployments, this generally does not matter; upper scalability is the primary issue in OpenBSD. One pfSense® developer managing several hundred OpenBSD firewalls using PF was forced to switch his high load systems to PF on FreeBSD to handle the high packets per second rate required by portions of his network. The network performance in OpenBSD has improved since 2004, but limitations still exist.

Multiprocessor capabilities for PF in FreeBSD allows for greater scalability and is utilized by pfSense software as seen in this network performance analysis: https://github.com/gvnn3/netperf/blob/master/Documentation/netperf.pdf.

2.3.3 Familiarity and ease of fork

The code for m0n0wall was based on FreeBSD, and pfSense software forked from m0n0wall. Changing the base operating system would require prohibitively large modifications and could have introduced limitations from other operating systems, requiring features to be removed or altered.

2.4 Common Deployments

pfSense® software can meet the needs of nearly any type and size of network environment, from a SOHO to datacenter environments. This section outlines the most common deployments.

2.4.1 Perimeter Firewall

The most common deployment of pfSense software is a perimeter firewall. pfSense software accommodates networks requiring multiple Internet connections, multiple LAN networks, and multiple DMZ networks. BGP (Border Gateway Protocol), connection redundancy, and load balancing capabilities are configurable as well.

See also:

These advanced features are further described in *Routing* and *Multiple WAN Connections*.

2.4.2 LAN or WAN Router

pfSense software configured as a LAN or WAN router and perimeter firewall is a common deployment in small networks. LAN and WAN routing are separate roles in larger networks.

LAN Router

pfSense software is a proven solution for connecting multiple internal network segments. This is most commonly deployed with VLANs configured with 802.1Q trunking, described more in *Virtual LANs (VLANs)*. Multiple Ethernet interfaces are also used in some environments. High-volume LAN traffic environments with fewer filtering requirements may need layer 3 switches or ASIC-based routers instead.

WAN Router

pfSense software is a great solution for Internet Service Providers. It offers all the functionality required by most networks at a much lower price point than other commercial offerings.

2.4.3 Special Purpose Appliances

pfSense software can be utilized for less common deployment scenarios as a stand-alone appliance. Examples include: VPN appliance, Sniffer appliance, and DHCP server appliance.

VPN Appliance

pfSense software installed as a separate Virtual Private Network appliance adds VPN capabilities without disrupting the existing firewall infrastructure, and includes multiple VPN protocols.

Sniffer Appliance

pfSense software offers a web interface for the `tcpdump` packet analyzer. The captured `.cap` files are downloaded and analyzed in Wireshark.

See also:

For more information on using the packet capture features, see *Packet Capturing*.

DHCP Server Appliance

pfSense software can be deployed strictly as a Dynamic Host Configuration Protocol server, however, there are limitations of the pfSense software GUI for advanced configuration of the ISC DHCP daemon.

See also:

For more information on configuring the DHCP service on pfSense, see *DHCP*.

2.5 Interface Naming Terminology

All interfaces on pfSense® software can be assigned any name desired, but they all start with default names: WAN, LAN, and OPT.

2.5.1 WAN

Short for *Wide Area Network*, WAN is the untrusted public network outside the firewall. In other words, the WAN interface is the firewall's connection to the Internet or other upstream network. In a multi-WAN deployment, WAN is the first or primary Internet connection.

At a minimum, the firewall must have one interface, and that is WAN.

2.5.2 LAN

Short for *Local Area Network*, LAN is commonly the private side of a firewall. It typically utilizes a *private IP address* scheme for local clients. In small deployments, LAN is typically the only internal interface.

2.5.3 OPT

OPT or *Optional* interfaces refer to any additional interfaces other than WAN and LAN. OPT interfaces can be additional LAN segments, WAN connections, DMZ segments, interconnections to other private networks, and so on.

2.5.4 DMZ

Short for the military term *demilitarized zone*, DMZ refers to the buffer between a protected area and a war zone. In networking, it is an area where public servers are reachable from the Internet via the WAN but isolated from the LAN. The DMZ keeps the systems in other segments from being endangered if the network is compromised, while also protecting hosts in the DMZ from other local segments and the Internet in general.

Warning: Some companies misuse the term "DMZ" in their firewall products as a reference to 1:1 NAT on the WAN IP address which exposes a host on the LAN. For more information, see *1:1 NAT on the WAN IP, aka "DMZ" on Consumer Gateways*.

2.5.5 FreeBSD interface naming

The name of a FreeBSD interface starts with the name of its network driver. It is then followed by a number starting at `0` that increases incrementally by one for each additional interface sharing that driver. For example, a common driver used by Intel gigabit network interface cards is `igb`. The first such card in a firewall will be `igb0`, the second is `igb1`, and so on. Other common driver names include `cxl` (Chelsio 10G), `em` (Also Intel 1G), `ix` (Intel 10G), `bge` (various Broadcom chipsets), amongst numerous others. If a system mixes an Intel card and a Chelsio card, the interfaces will be `igb0` and `cxl0` respectively.

See also:

Interface assignments and naming are further covered in *Installing and Upgrading*.

2.6 Finding Information and Getting Help

This section offers guidance on finding information in this documentation, on pfSense® software in general, as well as providing further resources.

2.6.1 Finding Information

The search function on the documentation is the easiest way to find information on a specific topic. The most common features and deployments of pfSense software are covered in this documentation. When reading the HTML version of the documentation, the search function is in the upper left of the page. When reading an eBook style copy, consult the documentation for the book reader software for information on how to search.

There is a wealth of additional information and user experiences available on the various Netgate websites. The best way to search the sites is a Google search appending `site:netgate.com` to the query. This will search the website, forum, documentation, etc. which are all official sources of information.

2.6.2 Getting Help

A help icon is available on almost every page,  , and links to the associated page in documentation.

Netgate offers several other ways to get help with pfSense software, including the Netgate Forum and this documentation. There is also a pfSense subreddit where community members can assist each other. More information can be found on the Netgate website at Obtaining Support. Many of these links are reachable from the Help menu in the GUI.

2.7 Comparison to Commercial Alternatives

The question of security and assistance vs. commercial alternatives comes up from time to time. The history of this project since its inception in 2004 proves we're as secure as any, and better than many, commercial alternatives. The experiences of our customers proves not only can we match the service of any commercial firewall vendor, we exceed it. This page serves to debunk the common myths when comparing to commercial alternatives.

2.7.1 "Hardware" firewalls are better myth

Commercial firewall companies' marketing departments have done a fine job ingraining the myth of "hardware firewalls" into some people's minds. The reality is there is no such thing as a "hardware firewall." All firewalls are hardware that runs software. Most commercial firewalls are based on BSD (same as pfSense®) or Linux. Numerous commercial firewalls run many of the same underlying software programs that pfSense software uses. Many commercial alternatives run on x86 hardware that's no different from what people use for pfSense software. In fact many people have loaded pfSense software on hardware that used to run their commercial firewall, including WatchGuard, Nortel, Barracuda and more.

2.7.2 Open source is insecure myth

Some people are of the mindset that because the source is open, it's insecure because everyone can see how it works. Anyone who has paid any attention to security over the past 20 years knows the absurdity of that statement. No software relies on the obscurity of source code for security. If there was any truth in that, Microsoft Windows would be the most secure OS ever created, when the reality is all the open source operating systems (all the BSDs and Linux) have security track records that are worlds better than Windows'. History proves the same applies to any software. Internet Explorer is continually hit with major security holes that many times take weeks to patch while they're being exploited in the wild, while open source browsers Firefox, Chrome and others have had significantly better security track records.

The widespread UPnP vulnerabilities announced in 2013 affecting over 300 commercial products is another good example. The vendors of hundreds of commercial products made extremely basic security mistakes, shipping with absurdly insecure defaults, and shipping outdated software. That's never been an issue with pfSense software. That's only one example of where pfSense software has done a better job than many commercial vendors.

2.7.3 Commercial alternatives have better technical assistance myth

With some open source projects, it's true that a user is stuck if they need help. However, Netgate offers commercial assistance for pfSense software, Netgate TAC, that rivals anything other commercial vendors offer.

2.8 Can pfSense software meet regulatory requirements

Prospective pfSense® users commonly inquire about the ability to meet security requirements applicable to their specific environments. Some of those include PCI, SOX, GLBA, HIPAA, amongst numerous other similar regulations for publicly traded companies, financial institutions, healthcare institutions, and others.

There are numerous companies in many regulated industries using pfSense software that pass their audits with no problems, including all the aforementioned regulations and standards, amongst others. However, it's important to keep in mind that a firewall is a small portion of the security infrastructure, and those regulations are more about policies, procedures, and configuration than the actual products being used.

So yes, pfSense software *can* meet regulatory requirements, but that is dependent on configuration, policies, procedures, amongst other things - there is no compliance silver bullet. There may be circumstances specific to one company that

make another product a better fit for compliance (or other) reasons, but that's true of all commercial and open source solutions, there is no one product that is a perfect fit for everyone.

2.9 Can I sell pfSense software

Many consulting companies offer solutions based on pfSense® software to their customers. A business or individual can load pfSense software for themselves, friends, relatives, employers, and, yes, even customers, so long as the Trademark Guidelines and Apache 2.0 license requirements as detailed on the website are obeyed by all parties involved.

What can not be offered is a commercial redistribution of pfSense® software, for example the guidelines do not permit someone to offer "Installation of pfSense® software" as a service or to sell a device preloaded with pfSense® software to customers without the prior express written permission of ESF pursuant to the trademark policy.

Example 1
A consultant may offer firewall services (e.g. "Fred's Firewalls"), without mentioning pfSense® software or using the logo in their advertising, marketing material, and so on. They can install pfSense® software and manage it for their customers.

Example 2
Fred's Firewalls may make a customized distribution pfSense® software with their own name and logo used in place of the pfSense marks. They can use the pfSense marks to truthfully describe the origin of the software, such as "Fred's Firewall software is derived from the pfSense CE source code." Even though Fred's Firewall is based on pfSense® software, it **cannot** be referred to as "pfSense® software" since it has been modified.

Example 3
Fred's Firewalls may sell their customized firewall distribution preloaded on systems to customers, so long as the relationship to pfSense software is clearly stated.

The Apache 2.0 license only applies to the software and **not** the pfSense name and logo, which are trademarks and may not be used without a license. Reading and understanding the trademark policy document is required before one considers selling pfSense software.

2.9.1 Contributing Back to the Project

We ask anyone profiting by using pfSense software to contribute to the project in some fashion. Ideally with the level of contributions from a business or individual corresponding to the amount of financial gain received from use of pfSense software. Many paths exist for resellers and consultants to contribute. For the long term success of the project this support is critically important.

- Purchase hardware and merchandise from the Netgate Store.
- Become a Netgate Partner to resell Netgate hardware preloaded with pfSense software.
- Development contributions - Dedicate a portion of internal developers' time to open source development.
- Help with support and documentation - Assisting users on the forum and mailing list, or contributing documentation changes, aides the overall project.
- Netgate TAC subscription - Having direct access to the Netgate team for questions or deployment assistance helps ensure success.

2.9.2 Using the pfSense Name and Logo

The "pfSense" name and logo are trademarks of Electric Sheep Fencing, LLC.

The pfSense software source code is open source and covered by the Apache 2.0 license. That license only covers the source code and **not** our name and trademarks, which have restricted usage.

We think it is great that people want to promote and support the pfSense project. At the same time, we also need to verify that what is referred to as "pfSense" is a genuine instance of pfSense software and **not** modified in any way.

- The pfSense name and logo **MAY NOT** be used physically on a hardware device.
  - For example: A sticker, badge, etching, or similar rendering of the pfSense name or logo is NOT allowed.
- The pfSense logo **MAY NOT** be used on marketing materials or in other ways without a license, including references on websites.
- The pfSense name **MAY** be used to describe the case that a product is based on a pfSense distribution, but the designated product name may *not* include pfSense or a derivative. Basically stating facts regarding product origin is acceptable. Anything that implies that a product is endorsed by or made by ESF or the pfSense project is *not* allowed.

  Examples:
  - "Blahsoft Fireblah based on pfSense software" – Acceptable
  - "Blahsoft pfSense Firewall" – NOT Allowed
- **ONLY** an **UNMODIFIED** version of pfSense software can still be called "pfSense software".
  - If the source code has been changed, compiled/rebuilt separately, included extra file installations such as themes or add-on scripts, or any other customizations, it can **not** be called "pfSense software", it must be called something else.
  - Trademark protection aside, this requirement preserves the integrity and reputation of the pfSense project. It also prevents unverified changes that may be questionably implemented from being attributed to pfSense.
- If a pfSense distribution is modified, the resulting software **CANNOT** be called "pfSense" or anything *similar*. The new name must be distinct from pfSense. Trademark law does not allow use of names or trademarks that are confusingly similar to the pfSense Marks. This means, among other things, that law forbids using a variation of the pfSense Marks, their phonetic equivalents, mimicry, wordplay, or abbreviation with respect to similar or related projects, products, or services (for example, "pfSense Lifestyle," "PFsense Community," "pf-Sense Sensibility," "pfSensor", etc., **all** infringe on ESF's rights).

  Examples:
  - "pfSomething", or "somethingSense" – INFRINGING references
  - "ExampleWall", "FireWidget" – NON-Infringing references
- The "pfSense" name **MAY NOT** be used in a company name or similar. A company **CANNOT** be named "pfSense Support, Ltd" or "pfSense Experts, LLC", or use it in a domain name or subdomain reference. However, the company can state support for pfSense software, offer training for pfSense software, etc.
- There **MUST** be a distinction between a company name and pfSense or Electric Sheep Fencing, LLC. No relationship or endorsement can be stated or implied between the two companies, unless we have explicitly licensed and agreed to such a statement.

2.10 Contact with Netgate Servers

For several essential and optional services, such as checking for updates, installing updates and packages, and the Auto Config Backup service, pfSense® software makes outgoing connections to servers owned and operated by Netgate.

This document identifies these connections with an explanation of why it makes contact, what is transmitted in both directions, and how administrators can control these operations.

Each of the services described also generates server logs which Netgate discards after **365 days** (1 year). Netgate uses these logs to monitor performance, detect problems, and gather usage statistics while also satisfying GDPR requirements.

2.10.1 TAC Information

Netgate offers a number of assistance options for both pfSense and TNSR software users. See https://www.netgate.com/support/ for details.

If administrators activate the **Netgate Service and Support** widget on the pfSense software GUI Dashboard, pfSense software will query Netgate servers to obtain the current subscription status, a list of web links, and other information to help administrators make the best use of their options for assistance. This query occurs no more than once every 24 hours.

Note: This widget is active by default on installations of pfSense Plus software, but it can be removed at any time. It is not active by default on pfSense CE software.

Contracts are tied to individual devices so Netgate must positively identify these devices when providing service. This is not easy to do: Serial numbers are not consistent across different hardware manufacturers and building a firewall running pfSense software from scratch on "white box" hardware is common. To overcome this problem and allow purchasers of Netgate services to be sure that no one can "spoof" their contract, Netgate uses a "Netgate Device ID" (NDI) which the device displays on the dashboard and in console/ssh menu banner.

The NDI is a hashed SHA256 value based on available information on the system. Netgate **cannot** reverse engineer from the NDI how it was generated, what information was gathered, or what values were used. Even if it were practical to reverse-hash SHA256 (it is not) only a string of meaningless characters would result. Netgate knows only that NDI "xyz" has been assigned to contract "abc".

Information that is transmitted via the HTTPS connection to Netgate servers:

- Netgate ID

Server log information retained by Netgate for up to 1 year, then deleted:

- IP address

If administrators prefer not to have pfSense software transmit this information, remove or do not install the **Netgate Service and Support** Dashboard widget.

See also:

- *Managing Widgets*
- *Available Widgets*

2.10.2 Software Updates and Packages

To check for updates to pfSense software, navigate to **System > Update** or, if the **System Information** widget is active and configured to check for updates, visit the Dashboard. To install new packages or check for updates to installed packages, navigate to **System > Package Manager**.

These actions cause pfSense software to connect to Netgate servers and update its repository information and package metadata. The first connection is made to establish authorization level and request a list of any available updates.

To determine which updates are available for a given installation, pfSense software transmits the architecture (amd64, ARM64, etc.) to the web service along with a list of the installed software components and their versions. The server can then respond with a list of any available updates (For example "stable", "development", etc.) that are appropriate for the architecture type and authorized for the end-user installation (plus, community edition, etc.).

Similar queries are made when installing new packages.

pfSense software will request updated repository information any time an administrator navigates to **System > Package Manager**, **System > Update** or when the **System Information** Dashboard widget is active and the option to **Disable the dashboard auto-update check** is *unchecked* under **System > Update** on the **Update Settings** tab (*Dashboard Check*). This option is *unchecked* by default.

Additionally, pfSense software automatically updates this metadata once per day. The timing is randomized and can be anywhere from 1:00 A.M. local time up to 24 hours later.

Information that is transmitted via the HTTPS connection to Netgate servers:

- Netgate ID
- Serial Number
- pfSense software and installed package versions
- Architecture
- Configured language (locale)
- Platform information (Netgate products, virtualization platform, cloud provider, etc.)

Server log information retained by Netgate for up to 1 year, then deleted:

- IP address
- Netgate ID
- Serial Number
- pfSense software and installed package versions
- Architecture
- Platform information (Netgate products, virtualization platform, cloud provider, etc.)

See also:

- *Dashboard Check*
- *Upgrade Guide*
- *Packages*
- *Installation Feedback*

2.10.3 Auto Config Backup (ACB)

Netgate offers an automatic configuration backup service, known as Auto Config Backup or ACB, which is free for anyone to use.

When pfSense software makes a backup via ACB, either automatically or manually, it encrypts the `config.xml` configuration file using AES-256 and a key provided by the administrator with options suitable for file encryption.

pfSense software then transmits this encrypted file to https://acb.netgate.com where it is stored for future retrieval.

pfSense software also transmits a unique and random "device key" to privately identify which backups belong to a given installation.

This device key must be unique to each firewall backing up its configuration. To do this the key is automatically generated by taking the SHA256 hash of the SSH service keys on the installation. This allows ACB to reliably and properly identify ownership of each backup in an anonymous way.

The encryption key never leaves the firewall in plain text. It exists within the `config.xml` data, but has already been encrypted. Netgate **cannot** decrypt any stored configuration file, nor link it to any particular firewall serial number or identifier.

The ACB service is not active by default. If administrators prefer not to have this information transmitted, disable or do not configure the Automatic Configuration Backup feature.

Information that is transmitted via the HTTPS connection to Netgate servers:

- Netgate ID
- Device key
- AES-256 encrypted configuration
- Timestamp
- Optional hint (in case the device key is lost)
- pfSense software version (in case the ACB data format changes in future versions)

Server log information retained by Netgate for up to 1 year, then deleted:

- IP address
- Netgate ID
- pfSense software version

See also:

- *Automatic Configuration Backup Service*

2.10.4 Bogon Network List Updates

Bogon networks are those which should never be seen on the Internet, including reserved and unassigned IP address space.

If any interfaces are configured to **Block Bogon Networks**, then shortly after initial installation and once per month after that, pfSense software will contact Netgate servers to obtain an updated list of Bogon Networks.

The update script runs at 3:00 A.M. local time, and sleeps a random amount of time up to 12 hours before performing the update.

The frequency of Bogon network updates can be adjusted using the **Update Frequency** option for bogons under **System > Advanced** on the **Firewall & NAT** tab. Updates may be performed on a Daily, Weekly, or Monthly basis.

Information that is transmitted via the HTTPS connection to Netgate servers:

- pfSense software version
- Netgate ID (optional)
- Release (plus/community)

Server log information retained by Netgate for up to 365 days, then deleted:

- IP address
- pfSense software version
- Netgate ID (optional)
- Release

The Bogon update can either send or omit the Netgate ID based on the **Netgate Device ID** option under **Installation Feedback** on **System > Advanced**, **Miscellaneous** tab. The default is to send the ID.

See also:

- *Block Bogon Networks*
- *Interface Configuration*
- *Installation Feedback*

2.10.5 Copyright

When an administrator receives a new Netgate device, installs a new version of pfSense software, or performs a factory reset, the GUI displays a copyright notice. Once an administrator dismisses the notice, the GUI does not display it again until it detects an updated notice.

This notice is a reminder of what may and may not be done with pfSense software.

Before the GUI displays the notice, its HTML contents are retrieved from Netgate servers in the currently configured language (locale). There are different versions of the notice for Netgate provided hardware and for the "Community Edition" of pfSense software.

Information that is transmitted via the HTTPS connection to Netgate servers:

- Netgate ID
- Release (plus/community)

Server log information retained by Netgate for up to 365 days, then deleted:

- IP address
- Netgate ID
- Release

2.10.6 Contact

If there are any additional questions, please contact Netgate via telephone at +1 512 646-4100 or via email at sales@netgate.com.

The pfSense® Project is a free open source customized distribution of FreeBSD tailored for use as a firewall and router entirely managed by an easy-to-use web interface. No FreeBSD knowledge is required to deploy and use pfSense software. In fact, the majority of users have never used FreeBSD outside pfSense software. In addition to being a powerful, flexible firewall and routing platform, pfSense software includes a long list of related features. The pfSense software package system allows further expandability without adding bloat and potential security vulnerabilities to the

base distribution. pfSense software is a popular project with millions of downloads since its inception and hundreds of thousands of active installations. It has been proven successful in countless installations ranging from single computer protection in small home networks to thousands of network devices in large corporations, universities and other organizations.

To download the latest version, see previous versions, or to upgrade follow the guides located on the pfSense downloads page.

2.11 Project Inception

pfSense software was forked from the m0n0wall open source project in 2004. m0n0wall was focused specifically on providing a firewall/router for embedded devices and was sized for limited hardware resources. Initially pfSense software aimed at providing a firewall/router solution with an expanded set of capabilities on larger PC and server style hardware. pfSense software has continued to evolve over time, providing firewall, router, VPN, IDS/IPS, and more capabilities that work well on hardware from small home office size devices to large service provider size servers.

CHAPTER

THREE

RELEASES

This section contains information about past and present release of pfSense® software. This includes release notes and detailed version information.

- *General Release Information*
- *Current and Upcoming Supported Releases*
  - *pfSense Plus Software*
  - *pfSense CE Software*
- *Older/Unsupported Releases*
  - *pfSense Plus Software*
  - *pfSense CE Software*

3.1 25.11 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

Tip: Review the *Upgrade Guide* before performing any upgrade of pfSense software.

3.1.1 Errata

FreeBSD published security advisory FreeBSD-SA-25:12.rtsold for a remote command execution vulnerability in `rtsold`, which also affects pfSense software:

A fix for this issue is available via the *Recommended System Patches* function of the *System Patches Package* for users running pfSense Plus software version 25.11, pfSense Plus software version 25.07.1, and pfSense CE software version 2.8.1. Users of older versions can upgrade to a supported release or apply the patch on Redmine issue #16593.

3.1.2 General

- Base OS updated to FreeBSD 16-CURRENT
- OpenSSL upgraded to 3.5.3
- OpenSSH upgraded to 10.0p2
- PHP updated to 8.4
- VXLAN interface support has been re-added

3.1.3 Security

- Fixed anti-brute force protection bypass and potential denial of service #16312 #16314 pfSense-SA-25_09.sshguard

3.1.4 Endpoint-independent Port Restricted Cone Outbound NAT

This version includes partial experimental support for "Port Restricted Cone" endpoint-independent outbound NAT. This functionality must be manually enabled on a per-rule basis.

"Port Restricted Cone" NAT mappings attempt to preserve port and external address mappings for clients when speaking to multiple remote hosts, but in a dynamic way that does not rely on static port NAT. This helps avoid issues with multiple local clients using the same source port to the same remote host. These rules enable a client communicating with multiple remote hosts using the same source port to receive the same external IP address and port on outbound connections to any destination. This behavior facilitates use cases such as online gaming, peer-to-peer connections, and VoIP.

Inbound communication from a remote host and port is only possible after a local client initiates first contact to that remote host and port. While this is more secure, it is not yet capable of "full cone" NAT which some use cases may require such as certain types of online gaming.

See also:

- *Outbound NAT*
- *Endpoint-independent Port Restricted Cone NAT*
- *Configuring pfSense Software for Online Gaming*

3.1.5 pfSense Plus

Changes in this version of pfSense Plus software.

Backup / Restore

- Fixed: RRD data fails to restore via the ECL #16141

Captive Portal

- Fixed: Captive Portal Ethernet rules can block ARP #16264
- Fixed: Reserved DUMMYNET pipes for Captive Portal can overlap #16540

Configuration Backend

- Changed: Improve file handling of the configuration cache #16469

DHCP (IPv4)

- Changed: Upgrade to Kea 3.0.2 #16388
- Changed: Kea configuration parameter `client-class` is deprecated #16468

DHCP (IPv6)

- Fixed: Hostnames in Kea static leases may not be registered with DNS #16552

DNS Forwarder

- Fixed: PHP error in DNS Forwarder host overrides when the language is set to French #14741

DNS Resolver

- Changed: Update Unbound to 1.24.2 to address CVE-2025-11411 #16503

Dashboard

- Fixed: Manually verifying the boot environment makes config changes #15499
- Fixed: Thermal Sensors widget does not respect per-sensor threshold vales #16266

Diagnostics

- Fixed: Captive Portal `backwardsyncpassword` value not sanitized in status output #16339

Dynamic DNS

- Added: Preserve other record types when updating IPv4 or IPv6 using deSEC DDNS #12495
- Fixed: Dynamic DNS does not use preferred VIP in Gateway Group #16326
- Fixed: Custom Dynamic DNS services ignore the monitor interface #16368

Gateway Monitoring

- Fixed: Gateway monitoring daemon can unexpectedly use a CARP VIP as the source IP address #16322

Gateways

- Fixed: Gateway list order is incorrect until reloading page after moving entries and saving #16495

Hardware / Drivers

- Fixed: Netgate 2100/3100 LED controller not responding to `gpioctl` #16526
- Fixed: QLink/Marvell 41000 NIC bug #16248
- Added: Support 2.5G SGMII (SFP GPON ONT) in `bxe` driver (QLogic NetXtreme II BCM57810) #16321
- Fixed: e1000 network interfaces unexpectedly link at half-duplex #16449

IPsec

- Changed: Update strongSwan to 6.0.3 #16509

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: Cannot set RADVD router lifetime to `0` #16472

Installer

- Fixed: Configuration data restored during installation can be overwritten by hardware-specific default values #16176

Interfaces

- Added: VXLAN Interfaces #11732
- Added: Option to change QinQ ethertype to Service VLAN Tag #13340
- Fixed: Retain previous QinQ VLAN tag type value for existing entries on upgrade #13622

Logging

- Added: Option to disable logging of packets blocked due to unmatched IP options #16068
- Fixed: `syslogd` daemon can terminate when a remote log server refuses connections #16362

OpenVPN

- Fixed: Automatic IPv6 gateways for OpenVPN servers are created with the wrong gateway address #16351
- Fixed: OpenVPN servers will not start with DH parameter lengths less than 2048 #16421
- Fixed: OpenVPN does not include `client-to-client` in generated configuration for Peer-to-Peer SSL/TLS servers #16428

Operating System

- Fixed: `rc.savecore` errors prevent boot in ZFS #15613
- Fixed: Swap fails to activate when multiple swap partitions exist #16232

PHP Interpreter

- Changed: Upgrade PHP to 8.4 #16471

PPP Interfaces

- Changed: Sanitize PPPoE configuration parameters #16128
- Fixed: PPPoE interfaces using `if_pppoe` increase error counters due to normal ALTQ traffic shaping operations #16216
- Fixed: Virtual IP addresses on PPPoE interfaces using `if_pppoe` can prevent PPP session termination #16487

Package System

- Fixed: Error notification and log message `"Updating repositories metadata" returned error code 1` at boot due to `certctl` race condition #16341

Rules / NAT

- Added: Allow floating rules using the "match" action to match based on IP Options #16215
- Added: Block non-global NAT64 addresses by default #16241
- Changed: Refactor PF ruleset generation #16307
- Added: Avoid traffic stalls from unnecessary filter reloads #16308
- Fixed: NAT64 rules using `reply-to` do not forward packets #16429
- Fixed: Filter rule evaluation continues after matching a `match quick` rule #16475
- Added: Support state killing on gateway recovery for policy-routed traffic from the firewall itself #16502
- Added: Endpoint-independent Port Restricted Cone Outbound NAT rules #16517
- Fixed: NAT64 rules do not pass traffic when a gateway is specified for the rule #16546
- Changed: Update output and parsing behavior for PHP shell `pfanchordrill` #16551

System Logs

- Fixed: Log entries without a hostname can cause the system log to display in an unexpected manner #15411

Traffic Shaper (Limiters)

- Fixed: Using a Limiter on a rule with a gateway group limits all traffic through that gateway instead of the host IP address #15770

Translations

- Fixed: Korean locale configuration name is incorrect #16505

Unknown

- Fixed: pfSense Plus does not work with AWS new Instance Metadata Service (IMDSv2) #14772

Upgrade

- Fixed: PHP shell playback script `upgradeconfig` incorrectly replaces running configuration when Nexus is enabled #16179
- Added: Fix configuration artifacts on upgrade #16253

User Manager / Privileges

- Fixed: `sshguard` does not trigger for GUI logins from usernames containing unexpected characters #16312
- Fixed: GUI login events from usernames containing special characters or long strings can cause ambiguous or confusing log messages #16314

Virtual IP Addresses

- Fixed: Input validation text for deleting an IP Alias VIP within a CARP VIP subnet may reference incorrect VIP #16272

Web Interface

- Fixed: Boot Environment page fails to load if `pfsense:version` ZFS property contains newlines #16375
- Changed: Apple TouchID/FaceID probes for site icon files that do not exist #6727

XMLRPC

- Fixed: Membership to `admins` group is lost when synchronizing user changes via XMLRPC #16392

3.2 2.8.1 New Features and Changes

This is a maintenance software release primarily containing bug fixes for issues discovered in 2.8.0.

Warning: If upgrading from versions before CE 2.8.0, refer to *2.8.0 New Features and Changes* for additional important information.

Tip: Review the *Upgrade Guide* before performing any upgrade of pfSense software.

3.2.1 General

3.2.2 pfSense CE

Changes in this version of pfSense CE software.

Auto Configuration Backup

- Fixed: AutoConfigBackup entries show incorrect timestamps #16209

Dynamic DNS

- Fixed: Dynamic DNS fails to update AAAA record for Route53 and No-IP services #16249
- Changed: Update list of Dynamic DNS services that are only reachable over IPv4 for AAAA record updates #16251

Interfaces

- Fixed: Using IPv4 VIPs on PPPoE interfaces with `if_pppoe` causes looping #16235
- Fixed: PPPoE IPv6 gateways are unavailable without Router Advertisements when using `if_pppoe` #16265

OpenVPN

- Fixed: Kernel panic in OpenVPN `if_ovpn` due to use-after-free of mbuf #16319

Operating System

- Fixed: Kernel Panic in `if_qlnxe` kernel module when loaded via `kldload` #16101
- Changed: Reduce writes to disk when using ZFS #16210
- Added: Allow custom ZFS pool names #16212
- Fixed: Panic in PF if the declared packet length is longer than the actual packet #16318
- Fixed: Multiple removable filesystem panics #16320
- Added: Support 2.5G SGMII (SFP GPON ONT) in `bxe` driver (QLogic NetXtreme II BCM57810) #16321

Rules / NAT

- Fixed: Firewall rules with an interface address for the NAT64 source always use the interface itself #16250
- Fixed: Firewall rules are not performing source tracking when Sticky Connections option is enabled #16282

System Logs

- Fixed: `syslogd` does not respect the configured IP address binding #16285

UPnP IGD & PCP

- Fixed: UPnP IGD & PCP status page does not show "disabled" message after the service has been enabled then disabled #16274

3.3 25.07.1 New Features and Changes

This is a maintenance release of pfSense® Plus software which fixes issues affecting certain hardware configurations.

Warning: Users who have not yet upgraded to pfSense Plus software version 25.07 should review the *25.07 New Features and Changes* document before upgrading to this release.

Tip: Review the *Upgrade Guide* before performing any upgrade of pfSense software.

3.3.1 pfSense Plus

Changes in this version of pfSense Plus software.

Operating System

- Fixed: Go-based software crashes on hardware with 5-level paging (LA57) #16369
- Fixed: EFI loader fails to boot on some devices #16381

3.4 25.07 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

Tip: Review the *Upgrade Guide* before performing any upgrade of pfSense software.

3.4.1 General

- Older devices with ISA-based serial console ports may not fully detect their console due to changes in how FreeBSD probes serial ports. pfSense plus software attempts to detect known affected models of hardware from Netgate. Other devices may require manual intervention.

  See *ISA Serial Console not Fully Functional* for details and a workaround.

- This version of pfSense Plus software includes a new kernel-based PPPoE backend, `if_pppoe`. This will replace the current MPD-based implementation. This new backend is more efficient and enables much faster speeds over PPPoE interfaces.

  This new PPPoE backend is not active by default in this version, but can be enabled with *the global option* under **System > Advanced** on the **Networking** tab. This backend will be enabled by default on future versions of pfSense Plus software.

  The `if_pppoe` backend does not support all advanced features of the MPD implementation. For example, it does not support MLPPP.

- This release includes support for DHCPv6 Prefix Delegation in the Kea DHCP daemon.

  Warning: Prefix Delegation settings in Kea use a different format than the ISC DHCPv6 daemon, so Kea **cannot** use existing settings for Prefix Delegation. Settings for Prefix Delegation must be re-created manually when switching from ISC DHCPv6 to Kea DHCPv6. For details, see *DHCPv6 Prefix Delegation*.

- Users of the Gandi Dynamic DNS service **must** change their current API token to a Personal Access Token (PAT) as Gandi now requires this authentication method for Dynamic DNS updates. For uninterrupted Dynamic DNS service, create a new PAT and save that PAT value in Gandi Dynamic DNS entries before upgrading to this release.

3.4.2 Security / Errata

This release fixes several security issues in pfSense Plus software, including:

- pfSense-SA-25_01.webgui Multiple problems in Dashboard widget key handling which could lead to XSS, denial of service, or configuration corruption.
- pfSense-SA-25_02.webgui OpenVPN management interface command injection from OpenVPN status and Dashboard widget.
- pfSense-SA-25_03.webgui Potential XSS in the AutoConfigBackup backup list.
- pfSense-SA-25_04.webgui Potential disclosure of AutoConfigBackup Device Key if SSH service is enabled and exposed to untrusted networks.
- pfSense-SA-25_05.webgui Stored XSS in Firewall Schedules.
- pfSense-SA-25_06.webgui Stored XSS in IPsec tunnel Phase 1 list.
- pfSense-SA-25_07.webgui Stored XSS in Wake on LAN pages and Dashboard widget.
- pfSense-SA-25_08.webgui Stored XSS in OpenVPN Dashboard widget.

Fixes for these security issues are available via the *Recommended System Patches* function of the *System Patches Package* for users running pfSense Plus software version 24.11.

3.4.3 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Added: System Aliases for various reserved networks #15776
- Changed: Exclude the WireGuard and Tailscale interface group system aliases from rules #15848

Auto Configuration Backup

- Fixed: Long configuration revision reasons can cause AutoConfigBackup upload to fail #12249
- Fixed: Potential XSS in AutoConfigBackup backup list on `services_acb.php` #15927
- Fixed: AutoConfigBackup remote revision timestamps may not be unique due to batch uploads #16011
- Fixed: "Reset" button on AutoConfigBackup Restore tab does not submit the form #16012
- Changed: AutoConfigBackup code cleanup and GUI refresh #16013
- Added: Download function for AutoConfigBackup entries #16014
- Added: Method to change the AutoConfigBackup device key #16015
- Changed: Change AutoConfigBackup default key generation format #16016
- Fixed: AutoConfigBackup entries show incorrect timestamps #16209

Backup / Restore

- Fixed: Reinstall Packages button reports another instance of `pfSense-upgrade` is running #15494
- Fixed: Backup configuration cache is not cleaned automatically #15994

Captive Portal

- Fixed: PHP error in Captive Portal with undefined zone interface list #15907
- Fixed: Captive Portal does not function with MAC filtering disabled #15926
- Fixed: Captive Portal service management via `pfSsh.php svc` fails when the zone name contains uppercase letters #16030
- Fixed: Creating a Captive Portal zone with uppercase letters overwrites existing zones of the same name #16032

Certificates

- Added: Certificate Authorities created in the GUI do not have the Basic Constraints extension marked critical #15818
- Changed: Additional error handling for invalid certificate configuration #15975

Configuration Backend

- Fixed: Contents of the configuration may still be stored even when the XML is not valid #15860
- Fixed: PHP error on save with very long configuration change descriptions #15911

DHCP (IPv4)

- Added: Kea DHCP Custom Configuration Support (IPv4 and IPv6) #15321
- Fixed: Kea fails to start if DHCP pool configuration contains default lease time or max lease time #15332
- Added: Kea Static ARP Support (IPv4 only) #15654
- Fixed: Kea can unintentionally attempt to spawn multiple processes and fail #16019
- Fixed: Static lease DNS records are incorrectly removed when backing lease expires #16022

DHCP (IPv6)

- Fixed: Old IPv6 addresses may continue to be used after DHCP or RA changes #12947
- Added: Kea DHCPv6 Prefix Delegation Support (IPv6 Only) #15652

DNS Forwarder

- Fixed: Unable to change DNS Forwarder domain overrides #15890

DNS Resolver

- Fixed: DNS Resolver option for Query Name Minimization cannot be disabled #15925

Dashboard

- Fixed: Clicking the picture widget image downloads the image with an invalid filename instead of showing it inline #15767
- Fixed: Dashboard `widgetkey` values are not validated on save or load, can lead to configuration corruption or other problems #15844
- Changed: Improve the system load impact from Dashboard widgets #15969
- Fixed: Potential XSS in OpenVPN Widget #16258

Diagnostics

- Fixed: Adding Wake-On-LAN entry from ARP table view can incorrectly include OEM text in MAC address field #15162
- Fixed: PHP error from invalid IPv6 address on `diagnostics_ping.php` #16005
- Fixed: The filtered states shown may include states for interfaces other than the selected interface #16043
- Fixed: Cannot kill states using the post-NAT address #16047
- Fixed: Memory leak in `pftop` #16347

Dynamic DNS

- Added: Improve Dynamic DNS client IPv6 support #11177
- Added: Per-instance options to control Dynamic DNS client Check IP Service behavior #14067
- Fixed: Dynamic DNS uses the default gateway interface instead of the specified interface #14605
- Added: Support LuaDNS provider #15089
- Changed: Update Gandi LiveDNS service with API changes #15258
- Fixed: RFC 2136 Dynamic DNS cannot update AAAA records over IPv6 #16028
- Fixed: Dynamic DNS IP address may not be updated after changing the interface of a Dynamic DNS entry #16046
- Fixed: Dynamic DNS fails to update AAAA record for Route53 and No-IP services #16249
- Changed: Update list of Dynamic DNS services that are only reachable over IPv4 for AAAA record updates #16251

Gateway Monitoring

- Fixed: The monitoring IP address for dynamic gateways may be unexpectedly routed via a different gateway #16069
- Fixed: Improve gateway status detection with routed monitoring addresses #16180

Gateways

- Changed: Clarify descriptions for gateway recovery options #15429
- Fixed: Cannot set a new name when duplicating an existing gateway group #16036

IPsec

- Fixed: Input validation for duplicate remote gateways does not work when using the duplicate P1 button #15598
- Fixed: Firewall generates invalid rules for IPsec tunnels with descriptions containing special symbols #16095
- Fixed: Potential XSS in IPsec Phase 1 #16115
- Fixed: IPsec unnecessarily prompts to apply changes after input errors #16162

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: Incorrect warning from `radvd` about `AdvRDNSSLifetime` value #12938
- Added: PREF64 support in Router Advertisements #15808
- Fixed: Routing Advertisements daemon fails to start when configured with more than 3 RDNSS entries in a prefix #15876

Interfaces

- Fixed: Config access error with null static routes #16104
- Fixed: Config access error after changing an interface from DHCP to Static #16105
- Fixed: Using IPv4 VIPs on PPPoE interfaces with `if_pppoe` causes looping #16235
- Fixed: PPPoE IPv6 gateways are unavailable without Router Advertisements when using `if_pppoe` #16265

L2TP

- Fixed: L2TP server settings are not saved correctly #15882

Logging

- Added: Enhanced firewall log action information display #15415
- Fixed: PHP error when saving System Log settings #15988

Multi-Instance Management

- Fixed: MIM GUI is unable to write IPv6 aliases #15959
- Fixed: Renaming an alias in MIM does not update firewall and NAT rules with the new alias name #15989

NTPD

- Fixed: PHP error after saving NTP settings with an interface selected #16063

OpenVPN

- Fixed: OpenVPN Status Page and Dashboard Widget use input values without validation #15856
- Fixed: Configuration upgrade from before revision 19.1 removes OpenVPN settings #15895
- Fixed: Kernel panic in OpenVPN `if_ovpn` due to use-after-free of mbuf #16319

Operating System

- Fixed: `pftop` core dump with ICMP states #15595
- Fixed: Azure: User credentials entered during new VM deployments are not applied to the system #15871
- Fixed: Values obtained from `sysctl` are sometimes unexpectedly empty, leading to PHP and other math errors #14648
- Fixed: Errors on the console when starting/stopping services #15912
- Fixed: RAM disk configuration check fails at boot #16023
- Fixed: RAM Disk cron jobs are not saved correctly #16059
- Fixed: Panic accessing `sysctl` OID `net.inet.ip.nhdispatch` with an INVARIANTS kernel #16081
- Fixed: Kernel Panic in `if_qlnxe` kernel module when loaded via `kldload` #16101
- Changed: Reduce writes to disk when using ZFS #16210
- Added: Allow custom ZFS pool names #16212
- Fixed: Panic in PF if the declared packet length is longer than the actual packet #16318
- Fixed: Multiple removable filesystem panics #16320

PHP Interpreter

- Fixed: Cookie named `id` prevents some forms from being loaded or saved properly #11268
- Changed: Limit PHP request order processing to only GET and POST #15893

PPP Interfaces

- Fixed: PPPoE WAN loses IPv4 addresses on `IPV6CP` `LayerDown` events #16103
- Added: Support `if_pppoe` backend for PPPoE WAN interfaces #16134

Package System

- Fixed: Deleting one pre-installed package may delete other pre-installed packages #15643
- Fixed: The package `post-install` script does not run with a system upgrade on ZFS #16057
- Changed: `pkg` no longer supports setting `ALTABI` manually at run-time #16060

Rules / NAT

- Fixed: Separators for Ethernet rules span past the actions column #16079
- Added: NAT64 support #2358
- Fixed: SCTP states not purged causing subsequent SCTP INIT to be blocked #15924
- Fixed: Incorrect rule may be opened for editing after rule order has changed #15935
- Fixed: Tracking information for firewall rules is not shown when editing the rule #15936
- Fixed: Warning message in logs when changing firewall rules after setting Require Firewall Interface #15961
- Fixed: Deleting or adding a firewall rule may result in an unexpected rule order #16076
- Fixed: Potential XSS in Firewall Schedules #16114
- Fixed: Input validation prevents creating port forwards for the same port using a different address family #16130
- Fixed: Firewall rules using interface subnet aliases may prevent filter rules from loading after upgrades #16182
- Fixed: Firewall rules with an interface address for the NAT64 source always use the interface itself #16250
- Fixed: Firewall rules are not performing source tracking when Sticky Connections option is enabled #16282

System Logs

- Added: Separate IDS/IPS and link-local firewall log entries from default block logging #16092
- Fixed: `syslogd` does not respect the configured IP address binding #16285

Traffic Shaper (ALTQ)

- Fixed: Error when viewing ALTQ Traffic Shaper queue status #15885

Traffic Shaper (Limiters)

- Fixed: Limiters saved while MIM is enabled disappear after reboot #16051
- Fixed: Input validation error when applying limiter changes #13158
- Fixed: Setting a limiter queue length greater than 100 prevents the limiter from loading #13662
- Fixed: Cannot add limiters named `new` #13687
- Fixed: PHP error when a queue is added with the same name as a limiter #15914

UPnP IGD & PCP

- Changed: Update UPnP IGD & PCP GUI text #15864
- Changed: Make the UPnP IGD & PCP STUN port optional #15865
- Fixed: UPnP IGD & PCP status page does not show "disabled" message after the service has been enabled then disabled #16274

Upgrade

- Fixed: Upgrade available LED not set before branch is selected. #15880
- Changed: Link to release information on the system update page #15953
- Fixed: Boot loader is not upgraded on UFS installs #16064

User Manager / Privileges

- Fixed: Users with Deny Config Write privilege can trigger some VLAN interface operations #15282
- Fixed: Users with Deny Config Write privilege can trigger some QinQ interface operations #15318
- Fixed: PHP error when a user is denied access to the dashboard #15873
- Fixed: Users with Deny Config Write privilege can trigger logging operations #15874
- Fixed: Users with Deny Config Write privilege can change their own password #15908

Wake on LAN

- Fixed: Potential XSS in Wake on LAN page and widget #16116

Web Interface

- Added: Custom message text for the login screen #9293
- Changed: Update nginx HTTP2 syntax #15863
- Fixed: Incorrect color in button text within disabled rows #15977

3.5 2.8.0 New Features and Changes

This pfSense® CE software release includes new features and bug fixes.

3.5.1 Upgrade Notes

> **Warning:** Due to major changes in PHP and base OS versions, there is a higher than usual chance that packages will interfere with the upgrade process.
>
> To give an upgrade the best possible chance of going smoothly, uninstall **all** packages before starting the upgrade.

Before upgrading, pay particular attention to the *Pre-Upgrade Tasks* section of the *Upgrade Guide*. The most crucial points are noted in this section, but the best practice is to follow all of the precautions noted in the Upgrade Guide.

Boot Loader

This version requires an updated boot loader, which is automatically handled by the upgrade process for nearly all cases. However, there may be some edge cases where the automatic update does not update the loader currently used by the device. For example, if there are multiple unmirrored disks and the BIOS/EFI Firmware is not booting from the disk containing the updated loader, but an older unrelated installation on a separate disk. One particular case where this can happen is when there is a previous installation to MMC which has been followed by an installation to an add-on SSD without clearing the MMC contents.

In these cases the best practice is to wipe the unused disk so it cannot interfere. See *Troubleshooting Multiple Disks* for details.

Legacy Serial Console

After upgrading, older devices with ISA-based serial console ports may not fully detect their console due to changes in how FreeBSD probes serial ports. Devices may require manual intervention.

Devices affected by this are primarily older non-EFI hardware running pfSense CE software. This includes devices such as the APU (1), RCC-DFF 2220, and RCC-VE models such as the 2440, 4860, and 8860.

See also:

See *ISA Serial Console not Fully Functional* for details and a workaround.

Note: Netgate devices running pfSense Plus software and which are known to be affected automatically handle this configuration and do not require manual changes.

Low Memory Hardware

Hardware with **1 GiB or less** available memory may have issues upgrading depending on which features, services, or packages are running.

Tip: For devices running ZFS, see *ZFS Tuning* for information on reducing ZFS memory usage.

For the best chance of success in these cases, temporarily disable any non-critical services before starting the upgrade. Rebooting before attempting the upgrade can also be beneficial.

3.5.2 General

- PHP has been upgraded from 8.2.x to 8.3.x
- The base operating system has been upgraded to FreeBSD 15-CURRENT
- This version of pfSense CE software includes a new kernel-based PPPoE backend, `if_pppoe`. This will replace the current MPD-based implementation. This new backend is more efficient and enables much faster speeds over PPPoE interfaces.

  This new PPPoE backend is not active by default in this version, but can be enabled with *the global option* under **System > Advanced** on the **Networking** tab. This backend will be enabled by default on future versions of pfSense Plus software.

  The `if_pppoe` backend does not support all advanced features of the MPD implementation. For example, it does not support MLPPP.
- The default State Policy has been changed from **Floating** to **Interface Bound** for increased security. However, **Interface Bound** states may have issues in certain cases with IPsec VTI, Multi-WAN policy routing (`route-to`), `reply-to`, as well as with High Availability state synchronization (pfsync) on non-identical hardware. Workarounds are in place to fall back to **Floating** states in certain cases, such as IPsec/VTI.

  The default policy can be toggled back to **Floating** using the **State Policy** option under **System > Advanced** on the **Firewall & NAT** tab.

  There is also an option to override this behavior on a per-rule basis in the advanced options when editing a firewall rule.

  See also:

  - *State Policy History*
  - *Firewall State Policy*
  - *State Policy*
- This release includes support for enhanced gateway recovery "fail back" by optionally clearing states from lower tier gateways when a more preferred gateway recovers.

  See also:

  - *State Killing on Gateway Recovery*
  - *Gateway Group Options*
- This release includes support for High Availability in the Kea DHCP daemon.

  This implementation has several advantages over the older ISC DHCP implementation, including:

  - Supports HA for DHCPv4 **and** DHCPv6.

  - Simplified HA setup, all in one place on each node for each type.
  - Works in hot standby mode, which is more reliable.
  - Can synchronize lease data over the SYNC interface for security and ease of use, and can optionally encrypt the sync data for added protection.

  See also:

  For in-depth information on this feature, see https://www.netgate.com/blog/improvements-to-kea-dhcp

- This release includes support for DNS Registration of DHCP client hostnames from the Kea DHCP daemon to the Unbound DNS Resolver
  - DNS records are updated dynamically on-the-fly, they do not require a resolver restart and are not disruptive.
  - Supports DNS Registration for DHCPv4 **and** DHCPv6
  - DNS Registration can be configured on a per-interface or global manner, with the ability to enable or disable specific interfaces as needed.
  - DNS records are not limited to the system domain name. DNS Registration honors the domain name on the DHCP settings for each interface **and** on static mappings.
  - DNS records are accurate/updated on both high availability peers
  - Static mappings can be registered when Kea starts (similar to ISC) or when a static mapping client obtains a lease.
- This release includes support for DHCPv6 Prefix Delegation in the Kea DHCP daemon.

  Warning: Prefix Delegation settings in Kea use a different format than the ISC DHCPv6 daemon, so Kea **cannot** use existing settings for Prefix Delegation. Settings for Prefix Delegation must be re-created manually when switching from ISC DHCPv6 to Kea DHCPv6. For details, see *DHCPv6 Prefix Delegation*.

See also:

For additional details on implementation of Kea DHCP features see https://redmine.pfsense.org/issues/15650

- Users of the Gandi Dynamic DNS service **must** change their current API token to a Personal Access Token (PAT) as Gandi now requires this authentication method for Dynamic DNS updates. For uninterrupted Dynamic DNS service, create a new PAT and save that PAT value in Gandi Dynamic DNS entries before upgrading to this release.

3.5.3 Security / Errata

This release fixes several security issues in pfSense CE software, including:

- pfSense-SA-25_01.webgui Multiple problems in Dashboard widget key handling which could lead to XSS, denial of service, or configuration corruption.
- pfSense-SA-25_02.webgui OpenVPN management interface command injection from OpenVPN status and Dashboard widget.
- pfSense-SA-25_03.webgui Potential XSS in the AutoConfigBackup backup list.
- pfSense-SA-25_04.webgui Potential disclosure of AutoConfigBackup Device Key if SSH service is enabled and exposed to untrusted networks.
- pfSense-SA-25_05.webgui Stored XSS in Firewall Schedules.
- pfSense-SA-25_06.webgui Stored XSS in IPsec tunnel Phase 1 list.

- pfSense-SA-25_07.webgui Stored XSS in Wake on LAN pages and Dashboard widget.

Fixes for these security issues are available via the *Recommended System Patches* function of the *System Patches Package* for users running pfSense CE software version 2.7.2.

3.5.4 pfSense CE

Changes in this version of pfSense CE software.

Aliases / Tables

- Added: Allow user-defined rules to utilize built-in system aliases #1979
- Fixed: Interface subnet aliases do not contain IPv6 VIPs #15096
- Added: System Aliases for various reserved networks #15776
- Changed: Exclude the WireGuard and Tailscale interface group system aliases from rules #15848

Authentication

- Fixed: PHP errors in LDAP server prevent it from falling back to Local Database #15122
- Fixed: GUI logout messages do not use the `auth` log facility #15719

Auto Configuration Backup

- Fixed: Long configuration revision reasons can cause AutoConfigBackup upload to fail #12249
- Fixed: `services_acb_settings.php` does not fully validate value of `frequency`, uses value without encoding #15224
- Fixed: Special characters in the ACB configuration change description can cause PHP errors #15711
- Fixed: AutoConfigBackup tries to upload backups before the system has finished booting #15718
- Fixed: Potential XSS in AutoConfigBackup backup list on `services_acb.php` #15927
- Fixed: AutoConfigBackup scheduled backups always upload even when the configuration has not changed #16010
- Fixed: AutoConfigBackup remote revision timestamps may not be unique due to batch uploads #16011
- Fixed: "Reset" button on AutoConfigBackup Restore tab does not submit the form #16012
- Changed: AutoConfigBackup code cleanup and GUI refresh #16013
- Added: Download function for AutoConfigBackup entries #16014
- Added: Method to change the AutoConfigBackup device key #16015
- Changed: Change AutoConfigBackup default key generation format #16016

Backup / Restore

- Added: Support for CD/DVD drives in the External Configuration Locator (ECL) #14728
- Fixed: DHCP leases may not be restored from older configuration backups #15076
- Fixed: PHP error when generating a notification after detecting a malformed configuration #15157
- Fixed: Skip Packages option for Configuration Backups fails with large configurations #15624

CARP

- Fixed: HA node with CARP VIP in backup state is unable to ping the active node using that CARP VIP address #14026

Captive Portal

- Fixed: Disconnecting a user from Captive Portal may allow previously established connections to continue #13226
- Added: Support using a mask to block MAC addresses in Captive Portal #15257
- Fixed: Old auto-added MAC addresses are not pruned for non-concurrent Captive Portal sessions #15299
- Fixed: Captive Portal logo fails to load after authenticated redirect #15404
- Fixed: Captive Portal zones can fail to start due to ID conflict #15772
- Fixed: PHP error in Captive Portal with undefined zone interface list #15907
- Fixed: Captive Portal service management via `pfSsh.php svc` fails when the zone name contains uppercase letters #16030
- Fixed: Creating a Captive Portal zone with uppercase letters overwrites existing zones of the same name #16032

Certificates

- Fixed: Certificate Manager GUI inconsistency in Revocation tab titles #15454
- Added: Certificate Authorities created in the GUI do not have the Basic Constraints extension marked critical #15818
- Changed: Additional error handling for invalid certificate configuration #15975

Configuration Backend

- Fixed: System proxy credentials with certain characters may fail to authenticate #15565

Console Menu

- Changed: Dynamically adjust the interface name maximum width in the login banner #13268
- Fixed: Declining to reset the admin account via the console menu still prompts to change the password #15751

DHCP (IPv4)

- Added: Settings tab for global Kea DHCP server options #5080
- Added: Better handling of duplicate IP addresses in static DHCP assignments #13256
- Changed: Reduce log spam when deleting a static DHCP entry #13263
- Added: Explicitly enable/disable DHCP Dynamic DNS updates in each scope #13894
- Fixed: Kea fails to restart due to race between process termination and startup #14977
- Fixed: Kea does not allow FQDNs for NTP servers but input validation does not prevent them from being added #14991
- Fixed: Kea DHCP PHP error from WINS server value #14996
- Fixed: Kea DHCP sends wrong bootloader file for UEFI #15032
- Fixed: Kea will not start with identical MAC address filters on multiple interfaces #15130
- Added: Kea DHCP Custom Configuration Support (IPv4 and IPv6) #15321
- Fixed: Changes in Kea DHCP interface pools may invalidate lease database content #15328
- Fixed: Kea fails to start if DHCP pool configuration contains default lease time or max lease time #15332
- Added: Kea High Availability Support (IPv4 and IPv6) #15575
- Added: Kea DNS Resolver (Unbound) Integration (IPv4 and IPv6) #15651
- Added: Kea Static ARP Support (IPv4 only) #15654
- Fixed: IPv4 DHCP client responses may be routed unexpectedly out unrelated WANs #15702
- Added: Kea DHCP lease database RAM disk support (IPv4 and IPv6) #15828
- Fixed: Kea can unintentionally attempt to spawn multiple processes and fail #16019

DHCP (IPv6)

- Fixed: Old IPv6 addresses may continue to be used after DHCP or RA changes #12947
- Fixed: Shortcut bar on DHCPv6 leases (`status_dhcpv6_leases.php`) navigates to DHCPv4 destinations, not DHCPv6 #15117
- Fixed: DHCPv6 settings page “DDNS Reverse” check box not showing current state #15118
- Added: Kea DHCPv6 Prefix Delegation Support (IPv6 Only) #15652

DNS Forwarder

- Added: Option to allow the DNS Forwarder to ignore system DNS servers #14165
- Fixed: DNS Forwarder ignores "Use remote DNS Servers, ignore local DNS" setting #15434
- Changed: Update dnsmasq to version 2.90 #15465

DNS Resolver

- Fixed: DNS Resolver host overrides ignore all aliases if first entry has a domain set but no hostname #14942
- Fixed: Applying interface changes may not update default ACLs for the DNS Resolver #15071
- Fixed: Potential local file include vulnerability via DNS Resolver Python Module Script include mechanism #15135
- Fixed: Local DNS resolution behavior does not add an IPv6 nameserver #15139
- Changed: Update Unbound to 1.22.0 #15483
- Fixed: Automatic EDNS value may be lower than expected #15704
- Fixed: Unbound configuration file contains Localhost address in forwarding mode with TLS enabled #15722
- Fixed: `unbound-checkconf` fails with python mode enabled #15723

Dashboard

- Fixed: Firewall Logs Dashboard Widget is slow and may fail to update #12673
- Added: Improve Thermal Sensors Dashboard widget readability #13520
- Fixed: Traffic Graph widget displays bandwidth usage values which are half the actual usage amount #14933
- Fixed: Firewall Logs Dashboard widget update interval does not behave as expected #15373
- Added: Show current boot method in System Information Dashboard widget #15422
- Fixed: Incorrect icon on collapsed dashboard widgets #15439
- Fixed: Dashboard widgets refresh at unintended intervals #15725
- Changed: Improve Thermal Sensors Dashboard widget refresh code #15728
- Fixed: Session cookie warnings #15729
- Fixed: Clicking the picture widget image downloads the image with an invalid filename instead of showing it inline #15767
- Fixed: Dashboard `widgetkey` values are not validated on save or load, can lead to configuration corruption or other problems #15844
- Changed: Improve the system load impact from Dashboard widgets #15969

Diagnostics

- Added: Add Kea information to `status.php` #14953
- Fixed: Adding Wake-On-LAN entry from ARP table view can incorrectly include OEM text in MAC address field #15162
- Fixed: `crash_reporter.php` displays PHP Error log without encoding #15264
- Added: Add EFI boot information to `status.php` #15297
- Added: Add `loader.conf.lua` contents to `status.php` #15298
- Fixed: Errors in `status.php` IPsec sections when IPsec is not configured #15310
- Fixed: Sanitize RFC 2136 Dynamic DNS update keys in `status.php` output #15490
- Fixed: File browser on `diag_edit.php` does not encode directory names before display #15525
- Fixed: State table entries printed on `diag_dump_states.php` may contain an unexpected interface #15657
- Fixed: PHP error from invalid IPv6 address on `diagnostics_ping.php` #16005
- Fixed: Cannot kill states using the post-NAT address #16047

Dynamic DNS

- Added: Enable `@` support for Azure in Dynamic DNS #10000
- Added: Improve Dynamic DNS client IPv6 support #11177
- Added: Per-instance options to control Dynamic DNS client Check IP Service behavior #14067
- Added: Enable `@` support for name.com in Dynamic DNS #14289
- Fixed: Dynamic DNS uses the default gateway interface instead of the specified interface #14605
- Added: Support LuaDNS provider #15089
- Changed: Update Gandi LiveDNS service with API changes #15258
- Changed: Update Dynamic DNS API URL for porkbun.com #15779
- Fixed: Dynamic DNS attempts to resolve entries with disabled interfaces #15802
- Fixed: RFC 2136 Dynamic DNS cannot update AAAA records over IPv6 #16028
- Fixed: Dynamic DNS IP address may not be updated after changing the interface of a Dynamic DNS entry #16046

FreeBSD

- Fixed: Kernel panic in HA nodes when under high load #15413

Gateway Monitoring

- Fixed: Gateway behavior differs when the gateway does not exist in the configuration #12920
- Fixed: Gateway monitoring includes disabled gateways #15635
- Fixed: The monitoring IP address for dynamic gateways may be unexpectedly routed via a different gateway #16069
- Fixed: Improve gateway status detection with routed monitoring addresses #16180

Gateways

- Fixed: Killing states on downed gateways breaks when `Skip rules when gateway is down` is enabled #15223
- Fixed: Killing states on downed gateways breaks for static interface configurations #15225
- Fixed: Removing a gateway group used as the default gateway results in no default route #15248
- Changed: Clarify descriptions for gateway recovery options #15429
- Fixed: Saving an IPv6 gateway overrides the IPv4 gateway #15589
- Fixed: No default route after boot #15791

Hardware / Drivers

- Fixed: Newer variant models within the PC Engines APU2 platform are not recognized, causing garbled early serial console output #13498
- Added: Recognize QAT 4xxx devices in System Information Widget #15233

High Availability

- Fixed: Removing a route from the High Availability primary node does not remove the entry from the routing table on the secondary node #15795

IGMP Proxy

- Fixed: IGMP proxy works intermittently #15043
- Fixed: Kernel Panic when IGMPProxy gets CIDR Removed #15831

IPsec

- Fixed: MSS clamping on VPN traffic does not work on IPsec IPv6 mobile VPNs #14312
- Fixed: Large number of IPsec tunnels causes long filter reload times #14893
- Fixed: IPsec VTI is not created correctly when using a Phase 2 remote type of `Network` #15124
- Fixed: Cannot configure dual stack IPsec tunnel to accept connections from any remote address on both address families #15147
- Fixed: Removing an IPsec Phase 1 entry can either remove the wrong Phase 2 entries or leave orphaned Phase 2 entries in the configuration #15171

- Fixed: Change Mobile IPsec RADIUS accounting to use `accounting_requires_vip` so accounting will not activate for non-mobile VPNs #15176
- Added: Show interface subnet details in a tooltip on the IPsec Phase 2 list #15245
- Fixed: Reordering IPsec Phase 2 entries may result in a malformed configuration #15384
- Fixed: Input validation for duplicate remote gateways does not work when using the duplicate P1 button #15598
- Fixed: Mobile IPsec does not automatically switch to failover gateway #15685
- Fixed: Mobile IPsec sends incorrect DNS attribute IDs #15755
- Fixed: Firewall generates invalid rules for IPsec tunnels with descriptions containing special symbols #16095
- Fixed: Potential XSS in IPsec Phase 1 #16115
- Fixed: IPsec allows deleting P1/P2 entries with an assigned VTI #16158
- Fixed: IPsec unnecessarily prompts to apply changes after input errors #16162

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: Non Link-Local IPv6 CARP address does not get advertised to endpoints with RADVD #12581
- Fixed: Incorrect warning from `radvd` about `AdvRDNSSLifetime` value #12938
- Fixed: `radvd` service shows as stopped in services list when it should be disabled and hidden from that list #14936
- Fixed: Cannot disable Router Advertisements when the interface IPv6 configuration is set to `None` #14967
- Fixed: Router Advertisement daemon does not prioritize IPv6 GUA over ULA #15057
- Added: PREF64 support in Router Advertisements #15808
- Fixed: Routing Advertisements daemon fails to start when configured with more than 3 RDNSS entries in a prefix #15876

Installer

- Fixed: Clean installation using Auto (ZFS) + MBR (BIOS) does not boot #14930
- Fixed: Installing to ZFS mirror does not format or populate EFI partition on additional disks #15083

Interfaces

- Fixed: Adding MSS and MTU values on a LAGG VLAN interface breaks connectivity #14083
- Fixed: Sending IPv6 traffic on a disabled interface can trigger a kernel panic #14431
- Fixed: PHP error in `interfaces_qinq_edit.php` when creating a QinQ interface #15181
- Fixed: PHP error when applying interface settings if the `/tmp/.interfaces.apply` file is present but empty #15423
- Added: Use natural sorting when sorting interfaces #15437
- Fixed: OpenVPN QinQ interface creation fails #15692
- Fixed: Interface group members are not validated on load/save on `interfaces_groups_edit.php`, and are printed without encoding on `interfaces_groups.php` #15778

- Fixed: Config access error with null static routes #16104
- Fixed: Config access error after changing an interface from DHCP to Static #16105

LAGG Interfaces

- Fixed: Reconfiguring a parent LAGG interface breaks its VLANs #9453

Logging

- Fixed: Restarting the logging daemon during rotation also restarts `sshguard`, leading to frequent log messages #12747
- Changed: Remove `Time` column from OS Boot logs #15106
- Added: Enhanced firewall log action information display #15415
- Fixed: PHP error when saving System Log settings #15988

Multi-WAN

- Added: Ability to selectively kill states on gateway recovery #855

NTPD

- Added: NTP authentication support #8794

OpenVPN

- Added: More GUI options for OpenVPN Client-Specific Overrides #12522
- Added: OpenVPN NBDD server options #13085
- Fixed: OpenVPN WINS options may be visible even when NetBIOS is disabled #13087
- Fixed: Some OpenVPN NetBIOS settings are kept even when NetBIOS is disabled #13089
- Fixed: OpenVPN NetBIOS Node Type and Scope ID options are not pushed to clients #13090
- Fixed: `openvpn.auth-user.php` gets stuck at 100% CPU usage when RADIUS authentication times out #14386
- Fixed: OpenVPN forms invalid `route` statements for empty local networks #14919
- Fixed: PHP error with OpenVPN server certificate verification if the certificate has multiple `CN` attributes #15133
- Fixed: OpenVPN Wizard fails when a VIP is used #15148
- Changed: Remove deprecated OpenVPN hardware crypto engine option #15188
- Fixed: OpenVPN Status Page and Dashboard Widget use input values without validation #15856

Operating System

- Fixed: `/etc/rc.local` script content is executed at login instead of during boot sequence #10980
- Fixed: Values obtained from `sysctl` are sometimes unexpectedly empty, leading to PHP and other math errors #14648
- Fixed: Static ARP assignments lose `permanent` flag in ARP table #14970
- Fixed: Permissions on tmpfs RAM disk for `/var` are too lenient #15054
- Fixed: `pfctl` is unable to retrieve state creator list in certain circumstances #15108
- Fixed: `loader.conf` may be missing `loader_conf_files` so `loader.conf.lua` may not be parsed #15288
- Fixed: Proxy variables in `crontab` contents are improperly formatted #15502
- Fixed: `resizewin` occasionally gets fed a spurious line feed over certain serial console+client combinations #15777
- Fixed: Panic accessing `sysctl` OID `net.inet.ip.nhdispatch` with an INVARIANTS kernel #16081

PHP Interpreter

- Fixed: Cookie named `id` prevents some forms from being loaded or saved properly #11268
- Fixed: Extensions directory is not set in `rc.php_ini_setup` #14488
- Changed: Update PHP to 8.3.x #15053
- Fixed: `check_dnsavailable()` failing even when DNS is available #15127
- Fixed: PHP error display formatting issues #15263
- Fixed: Memory leak in pfSense module function `pfSense_get_ifaddrs()` #15471

PPP Interfaces

- Fixed: PPPoE WAN loses IPv4 addresses on `IPV6CP` `LayerDown` events #16103
- Added: Support `if_pppoe` backend for PPPoE WAN interfaces #16134

Package System

- Added: Allow overriding text scrolling during package install/uninstall #15022
- Fixed: Extra space in `pkg` configuration file `FreeBSD.conf` #15069
- Fixed: Updates fail against an authenticated upstream proxy #15094
- Fixed: Package navigation menus can be duplicated when reinstalling the package #15700
- Fixed: The package `post-install` script does not run with a system upgrade on ZFS #16057
- Changed: `pkg` no longer supports setting `ALTABI` manually at run-time #16060

Packet Capture

- Fixed: Unable to perform Packet Captures on a tailscale interface in GUI with default settings #15145
- Added: Allow filtering packet captures by system-defined protocols #15609

Routing

- Fixed: ICMPv6 Path MTU Discovery breaks with NPT #14290
- Fixed: IPsec VTI static routes may not be added after the system boots #15449
- Fixed: Routes with IPv6 Address as Next Hop for IPv4 Destination Causes Kernel Panic #15601

Rules / NAT

- Added: NAT64 support #2358
- Added: Kill states using the pre-NAT address #11556
- Changed: Add global option to set default PF State Policy (if-bound vs floating) #15173
- Added: Add per-rule option to set PF State Policy (if-bound vs floating) #15183
- Fixed: Outbound NAT rules using an alias without a matching address family create unexpected PF rules #15197
- Fixed: Advanced rule options tooltip does not show negated Tag option #15214
- Added: Show details of system aliases in tooltip on firewall and NAT rule lists #15234
- Fixed: Egress states remain when killing states for scheduled rules #15252
- Fixed: Interface-bound state policy does not handle IPsec VTI traffic as expected when filtering on `enc0` interface #15430
- Fixed: Per-rule byte counter values lost across a filter reload #15516
- Fixed: Separator positions are incorrect when copying interface group rules #15537
- Added: GUI options to change default SCTP state timeouts #15661
- Fixed: Setting the Port Forward interface to an interface group selects an invalid destination #15671
- Fixed: SCTP states not purged causing subsequent SCTP INIT to be blocked #15924
- Fixed: Incorrect rule may be opened for editing after rule order has changed #15935
- Fixed: Deleting or adding a firewall rule may result in an unexpected rule order #16076
- Fixed: Potential XSS in Firewall Schedules #16114
- Fixed: Input validation prevents creating port forwards for the same port using a different address family #16130
- Fixed: Firewall rules using interface subnet aliases may prevent filter rules from loading after upgrades #16182

S.M.A.R.T.

- Changed: Query for SMART data only on root disk devices #15586

SNMP

- Fixed: File descriptor leak in `bsnmpd` #15481

Services

- Fixed: NTP option "DNS Resolution" has no effect when using NTP pool hostnames #15552

Setup Wizard

- Changed: Error handling in the Setup Wizard is very user-unfriendly #15302

System Logs

- Added: Separate IDS/IPS and link-local firewall log entries from default block logging #16092

Traffic Shaper (Limiters)

- Fixed: Input validation error when applying limiter changes #13158
- Fixed: Setting a limiter queue length greater than 100 prevents the limiter from loading #13662
- Fixed: Cannot add limiters named `new` #13687
- Fixed: Packets are passed through dummynet twice when using `route-to` leading to half the expected bandwidth #14854
- Fixed: Fragmented packets delayed by limiters are lost #15156
- Fixed: Reply traffic on a secondary WAN may be dropped when passed through dummynet #15363
- Fixed: PHP error when a queue is added with the same name as a limiter #15914

UPnP IGD & PCP

- Fixed: Port forward rules created by `miniupnpd` do not expire #15470
- Changed: Update UPnP IGD & PCP GUI text #15864
- Changed: Make the UPnP IGD & PCP STUN port optional #15865

Upgrade

- Fixed: Upgrading an EFI system installed to ZFS mirror does not upgrade EFI loader on additional disks #15084
- Changed: Link to release information on the system update page #15953
- Fixed: Boot loader is not upgraded on UFS installs #16064

User Manager / Privileges

- Fixed: Users with Deny Config Write privilege can trigger some VLAN interface operations #15282
- Fixed: Users with Deny Config Write privilege can trigger some QinQ interface operations #15318
- Fixed: CLI password check exits with a write access error when checking is a read-only operation #15442
- Fixed: PHP error when a user is denied access to the dashboard #15873
- Fixed: Users with Deny Config Write privilege can trigger logging operations #15874
- Fixed: Users with Deny Config Write privilege can change their own password #15908

Virtual IP Addresses

- Fixed: `choparp` service is not stopped after deleting Proxy ARP type Virtual IP addresses #14929
- Fixed: Network and broadcast address input validation is incorrectly applied to IPv6 VIPs #15361

Wake on LAN

- Fixed: Potential XSS in Wake on LAN page and widget #16116

Web Interface

- Added: Overflow scrolling for top navigation drop-down menus in Fixed mode #7943
- Added: Custom message text for the login screen #9293
- Fixed: Some messages presented to users contain relative links to pages which may be invalid when triggered from certain packages #13413
- Changed: Update vendor files #13537
- Fixed: `status_interfaces.php` is missing several values for SFP modules #15112
- Changed: Remove `jquery-treegrid` unit testing files #15265
- Added: 50x and 404 error handling to GUI web server configuration #15322
- Changed: Remove deprecated HTTP/1.0 Pragma header #15781
- Changed: Use minified nvd3 vendor files #15782
- Changed: Update nginx HTTP2 syntax #15863
- Fixed: Incorrect color in button text within disabled rows #15977

XMLRPC

- Fixed: Secondary node attempts to delete the `admins` group when synchronizing accounts via XMLRPC #15067
- Fixed: Changes to the `admins` user group are not synced to the secondary node #15898

3.6 24.11 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

3.6.1 General

- This release includes support for High Availability in the Kea DHCP daemon.

  This implementation has several advantages over the older ISC DHCP implementation, including:

  - Supports HA for DHCPv4 **and** DHCPv6.
  - Simplified HA setup, all in one place on each node for each type.
  - Works in hot standby mode, which is more reliable.
  - Can synchronize lease data over the SYNC interface for security and ease of use, and can optionally encrypt the sync data for added protection.

  See also:

  For in-depth information on this feature, see https://www.netgate.com/blog/improvements-to-kea-dhcp

- This release includes support for DNS Registration of DHCP client hostnames from the Kea DHCP daemon to the Unbound DNS Resolver
  - DNS records are updated dynamically on-the-fly, they do not require a resolver restart and are not disruptive.
  - Supports DNS Registration for DHCPv4 **and** DHCPv6
  - DNS Registration can be configured on a per-interface or global manner, with the ability to enable or disable specific interfaces as needed.
  - DNS records are not limited to the system domain name. DNS Registration honors the domain name on the DHCP settings for each interface **and** on static mappings.
  - DNS records are accurate/updated on both high availability peers
  - Static mappings can be registered when Kea starts (similar to ISC) or when a static mapping client obtains a lease.

See also:

For additional details on implementation of Kea DHCP features see https://redmine.pfsense.org/issues/15650

3.6.2 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Added: Allow user-defined rules to utilize built-in system aliases #1979

Authentication

- Fixed: `sshguard` is not properly detecting GUI login failures #15687
- Fixed: GUI logout messages do not use the `auth` log facility #15719

Auto Configuration Backup

- Fixed: Special characters in the ACB configuration change description can cause PHP errors #15711
- Fixed: AutoConfigBackup tries to upload backups before the system has finished booting #15718

Backup / Restore

- Fixed: Factory resetting the configuration removes WireGuard #15511
- Fixed: Skip Packages option for Configuration Backups fails with large configurations #15624

CARP

- Fixed: HA node with CARP VIP in backup state is unable to ping the active node using that CARP VIP address #14026

Captive Portal

- Fixed: Captive Portal logo fails to load after authenticated redirect #15404
- Fixed: Captive Portal zones can fail to start due to ID conflict #15772

Certificates

- Fixed: CA certificates are not added to the Trust Store #15440
- Fixed: Certificate Manager GUI inconsistency in Revocation tab titles #15454

Configuration Backend

- Fixed: System proxy credentials with certain characters may fail to authenticate #15565

Console Menu

- Fixed: Declining to reset the admin account via the console menu still prompts to change the password #15751

DHCP (IPv4)

- Added: Settings tab for global Kea DHCP server options #5080
- Fixed: Kea fails to restart due to race between process termination and startup #14977
- Fixed: Kea will not start with identical MAC address filters on multiple interfaces #15130
- Fixed: Changes in Kea DHCP interface pools may invalidate lease database content #15328
- Fixed: Kea does not send configured TFTP server name #15518
- Added: Kea High Availability Support (IPv4 and IPv6) #15575
- Added: Kea DNS Resolver (Unbound) Integration (IPv4 and IPv6) #15651
- Fixed: IPv4 DHCP client responses may be routed unexpectedly out unrelated WANs #15702
- Fixed: Hostnames for ISC DHCP leases are not removed from Unbound when switching to Kea #15750
- Added: Kea DHCP lease database RAM disk support (IPv4 and IPv6) #15828

DNS Forwarder

- Fixed: DNS Forwarder ignores "Use remote DNS Servers, ignore local DNS" setting #15434
- Changed: Update dnsmasq to version 2.90 #15465

DNS Resolver

- Fixed: Reduce disruptions when changing DNS records from DHCP leases in Unbound #5413
- Changed: Update Unbound to 1.22.0 #15483
- Fixed: Automatic EDNS value may be lower than expected #15704
- Fixed: Unbound configuration file contains Localhost address in forwarding mode with TLS enabled #15722
- Fixed: `unbound-checkconf` fails with python mode enabled #15723

Dashboard

- Added: Improve Thermal Sensors Dashboard widget readability #13520
- Fixed: Traffic Graph widget displays bandwidth usage values which are half the actual usage amount #14933
- Fixed: Firewall Logs Dashboard widget update interval does not behave as expected #15373
- Added: Show current boot method in System Information Dashboard widget #15422
- Fixed: Incorrect icon on collapsed dashboard widgets #15439
- Fixed: Dashboard widgets refresh at unintended intervals #15725
- Changed: Improve Thermal Sensors Dashboard widget refresh code #15728
- Fixed: Session cookie warnings #15729

Diagnostics

- Fixed: Sanitize RFC 2136 Dynamic DNS update keys in `status.php` output #15490
- Fixed: File browser on `diag_edit.php` does not encode directory names before display #15525
- Fixed: State table entries printed on `diag_dump_states.php` may contain an unexpected interface #15657

Dynamic DNS

- Added: Enable `@` support for Azure in Dynamic DNS #10000
- Added: Enable `@` support for name.com in Dynamic DNS #14289
- Changed: Update Dynamic DNS API URL for porkbun.com #15779
- Fixed: Dynamic DNS attempts to resolve entries with disabled interfaces #15802

FreeBSD

- Fixed: Kernel panic in HA nodes when under high load #15413

Gateway Monitoring

- Fixed: Gateway monitoring includes disabled gateways #15635

Gateways

- Fixed: No default route after boot #15791

High Availability

- Fixed: Removing a route from the High Availability primary node does not remove the entry from the routing table on the secondary node #15795

IGMP Proxy

- Fixed: Kernel Panic when IGMPProxy gets CIDR Removed #15831

IPsec

- Fixed: Mobile IPsec does not automatically switch to failover gateway #15685
- Fixed: Mobile IPsec sends incorrect DNS attribute IDs #15755

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: Non Link-Local IPv6 CARP address does not get advertised to endpoints with RADVD #12581

Installer

- Fixed: Installing to ZFS mirror does not format or populate EFI partition on additional disks #15083

Interfaces

- Fixed: Adding MSS and MTU values on a LAGG VLAN interface breaks connectivity #14083
- Fixed: PHP error when applying interface settings if the `/tmp/.interfaces.apply` file is present but empty #15423
- Added: Use natural sorting when sorting interfaces #15437
- Fixed: OpenVPN QinQ interface creation fails #15692
- Fixed: Interface group members are not validated on load/save on `interfaces_groups_edit.php`, and are printed without encoding on `interfaces_groups.php` #15778

Logging

- Fixed: Restarting the logging daemon during rotation also restarts `sshguard`, leading to frequent log messages #12747

Multi-WAN

- Fixed: State Killing on Gateway Recovery fails for the default gateway group with the “Kill all” option selected #15694

NTPD

- Added: NTP authentication support #8794

OpenVPN

- Added: More GUI options for OpenVPN Client-Specific Overrides #12522
- Fixed: PHP error with OpenVPN server certificate verification if the certificate has multiple `CN` attributes #15133

Operating System

- Fixed: Kernel panic with pflow configured and active #15446
- Fixed: Proxy variables in `crontab` contents are improperly formatted #15502
- Fixed: `resizewin` occasionally gets fed a spurious line feed over certain serial console+client combinations #15777

PHP Interpreter

- Changed: Update PHP to 8.3.x #15053
- Fixed: Memory leak in pfSense module function `pfSense_get_ifaddrs()` #15471

Package System

- Fixed: Updates fail against an authenticated upstream proxy #15094
- Fixed: Package navigation menus can be duplicated when reinstalling the package #15700

Packet Capture

- Added: Allow filtering packet captures by system-defined protocols #15609

Routing

- Fixed: Interface-bound state policy does not handle IPsec VTI traffic as expected when filtering on `enc0` interface #15430
- Fixed: IPsec VTI static routes may not be added after the system boots #15449
- Fixed: Saving an IPv6 gateway overrides the IPv4 gateway #15589
- Fixed: Routes with IPv6 Address as Next Hop for IPv4 Destination Causes Kernel Panic #15601
- Fixed: Static routes using null gateways are not installed #15669

Rules / NAT

- Fixed: Per-rule byte counter values lost across a filter reload #15516
- Fixed: Separator positions are incorrect when copying interface group rules #15537
- Added: GUI options to change default SCTP state timeouts #15661
- Fixed: Setting the Port Forward interface to an interface group selects an invalid destination #15671

S.M.A.R.T.

- Changed: Query for SMART data only on root disk devices #15586

SNMP

- Fixed: File descriptor leak in `bsnmpd` #15481

Services

- Fixed: NTP option "DNS Resolution" has no effect when using NTP pool hostnames #15552

UPnP/NAT-PMP

- Fixed: Port forward rules created by `miniupnpd` do not expire #15470

Upgrade

- Fixed: Upgrading an EFI system installed to ZFS mirror does not upgrade EFI loader on additional disks #15084

User Manager / Privileges

- Fixed: CLI password check exits with a write access error when checking is a read-only operation #15442

Virtual IP Addresses

- Fixed: Network and broadcast address input validation is incorrectly applied to IPv6 VIPs #15361

Web Interface

- Changed: Remove deprecated HTTP/1.0 Pragma header #15781
- Changed: Use minified nvd3 vendor files #15782

3.7 24.03 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

Tip: Consult the *Upgrade Guide* before proceeding with any upgrade.

3.7.1 Pre-Upgrade Cautions

Before upgrading, pay particular attention to the *Pre-Upgrade Tasks* section of the *Upgrade Guide*. The most crucial points are noted in this section, but the best practice is to follow all of the precautions noted in the Upgrade Guide.

ZFS Boot Environment Space Usage

Before attempting the upgrade, check the list of current ZFS Boot Environments (**System > Boot Environments**) and clean up any older entries to ensure they do not consume space which may be needed during the upgrade. See *Check and Clean Up ZFS Boot Environments* for details.

Low Memory Hardware and AWS/Azure Instances

Hardware with **1 GiB or less** available memory may have issues upgrading depending on which features, services, or packages are running. This includes some Netgate hardware such as the Netgate 1100 when running with ZFS and/or certain services/packages. For the best chance of success in these cases, temporarily disable any non-critical services before starting the upgrade.

pfSense Plus software can no longer run on AWS ".nano" size instances as they lack sufficient RAM to upgrade properly. Attempting to upgrade a ".nano" instance to pfSense Plus software version 24.03 will fail before the upgrade is performed. Migrate the instance to a ".micro" or larger size **before** attempting to upgrade, or redeploy instead.

Similar to the above, pfSense Plus software can no longer run on Azure A0 instances. Migrate to instances with more memory.

Netgate 3100 (32-Bit ARM) Limitations

Support for the EOL Netgate 3100 device architecture, `armv7`, is being phased out upstream in FreeBSD. While this release still contains base system functionality for the Netgate 3100, several packages are unavailable as they can no longer build for that architecture. The list of packages unavailable for the Netgate 3100 now also includes **Suricata**, **Squid**, and **squidGuard**.

Users who wish to continue using those packages on a Netgate 3100 **should not** upgrade to this release.

3.7.2 General

- pfSense Plus software version 24.03 makes sure the user changes the `admin` account password in the user manager away from the default value. It also ensures that the password is not set to the same value as the username. This validation happens during the setup wizard for new installations, on login and loading any GUI page for existing users, and at the console/shell menu.

  Most users will not notice any difference since they have likely changed their `admin` account password to a secure custom value in the past.

  Resetting the password via the console menu now prompts the user to set a custom password rather than using a default value.

  Note: These restrictions apply to all accounts. Users are also prevented from changing passwords to problematic values.

- The default State Policy has been changed from **Floating** to **Interface Bound** for increased security. However, **Interface Bound** states may have issues in certain cases with IPsec VTI, Multi-WAN policy routing (`route-to`), `reply-to`, as well as with High Availability state synchronization (pfsync) on non-identical hardware.

  The default policy can be toggled back to **Floating** using the **State Policy** option under **System > Advanced** on the **Firewall & NAT** tab.

  There is also an option to override this behavior on a per-rule basis in the advanced options when editing a firewall rule.

  See also:

  - *State Policy History*
  - *Firewall State Policy*
  - *State Policy*

- This release adds support for Packet Flow Data export via `pflow` in PF. This feature natively exports NetFlow/IPFIX flow data to an external collector.

  See also:

  - *Firewall Packet Flow Data*

- This release includes support for enhanced gateway recovery "fail back" by optionally clearing states from lower tier gateways when a more preferred gateway recovers.

  See also:

  - *State Killing on Gateway Recovery*
  - *Gateway Group Options*

- This version requires an updated boot loader, which is automatically handled by the upgrade process for nearly all cases. However, there may be some edge cases where the automatic update does not update the loader currently used by the device. For example, if there are multiple unmirrored disks and the BIOS/EFI Firmware is not booting from the disk containing the updated loader, but an older unrelated installation on a separate disk. One particular case where this can happen is when there is a previous installation to MMC which has been followed by an installation to an add-on SSD without clearing the MMC contents.

  In these cases the best practice is to wipe the unused disk so it cannot interfere. See *Troubleshooting Multiple Disks* for details.

3.7.3 Minor Revision

Devices running pfSense Plus software version 24.03 may be seeing a "24.03_1" update available which is a very minor revision made to address a missing dependency on 64-bit ARM devices (https://redmine.pfsense.org/issues/15433). The revision is kept the same on all platforms for consistency.

Upgrading to this version is safe, but not necessary at this time unless users are running on 64-bit ARM devices and want access to S.M.A.R.T. disk data (e.g. Netgate 2100 devices which have an add-on SSD).

Using the GUI or pfSense-upgrade from the console or shell to upgrade from 24.03 to 24.03_1, the device will want to reboot, but in this case that is unnecessary. However, doing so is harmless except for the minimal downtime involved in the reboot during that upgrade process.

Manually updating from the shell via `pkg update; pkg upgrade` will pull in the new revision and fixed dependency as needed. Run those commands from a shell prompt and confirm that the proposed changes are OK. No additional action is necessary.

Devices which have not yet upgraded to 24.03 or those installed fresh via the Netgate Installer will obtain the latest version automatically and do not require any additional action after upgrading.

3.7.4 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: Interface subnet aliases do not contain IPv6 VIPs #15096

Authentication

- Changed: Prevent usage of the default password in User Manager accounts #15266
- Fixed: PHP errors in LDAP server prevent it from falling back to Local Database #15122

Auto Configuration Backup

- Fixed: `services_acb_settings.php` does not fully validate value of `frequency`, uses value without encoding #15224

Backup / Restore

- Added: Support for CD/DVD drives in the External Configuration Locator (ECL) #14728
- Fixed: DHCP leases may not be restored from older configuration backups #15076
- Fixed: PHP error when generating a notification after detecting a malformed configuration #15157

Captive Portal

- Fixed: Disconnecting a user from Captive Portal may allow previously established connections to continue #13226
- Added: Support using a mask to block MAC addresses in Captive Portal #15257
- Fixed: Old auto-added MAC addresses are not pruned for non-concurrent Captive Portal sessions #15299

Console Menu

- Changed: Dynamically adjust the interface name maximum width in the login banner #13268

DHCP (IPv4)

- Added: Better handling of duplicate IP addresses in static DHCP assignments #13256
- Changed: Reduce log spam when deleting a static DHCP entry #13263
- Added: Explicitly enable/disable DHCP Dynamic DNS updates in each scope #13894
- Fixed: Stale Kea control socket lock file can prevent Kea from starting #14977
- Fixed: Kea does not allow FQDNs for NTP servers but input validation does not prevent them from being added #14991
- Fixed: Kea DHCP PHP error from WINS server value #14996
- Fixed: Kea DHCP sends wrong bootloader file for UEFI #15032

DHCP (IPv6)

- Fixed: DHCP6 client does not take any action if the interface IPv6 address changes during renewal #12947
- Fixed: Shortcut bar on DHCPv6 leases (`status_dhcpv6_leases.php`) navigates to DHCPv4 destinations, not DHCPv6 #15117
- Fixed: DHCPv6 settings page "DDNS Reverse" check box not showing current state #15118

DNS Forwarder

- Added: Option to allow the DNS Forwarder to ignore system DNS servers #14165

DNS Resolver

- Fixed: DNS Resolver host overrides ignore all aliases if first entry has a domain set but no hostname #14942
- Fixed: Applying interface changes may not update default ACLs for the DNS Resolver #15071
- Fixed: Potential local file include vulnerability via DNS Resolver Python Module Script include mechanism #15135
- Fixed: Local DNS resolution behavior does not add an IPv6 nameserver #15139
- Changed: Upgrade Unbound to >= 1.19.1 #15256

Dashboard

- Fixed: Firewall Logs Dashboard Widget is slow and may fail to update #12673

Diagnostics

- Changed: Add ZFS Boot Environment list to status output #15164
- Added: Add Kea information to `status.php` #14953
- Fixed: `crash_reporter.php` displays PHP Error log without encoding #15264
- Added: Add EFI boot information to `status.php` #15297
- Added: Add `loader.conf.lua` contents to `status.php` #15298
- Fixed: Errors in `status.php` IPsec sections when IPsec is not configured #15310

Gateway Monitoring

- Fixed: Gateway behavior differs when the gateway does not exist in the configuration #12920

Gateways

- Fixed: Killing states on downed gateways breaks when `Skip rules when gateway is down` is enabled #15223
- Fixed: Killing states on downed gateways breaks for static interface configurations #15225
- Fixed: Removing a gateway group used as the default gateway results in no default route #15248

Hardware / Drivers

- Fixed: Newer variant models within the PC Engines APU2 platform are not recognized, causing garbled early serial console output #13498
- Added: Recognize QAT 4xxx devices in System Information Widget #15233

IGMP Proxy

- Fixed: IGMP proxy works intermittently #15043

IPsec

- Added: Group-based Mobile IPsec Virtual Address Pool assignment via RADIUS #13227
- Fixed: MSS clamping on VPN traffic does not work on IPsec IPv6 mobile VPNs #14312
- Fixed: Large number of IPsec tunnels causes long filter reload times #14893
- Fixed: IPsec VTI is not created correctly when using a Phase 2 remote type of `Network` #15124
- Fixed: Cannot configure dual stack IPsec tunnel to accept connections from any remote address on both address families #15147

- Fixed: Removing an IPsec Phase 1 entry can either remove the wrong Phase 2 entries or leave orphaned Phase 2 entries in the configuration #15171
- Fixed: Change Mobile IPsec RADIUS accounting to use `accounting_requires_vip` so accounting will not activate for non-mobile VPNs #15176
- Added: Show interface subnet details in a tooltip on the IPsec Phase 2 list #15245
- Fixed: Reordering IPsec Phase 2 entries may result in a malformed configuration #15384

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: `radvd` service shows as stopped in services list when it should be disabled and hidden from that list #14936
- Fixed: Cannot disable Router Advertisements when the interface IPv6 configuration is set to `None` #14967
- Fixed: Router Advertisement daemon does not prioritize IPv6 GUA over ULA #15057

Installer

- Fixed: Clean installation using Auto (ZFS) + MBR (BIOS) does not boot #14930

Interfaces

- Fixed: Sending IPv6 traffic on a disabled interface can trigger a kernel panic #14431
- Fixed: PHP error in `interfaces_qinq_edit.php` when creating a QinQ interface #15181

LAGG Interfaces

- Fixed: Reconfiguring a parent LAGG interface breaks its VLANs #9453

Logging

- Changed: Remove `Time` column from OS Boot logs #15106

Multi-WAN

- Added: Ability to selectively kill states on gateway recovery #855

OpenVPN

- Added: OpenVPN NBDD server options #13085
- Fixed: OpenVPN WINS options may be visible even when NetBIOS is disabled #13087
- Fixed: Some OpenVPN NetBIOS settings are kept even when NetBIOS is disabled #13089
- Fixed: OpenVPN NetBIOS Node Type and Scope ID options are not pushed to clients #13090
- Fixed: `openvpn.auth-user.php` gets stuck at 100% CPU usage when RADIUS authentication times out #14386
- Fixed: OpenVPN forms invalid `route` statements for empty local networks #14919

- Fixed: OpenVPN Wizard fails when a VIP is used #15148
- Changed: Remove deprecated OpenVPN hardware crypto engine option #15188

Operating System

- Added: Operating System support for PF `pflow` packet data flow export #15038
- Fixed: `/etc/rc.local` script content is executed at login instead of during boot sequence #10980
- Fixed: Static ARP assignments lose `permanent` flag in ARP table #14970
- Fixed: Permissions on tmpfs RAM disk for `/var` are too lenient #15054
- Fixed: `pfctl` is unable to retrieve state creator list in certain circumstances #15108
- Fixed: `loader.conf` may be missing `loader_conf_files` so `loader.conf.lua` may not be parsed #15288

PHP Interpreter

- Fixed: Extensions directory is not set in `rc.php_ini_setup` #14488
- Fixed: `check_dnsavailable()` failing even when DNS is available #15127
- Fixed: PHP error display formatting issues #15263

Package System

- Fixed: Extra space in `pkg` configuration file `FreeBSD.conf` #15069

Routing

- Fixed: ICMPv6 Path MTU Discovery breaks with NPT #14290

Rules / NAT

- Added: GUI to configure Packet Flow Data (`pflow`) export #15039
- Added: Kill states using the pre-NAT address #11556
- Changed: Add global option to set default PF State Policy (if-bound vs floating) #15173
- Added: Add per-rule option to set PF State Policy (if-bound vs floating) #15183
- Fixed: Outbound NAT rules using an alias without a matching address family create unexpected PF rules #15197
- Fixed: Advanced rule options tooltip does not show negated Tag option #15214
- Added: Show details of system aliases in tooltip on firewall and NAT rule lists #15234
- Fixed: Egress states remain when killing states for scheduled rules #15252

Setup Wizard

- Changed: Error handling in the Setup Wizard is very user-unfriendly #15302

Traffic Shaper (Limiters)

- Fixed: Packets are passed through dummynet twice when using `route-to` leading to half the expected bandwidth #14854
- Fixed: Fragmented packets delayed by limiters are lost #15156
- Fixed: Reply traffic on a secondary WAN may be dropped when passed through dummynet #15363

Upgrade

- Added: Boot Environments 2.0 #15280

Virtual IP Addresses

- Fixed: `choparp` service is not stopped after deleting Proxy ARP type Virtual IP addresses #14929

Web Interface

- Added: Overflow scrolling for top navigation drop-down menus in Fixed mode #7943
- Fixed: Some messages presented to users contain relative links to pages which may be invalid when triggered from certain packages #13413
- Changed: Update vendor files #13537
- Fixed: `status_interfaces.php` is missing several values for SFP modules #15112
- Changed: Remove `jquery-treegrid` unit testing files #15265
- Added: 50x and 404 error handling to GUI web server configuration #15322

XMLRPC

- Fixed: Secondary node attempts to delete the `admins` group when synchronizing accounts via XMLRPC #15067

3.8 23.09.1 New Features and Changes

This is a maintenance software release including new features and bug fixes.

Consult the *Upgrade Guide* before proceeding with any upgrade.

3.8.1 Security / Errata

FreeBSD Notices

This release includes corrections for several FreeBSD Errata Notices and Security Advisories, including:

- FreeBSD-SA-23:17.pf - TCP spoofing vulnerability in pf(4)
- FreeBSD-EN-23:16.openzfs - Potential ZFS Data Corruption

  For more information about ZFS data corruption, see *ZFS Data Corruption Details* later in this document.
- FreeBSD-EN-23:18.openzfs - High CPU usage by ZFS kernel threads
- FreeBSD-EN-23:17.ossl - ossl(4)'s AES-GCM implementation may give incorrect results
- FreeBSD-EN-23:20.vm - Incorrect results from the kernel physical memory allocator
- Performance issues in OpenSSL have also been identified and corrected, notably with acceleration such as AES-NI.

EFI Issue on Proxmox® VE

Some users of pfSense® software running under Proxmox VE 7.4 have had issues booting Virtual Machines via EFI. This may also affect other versions of Proxmox VE and pfSense software as well as FreeBSD.

Adding a serial port to the VM hardware appears to work around the issue for the time being. A fix for the root cause is under investigation and development for future versions.

At this time the best practice to avoid potential problems is to add a serial port to the VM, then shutdown the VM and start it back up **before** beginning the pfSense software upgrade process.

See also:

See *EFI Boot Issues* for additional recommendations.

ZFS Data Corruption Details

Two data corruption bugs were recently reported against ZFS, including the version of ZFS in recent releases of pfSense software. These bugs have been corrected upstream in FreeBSD and the fixes have been imported into this release.

One bug was in block cloning, which is disabled by default on pfSense software, and thus is unlikely to be a significant concern on this platform. The other bug has been present in ZFS for years and was difficult to trigger.

Given the history of data corruption problems due to hole reporting in files, the corrections for this issue include a preventive measure to disable hole reporting. The downside of disabling hole reporting is the possibly increased disk space usage.

Tip: Users on previous releases of pfSense software can reduce the likelihood of encountering the data corruption issue by creating a *System Tunable* for `vfs.zfs.dmu_offset_next_sync` with a value of `0`.

3.8.2 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: Rules using aliases of type `URL (IPs)` are not generated #14947

DHCP (IPv4)

- Fixed: ISC DHCP responds from a random port #15011

DHCP (IPv6)

- Fixed: PHP error on `services_dhcpv6.php` if the configuration contains an empty `dhcpv6` section #14978

DHCP Relay

- Fixed: Input validation prevents saving DHCPv6 Relay settings #14965

DNS Resolver

- Changed: Update Unbound to 1.18.0_1 to address looping UDP retries when ENOBUFS is returned #14980

IPsec

- Fixed: Mobile IPsec Group Authentication cannot be enabled #14963
- Fixed: Incorrect permissions on `ipsec.auth-user.php` #14974
- Fixed: IPsec log categories set to "Audit" do not function properly or save properly in the GUI #14990
- Changed: Update strongSwan to 5.9.11_3 #15050

Installer

- Added: Add an appropriately named file to install images to indicate what they are #14887

Interfaces

- Fixed: Mulicast traffic on a detached interface causes a panic #14917
- Fixed: PHP Error on `interfaces.php` when creating a PPP interface #14949
- Fixed: DHCP WAN with multiple (2+) IP Alias VIPs may show `0.0.0.0` as an interface address at boot #14966

OpenVPN

- Changed: Update OpenVPN to 2.6.8_1 #15049

Operating System

- Added: Method for users to customize shell initialization behavior #14746
- Fixed: Potential ZFS file corruption #15034

Rules / NAT

- Fixed: Invalid outbound NAT rules break the following rule #15024
- Fixed: Automatic outbound NAT rules show an empty NAT Address #15025

Traffic Graphs

- Fixed: Traffic graph filters apply incorrectly #14892

Upgrade

- Fixed: pfSense-boot does not update the EFI loader #15007

Web Interface

- Fixed: Firewall Maximum Table Entries "default size" is whatever is entered #11566

3.9 23.09 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

Consult the *Upgrade Guide* before proceeding with any upgrade.

Warning: Before attempting the upgrade, check the list of current ZFS Boot Environments (**System > Boot Environments**) and clean up any older entries to ensure they do not consume space which may be needed during the upgrade. See *Check and Clean Up ZFS Boot Environments* for details.

3.9.1 General

- PHP has been upgraded to 8.2.11
- The base operating system has been upgraded to a more recent point on FreeBSD 14-CURRENT
- Support for SCTP has been improved in PF for firewall rules, NAT, and logging. Rules can now act on SCTP packets by port number, previously it was only possible to filter on source or destination address.
- OpenSSL in the base system has been upgraded from 1.1.1t to 3.0.12.

  For details, see *OpenSSL upgraded to 3.0.12*.

- Kea DHCP Server has been added as an opt-in feature preview for IPv4 and IPv6 DHCP service. Kea will eventually replace the ISC DHCPD daemon which is EOL.

  Warning: Kea is not yet feature complete. For details, see *Kea DHCP Server feature preview now available*.

- IPv6 Router Advertisement configuration has been relocated to **Services > Router Advertisement** as a part of the ongoing DHCP Server changes.
- Certain parts of the base system are being migrated to packages rather than grouping them all together in an archive in the "base" package. For the most part this should be entirely transparent to users.

  Specifically, the code from the main pfSense software repository is now a part of the "pfSense" package. This lets management of files be handled entirely by `pkg` rather than carrying them in an archive. This migration is ongoing, so future versions will include additional portions of the system being packaged differently.
- The default driver for NVMe storage devices changed from `nvd(4)` to `nda(4)`. For most users this will be a transparent change since the majority of installations are mounted by label and do not reference a storage device by name.

  Some swap configurations may reference the old device name in `/etc/fstab`. Editing that file and correcting device names from `nvd` to `nda`, followed by a reboot, should restore swap functionality.

  If the new driver is problematic in certain environments the default can be changed back to `nvd(4)` by adding a *loader tunable* for `hw.nvme.use_nvd=1`.

OpenSSL upgraded to 3.0.12

OpenSSL has been upgraded to 3.0.12 from 1.1.1 in FreeBSD.

This change was necessary as OpenSSL 1.1.1 reached its End of Life (EOL) on September 11, 2023. This means there will be no security patches for vulnerabilities affecting OpenSSL 1.1.1.

The OpenSSL team decided to make an explicit jump in numbering from 1.1.x to 3.x to highlight that this new version included major structural, and more importantly, application programming interface (API) and application binary interface (ABI) changes compared to previous OpenSSL versions.

In addition to the differences in the library they also deprecated numerous weak algorithms of various types.

Due to these differences, changing from OpenSSL 1.1 to OpenSSL 3.0 is not a simple upgrade. Netgate developers have handled most of these changes as automatically as possible, though some things may still require manual adjustments. See the warnings in the next section for details.

OpenSSL 3.0.x Upgrade Warnings

Weak Certificate Digests such as SHA1 are Deprecated

Warning: OpenSSL 3.0.x no longer supports certificates signed with SHA1 or other older/weaker hashes. The minimum recommended hash strength is SHA256.

The upgrade process detects usage of weak certificates for the GUI, Captive Portal, and OpenVPN:

- If the **GUI** or a **Captive Portal** zone utilizes a weak CA or server certificate, the upgrade process generates a new self-signed certificate as a stopgap measure to allow the processes to start and let the user in to make any necessary corrections.

- If an **OpenVPN** instance is using a weak certificate, the instance is disabled as there is no viable general automated recovery method.
- **OpenVPN peers** using SHA1 certificates will fail, but such issues must be corrected on the peers. This may mean renewing or reissuing certificates or re-exporting clients for peers if they are currently using weak certificates.
- Other consumers of certificates, such as add-on packages, may be similarly affected but cannot be automatically adjusted.

The best practice is to reconfigure all services utilizing certificates with stronger certificates and to test these functions before performing an upgrade to ensure a smoother transition.

Numerous Deprecated Encryption and Digest Algorithms Removed

Warning: OpenSSL 3.0.x removes a large number of deprecated encryption and digest algorithms. This primarily affects OpenVPN, as other areas had not supported the affected algorithms in some time.

Encryption algorithms removed from OpenVPN:

- ARIA
- Blowfish (e.g. BF-CBC), which was formerly an OpenVPN default
- CAST5
- DES
- DESX
- IDEA
- RC2
- RC5
- SEED
- SM4

Hash algorithms removed from OpenVPN:

- MD4
- MDC2
- SM3
- Whirlpool

On upgrade, tunnels using these deprecated algorithms will be adjusted so they use more secure default values when necessary.

The best practice is to reconfigure tunnels using modern secure encryption and hashing, and to test tunnels before performing an upgrade to ensure a smoother transition.

Other OpenSSL-Related Concerns

- The certificate manager in the GUI can still read and generate certificates using weak hashes, but warns against their use. Avoid creating any new entries using weak hashes. This support will eventually be removed.
- The certificate manager no longer supports importing PKCS#12 archive files which were encrypted with weak ciphers, such as RC2-40. Some operating systems still export using such weak ciphers by default, including macOS and Windows.
- IPsec does not require any adjustments, it still supports SHA1 certificates for the time being and no additional algorithms have been deprecated or removed.
- Unbound does not require any adjustments, it still supports SHA1 certificates for the time being.
- Though the legacy provider for OpenSSL 3.0.x is built and included it does not help to work around the issues mentioned above in any meaningful way.

Kea DHCP Server feature preview now available

The ISC DHCPD server has reached its End of Life (EOL) as of October 5, 2022. Though ISC has stated they may continue to publish security fixes if they are warranted.

Netgate developers have started the migration to Kea DHCP server from ISC as a replacement for ISC DHCPD for IPv4 and IPv6 DHCP service. Basic functionality is present, but not all features are supported at this time.

Warning: Currently the Kea implementation lacks the following DHCP server features:

- Local DNS Resolver/Forwarder Registration for static and dynamic DHCP clients
- Remote DNS server registration
- DHCPv6 Prefix Delegation
- High Availability Failover
- Lease statistics/graphs
- Custom DHCP options

Kea is available as an opt-in preview feature on this release. The UI and settings for Kea are shared with the existing DHCP server. Administrators can easily switch between ISC DHCPD and Kea by navigating to **System > Advanced**, **Networking** tab and changing the new **Server Backend** setting in the **DHCP Options** section.

After Kea integration is complete it will become the default DHCP server on a future release of pfSense software and eventually the deprecated ISC DHCP server will be removed. The exact timing of these changes has not been finalized.

Note: As a part of changes to the DHCP server, IPv6 Router Advertisement configuration has been separated from the DHCP server UI and relocated to **Services > Router Advertisement**.

3.9.2 Security

In addition to OpenSSL and other concerns in the base OS and packages, this release addresses the following vulnerabilities in pfSense software:

- pfSense-SA-23_08.webgui (XSS in `getserviceproviders.php`, #14547)
- pfSense-SA-23_09.webgui (XSS in `status_logs_filter_dynamic.php`, #14548)
- pfSense-SA-23_10.webgui (Authenticated Command Execution in `interfaces_gif_edit.php` and `interfaces_gre_edit.php`, #14549)
- pfSense-SA-23_11.webgui (Authenticated Command Execution in `packet_capture.php`, #14809)

Tip: Patches for these issues are also available in the latest version of the *System Patches Package* for users of pfSense Plus software version 23.05.1 and pfSense CE software version 2.7.0.

3.9.3 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: Firewall rules fail to load when a URL table alias file does not exist #13068
- Added: Type column on Alias lists #13245
- Fixed: Static ARP entries are not configured at boot #14374
- Fixed: Firewall rules are not displayed properly when they reference a URL table alias and its file does not exist #14574

Authentication

- Added: Option to invalidate GUI login session if the client address changes #14265

Backup / Restore

- Changed: Increase timeout for password entry when restoring an encrypted configuration via ECL #14769

CARP

- Added: Add unicast CARP indication and peer address to CARP status #14348
- Fixed: Adding an IP Alias VIP using a unicast CARP VIP as its parent changes the CARP VIP to multicast at the OS level #14586
- Added: Prevent CARP status/maintenance mode from being erroneously toggled #13804
- Fixed: IPsec restart in CARP event scripts does not check VIP properly and never runs #14738

Captive Portal

- Fixed: Captive Portal incorrectly allows leading zeroes on voucher roll numbers #14325
- Fixed: Link to view Captive Portal custom HTML page content does not work #14598

Certificates

- Fixed: Cannot validate Certificates against Certificate Revocation Lists for Intermediate Certificate Authorities #9889
- Added: Improve System menu behavior for Certificate Manager privileges #14347
- Fixed: CA and Certificate renewal page does not properly list some SHA1 certificates as being weak #14678

Configuration Upgrade

- Fixed: PHP Error in `upgrade216_ipsec_create_vtimap()` #14400

Console Menu

- Fixed: Serial console output fails to render properly in certain cases on 4100, 6100, and 8200. #13455
- Fixed: PHP shell script `pfanchordrill` shows duplicate anchor content #14637

DHCP (IPv4)

- Added: Introduce Kea DHCP as an alternative DHCP server for IPv4 and IPv6 #6960

DNS Resolver

- Fixed: DNS Resolver experiences intermittent resolution failures with SSL over TLS due to ASLR #14056
- Added: Unbound Advanced Settings entry for `sock-queue-timeout` #14731
- Changed: Update Unbound to 1.18.0 #14732

Dashboard

- Fixed: System Information widget does not properly form list of active hardware crypto algorithms #14417
- Fixed: Gateway widget tooltip incorrectly indicates some gateways as being default #14542

Diagnostics

- Fixed: `diag_edit.php` warning is not cleared after picking non-directory to load #7589
- Changed: Combining Interface and Rule ID state table filter fields returns no results #14399
- Fixed: Improve error handling in `status.php` #14513
- Added: Status output plugin hook for packages to include their own data #14777

Dynamic DNS

- Added: Include hostname being updated in Dynamic DNS notifications #9504
- Added: Dynamic DNS support for Porkbun #14402
- Fixed: PHP error with One.com Dynamic DNS provider #14649
- Fixed: List of Dynamic DNS types with split host+domain name is missing several providers #14783
- Fixed: Correct name of Gandi LiveDNS #14784
- Fixed: Multi-WAN Dynamic DNS does not fail over when preferred WAN loses link #14829

FreeBSD

- Fixed: Kernel textdumps are not recovered properly on systems with multiple swap partitions #14767

Gateways

- Fixed: Misleading error message when adding/editing static routes which use a gateway on a disabled interface #8846
- Fixed: Cannot select IP Alias VIP with CARP VIP parent in Virtual IP drop-down on Gateway Groups #14524
- Fixed: A default route can remain after setting the default gateway to None #14717

Hardware / Drivers

- Fixed: Unnecessary delay when querying `ixgbe(4)` interfaces with SFP ports #13911
- Added: Options to control Intel Speed Shift #14047
- Fixed: Cavium `qlnxe` / `if_qlnxe` driver is not present #14534
- Fixed: `bnxt(4)` driver errors #14569
- Added: QAT 200xx devices are not recognized as supported #14844

IGMP Proxy

- Fixed: Kernel panic when running IGMP Proxy: Sleeping thread owns a non-sleepable lock #12079
- Fixed: Input validation error when saving IGMP Proxy settings #14301
- Fixed: IGMP Proxy cannot start on VirtIO (`vtnet`) interfaces #14665

IPsec

- Changed: Clarify that the IPsec keep alive check option ignores Child SA Start Action #12762
- Fixed: PHP error in `status_ipsec.php` after removing active IPsec tunnel configuration #14525
- Fixed: Multi-WAN IPsec does not fail over when preferred WAN loses link #14626
- Added: Show IPsec phase 1 authentication type in Mode column of tunnel list #14726
- Fixed: IPsec rejects certificate without any SANs #14831

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: IPv6 neighbor discovery protocol (NDP) fails in some cases #13423

Interfaces

- Fixed: GIF-based interface MTU is assigned to parent interface on boot when parent interface is a LAGG #13218
- Fixed: Cannot add a QinQ interface to a bridge #14377
- Fixed: `find_interface_ipv6_ll()` can return a VIP instead of the interface address #14392
- Fixed: Interface value is not properly validated when submitted on `interfaces_gif_edit.php` and `interfaces_gre_edit.php` #14549
- Fixed: Primary interface address is incorrectly set to the last address on the interface #14623
- Fixed: Link loss causes interfaces configured as Track Interface for IPv6 to lose their IPv4 addresses #14756
- Changed: Eliminate direct config access in `interfaces.php` #14790

Logging

- Fixed: Log rotation is not active if the configuration contains an empty <`syslog`> section or if that section is not present #14517
- Fixed: Per-log settings for file size and retention count are not honored #14545
- Added: Improve SCTP support in `filterlog` #14667

Notifications

- Added: Allow SMTP notifications from non-root processes #14337
- Fixed: PHP error when failing to write `config.cache` #14432

OpenVPN

- Fixed: DCO OpenVPN server bound to Localhost does not pass traffic as expected #14682
- Fixed: Rapidly clicking certain options on OpenVPN Client Overrides can cause hide/show field behavior to invert #13088
- Fixed: OpenVPN can select the wrong interface IP address when multiple addresses are present #14646
- Changed: Prevent weak SHA1 certificates from being used with OpenVPN clients and servers #14677
- Changed: Check for deprecated OpenVPN encryption and digest options on upgrade #14686

Operating System

- Fixed: Error when deleting ZFS Boot Environment created from duplicate of non-default entry #13348
- Fixed: Console and system log may contain unnecessary Netlink debug messages from IPsec #14370
- Added: Support receiving `EAPOL` frames on VLAN `0` in `wpa_supplicant` #14457
- Changed: Automatically configure PF states hash table size #14750
- Fixed: Panic when pfsync attempts to synchronize states between hosts with different rulesets #14804

PHP Interpreter

- Added: Option to configure a custom value for the PHP memory limit #13377
- Fixed: URL scheme is not properly validated in some cases #14356

PPP Interfaces

- Fixed: PPP interface default username/password are not being populated from provider data on `interfaces.php` and `interfaces_ppps_edit.php` #14544
- Fixed: `getserviceproviders.php` does not always validate value of `$connection`, displays without encoding #14547

PPPoE Server

- Fixed: PPPoE Server address input validation is incorrectly allowing IPv6 #13903

Packet Capture

- Added: Change default match modifier from "all of" to "any of" #14650
- Fixed: `packet_capture.php` uses `count` and `length` values in command execution without validation or encoding #14809

Rules / NAT

- Fixed: Ethernet rules using `(self)` as a source or destination make the ruleset fail to load #14478
- Fixed: Ethernet rule Action field hint text lists "reject" option which is not compatible with Ethernet rules #14515
- Fixed: Changes in Ethernet ruleset can lead to incorrect rule and separator order #14705
- Added: Support interface macros in Outbound NAT rules #3288
- Fixed: Negating `<interface> net` when a VIP exists on the interface results in unintended behavior #6799
- Added: Option to wait for interface selection before displaying firewall rules #13124
- Fixed: Default tab on `firewall_rules.php` is not selected if the configuration has no WAN interface #14345
- Added: Support interface groups in firewall rule source/destination fields #14448
- Fixed: "Convert interface definitions" option is not respected when bulk copying rules #14576
- Fixed: Rule separators are ordered incorrectly after removing rules in certain positions #14619
- Fixed: Rule separators are hidden when their index is greater than the number of rules #14621
- Added: Extend support for SCTP in firewall and NAT rules #14640
- Fixed: Separators get shifted when copying firewall rules between interfaces #14691
- Fixed: `ctype_digit()` returns unexpected result for values `<= 255` which can break some validation functions/usages #14702

System Logs

- Fixed: Firewall log parser does not handle SCTP log entries #13940
- Fixed: `status_logs_filter_dynamic.php` does not encode value of `interfacefilter` in raw mode #14548

Traffic Graphs

- Fixed: PHP Error when viewing Traffic Graphs in `iftop` mode #14500

Traffic Shaper (ALTQ)

- Added: Include `ixv` in ALTQ capable NIC list #14408
- Fixed: Kernel panic when using traffic shaping on a PPPoE interface #14497

Traffic Shaper (Limiters)

- Fixed: Limiters have no effect on upload traffic passed by policy routing rules #14039

Translations

- Fixed: Some functions fail if the Language does not exactly match an available Locale #13776
- Fixed: Polish translation contains an invalid `sprintf()` format in the text for `firewall_nat_out_edit.php` #13946

UPnP/NAT-PMP

- Changed: Update miniupnpd to 2.3.3 #14307
- Fixed: Remove broken `stun.sipgate.net` from UPnP STUN server list #14673

Upgrade

- Fixed: Update check in GUI does not always honor the configured proxy settings #14609

User Manager / Privileges

- Fixed: Copy function for User Manager Groups does not work for first group in list #14695

Web Interface

- Changed: GUI pages should use `POST` for AJAX calls, not `GET` #12431
- Fixed: Refactor IPsec code using config access functions #13704
- Fixed: PHP error in CSRF Magic from invalid time value #14394
- Fixed: Breadcrumb path missing on `system_register.php` #14462
- Changed: Prevent weak SHA1 certificates from being used with GUI and Captive Portal #14672
- Fixed: `status_carp.php` and `diag_dump_states.php` unresponsive with large state tables #14758
- Fixed: GUI TCP port is not updated in the configuration when saving with the field empty to remove an existing value #14820

Wireless

- Fixed: PHP error in `handle_wireless_post()` when toggling some wireless interface options #14579

3.10 2.7.2 New Features and Changes

This is a maintenance software release including new features and bug fixes.

Consult the *Upgrade Guide* before proceeding with any upgrade.

3.10.1 Security / Errata

FreeBSD Notices

This release includes corrections for several FreeBSD Errata Notices and Security Advisories, including:

- FreeBSD-SA-23:17.pf - TCP spoofing vulnerability in pf(4)
- FreeBSD-EN-23:16.openzfs - Potential ZFS Data Corruption

  For more information about ZFS data corruption, see *ZFS Data Corruption Details* later in this document.
- FreeBSD-EN-23:18.openzfs - High CPU usage by ZFS kernel threads
- FreeBSD-EN-23:17.ossl - ossl(4)'s AES-GCM implementation may give incorrect results
- FreeBSD-EN-23:20.vm - Incorrect results from the kernel physical memory allocator
- Performance issues in OpenSSL have also been identified and corrected, notably with acceleration such as AES-NI.

EFI Issue on Proxmox® VE

Some users of pfSense® software running under Proxmox VE 7.4 have had issues booting Virtual Machines via EFI. This may also affect other versions of Proxmox VE and pfSense software as well as FreeBSD.

Adding a serial port to the VM hardware appears to work around the issue for the time being. A fix for the root cause is under investigation and development for future versions.

At this time the best practice to avoid potential problems is to add a serial port to the VM, then shutdown the VM and start it back up **before** beginning the pfSense software upgrade process.

See also:

See *EFI Boot Issues* for additional recommendations.

ZFS Data Corruption Details

Two data corruption bugs were recently reported against ZFS, including the version of ZFS in recent releases of pfSense software. These bugs have been corrected upstream in FreeBSD and the fixes have been imported into this release.

One bug was in block cloning, which is disabled by default on pfSense software, and thus is unlikely to be a significant concern on this platform. The other bug has been present in ZFS for years and was difficult to trigger.

Given the history of data corruption problems due to hole reporting in files, the corrections for this issue include a preventive measure to disable hole reporting. The downside of disabling hole reporting is the possibly increased disk space usage.

Tip: Users on previous releases of pfSense software can reduce the likelihood of encountering the data corruption issue by creating a *System Tunable* for `vfs.zfs.dmu_offset_next_sync` with a value of `0`.

3.10.2 pfSense CE

Changes in this version of pfSense CE software.

DHCP (IPv4)

- Fixed: ISC DHCP responds from a random port #15011

DHCP (IPv6)

- Fixed: PHP error on `services_dhcpv6.php` if the configuration contains an empty `dhcpv6` section #14978

DHCP Relay

- Fixed: Input validation prevents saving DHCPv6 Relay settings #14965

IPsec

- Fixed: Mobile IPsec Group Authentication cannot be enabled #14963
- Fixed: Incorrect permissions on `ipsec.auth-user.php` #14974
- Fixed: IPsec log categories set to "Audit" do not function properly or save properly in the GUI #14990
- Changed: Update strongSwan to 5.9.11_3 #15050

Installer

- Added: Add an appropriately named file to install images to indicate what they are #14887

Interfaces

- Fixed: Mulicast traffic on a detached interface causes a panic #14917
- Fixed: PHP Error on `interfaces.php` when creating a PPP interface #14949

OpenVPN

- Changed: Update OpenVPN to 2.6.8_1 #15049

Operating System

- Fixed: Potential ZFS file corruption #15034

Rules / NAT

- Fixed: Invalid outbound NAT rules break the following rule #15024
- Fixed: Automatic outbound NAT rules show an empty NAT Address #15025

Upgrade

- Fixed: pfSense-boot does not update the EFI loader #15007

Web Interface

- Fixed: Firewall Maximum Table Entries "default size" is whatever is entered #11566

3.11 2.7.1 New Features and Changes

This is a maintenance software release including new features and bug fixes.

Consult the *Upgrade Guide* before proceeding with any upgrade.

3.11.1 General

- PHP has been upgraded to 8.2.11
- The base operating system has been upgraded to a more recent point on FreeBSD 14-CURRENT
- Support for SCTP has been improved in PF for firewall rules, NAT, and logging. Rules can now act on SCTP packets by port number, previously it was only possible to filter on source or destination address.
- OpenSSL in the base system has been upgraded from 1.1.1t to 3.0.12.

  For details, see *OpenSSL upgraded to 3.0.12*.
- Kea DHCP Server has been added as an opt-in feature preview for IPv4 and IPv6 DHCP service. Kea will eventually replace the ISC DHCPD daemon which is EOL.

  Warning: Kea is not yet feature complete. For details, see *Kea DHCP Server feature preview now available*.
- IPv6 Router Advertisement configuration has been relocated to **Services > Router Advertisement** as a part of the ongoing DHCP Server changes.

- Certain parts of the base system are being migrated to packages rather than grouping them all together in an archive in the "base" package. For the most part this should be entirely transparent to users.

  Specifically, the code from the main pfSense software repository is now a part of the "pfSense" package. This lets management of files be handled entirely by `pkg` rather than carrying them in an archive. This migration is ongoing, so future versions will include additional portions of the system being packaged differently.

- The default driver for NVMe storage devices changed from `nvd(4)` to `nda(4)`. For most users this will be a transparent change since the majority of installations are mounted by label and do not reference a storage device by name.

  Some swap configurations may reference the old device name in `/etc/fstab`. Editing that file and correcting device names from `nvd` to `nda`, followed by a reboot, should restore swap functionality.

  If the new driver is problematic in certain environments the default can be changed back to `nvd(4)` by adding a *loader tunable* for `hw.nvme.use_nvd=1`.

Troubleshooting

- Due to changes in `pkg`, the new version of `pkg` may not be able to properly locate and use the CA trust store when running on the previous version before upgrading.

  If the firewall is unable to load packages or check for updates after selecting the CE 2.7.1 upgrade branch, run `certctl rehash` from the console, a root shell prompt, or via **Diagnostics > Command Prompt**. This will allow `pkg` to utilize the system certificates until the next reboot.

OpenSSL upgraded to 3.0.12

OpenSSL has been upgraded to 3.0.12 from 1.1.1 in FreeBSD.

This change was necessary as OpenSSL 1.1.1 reached its End of Life (EOL) on September 11, 2023. This means there will be no security patches for vulnerabilities affecting OpenSSL 1.1.1.

The OpenSSL team decided to make an explicit jump in numbering from 1.1.x to 3.x to highlight that this new version included major structural, and more importantly, application programming interface (API) and application binary interface (ABI) changes compared to previous OpenSSL versions.

In addition to the differences in the library they also deprecated numerous weak algorithms of various types.

Due to these differences, changing from OpenSSL 1.1 to OpenSSL 3.0 is not a simple upgrade. Netgate developers have handled most of these changes as automatically as possible, though some things may still require manual adjustments. See the warnings in the next section for details.

OpenSSL 3.0.x Upgrade Warnings

Weak Certificate Digests such as SHA1 are Deprecated

Warning: OpenSSL 3.0.x no longer supports certificates signed with SHA1 or other older/weaker hashes. The minimum recommended hash strength is SHA256.

The upgrade process detects usage of weak certificates for the GUI, Captive Portal, and OpenVPN:

- If the **GUI** or a **Captive Portal** zone utilizes a weak CA or server certificate, the upgrade process generates a new self-signed certificate as a stopgap measure to allow the processes to start and let the user in to make any necessary corrections.

- If an **OpenVPN** instance is using a weak certificate, the instance is disabled as there is no viable general automated recovery method.
- **OpenVPN peers** using SHA1 certificates will fail, but such issues must be corrected on the peers. This may mean renewing or reissuing certificates or re-exporting clients for peers if they are currently using weak certificates.
- Other consumers of certificates, such as add-on packages, may be similarly affected but cannot be automatically adjusted.

The best practice is to reconfigure all services utilizing certificates with stronger certificates and to test these functions before performing an upgrade to ensure a smoother transition.

Numerous Deprecated Encryption and Digest Algorithms Removed

Warning: OpenSSL 3.0.x removes a large number of deprecated encryption and digest algorithms. This primarily affects OpenVPN, as other areas had not supported the affected algorithms in some time.

Encryption algorithms removed from OpenVPN:

- ARIA
- Blowfish (e.g. BF-CBC), which was formerly an OpenVPN default
- CAST5
- DES
- DESX
- IDEA
- RC2
- RC5
- SEED
- SM4

Hash algorithms removed from OpenVPN:

- MD4
- MDC2
- SM3
- Whirlpool

On upgrade, tunnels using these deprecated algorithms will be adjusted so they use more secure default values when necessary.

The best practice is to reconfigure tunnels using modern secure encryption and hashing, and to test tunnels before performing an upgrade to ensure a smoother transition.

Other OpenSSL-Related Concerns

- The certificate manager in the GUI can still read and generate certificates using weak hashes, but warns against their use. Avoid creating any new entries using weak hashes. This support will eventually be removed.
- The certificate manager no longer supports importing PKCS#12 archive files which were encrypted with weak ciphers, such as RC2-40. Some operating systems still export using such weak ciphers by default, including macOS and Windows.
- IPsec does not require any adjustments, it still supports SHA1 certificates for the time being and no additional algorithms have been deprecated or removed.
- Unbound does not require any adjustments, it still supports SHA1 certificates for the time being.
- Though the legacy provider for OpenSSL 3.0.x is built and included it does not help to work around the issues mentioned above in any meaningful way.

Kea DHCP Server feature preview now available

The ISC DHCPD server has reached its End of Life (EOL) as of October 5, 2022. Though ISC has stated they may continue to publish security fixes if they are warranted.

Netgate developers have started the migration to Kea DHCP server from ISC as a replacement for ISC DHCPD for IPv4 and IPv6 DHCP service. Basic functionality is present, but not all features are supported at this time.

Warning: Currently the Kea implementation lacks the following DHCP server features:

- Local DNS Resolver/Forwarder Registration for static and dynamic DHCP clients
- Remote DNS server registration
- DHCPv6 Prefix Delegation
- High Availability Failover
- Lease statistics/graphs
- Custom DHCP options

Kea is available as an opt-in preview feature on this release. The UI and settings for Kea are shared with the existing DHCP server. Administrators can easily switch between ISC DHCPD and Kea by navigating to **System > Advanced**, **Networking** tab and changing the new **Server Backend** setting in the **DHCP Options** section.

After Kea integration is complete it will become the default DHCP server on a future release of pfSense software and eventually the deprecated ISC DHCP server will be removed. The exact timing of these changes has not been finalized.

Note: As a part of changes to the DHCP server, IPv6 Router Advertisement configuration has been separated from the DHCP server UI and relocated to **Services > Router Advertisement**.

3.11.2 Security

In addition to OpenSSL and other concerns in the base OS and packages, this release addresses the following vulnerabilities in pfSense software:

- pfSense-SA-23_08.webgui (XSS in `getserviceproviders.php`, #14547)
- pfSense-SA-23_09.webgui (XSS in `status_logs_filter_dynamic.php`, #14548)
- pfSense-SA-23_10.webgui (Authenticated Command Execution in `interfaces_gif_edit.php` and `interfaces_gre_edit.php`, #14549)
- pfSense-SA-23_11.webgui (Authenticated Command Execution in `packet_capture.php`, #14809)

Tip: Patches for these issues are also available in the latest version of the *System Patches Package* for users of pfSense Plus software version 23.05.1 and pfSense CE software version 2.7.0.

3.11.3 pfSense CE

Changes in this version of pfSense CE software.

Aliases / Tables

- Fixed: Firewall rules fail to load when a URL table alias file does not exist #13068
- Added: Type column on Alias lists #13245
- Fixed: Static ARP entries are not configured at boot #14374
- Fixed: Firewall rules are not displayed properly when they reference a URL table alias and its file does not exist #14574

Authentication

- Added: Option to invalidate GUI login session if the client address changes #14265

Backup / Restore

- Changed: Increase timeout for password entry when restoring an encrypted configuration via ECL #14769

CARP

- Added: Prevent CARP status/maintenance mode from being erroneously toggled #13804
- Fixed: IPsec restart in CARP event scripts does not check VIP properly and never runs #14738

Captive Portal

- Fixed: Captive Portal incorrectly allows leading zeroes on voucher roll numbers #14325
- Fixed: Link to view Captive Portal custom HTML page content does not work #14598

Certificates

- Fixed: Cannot validate Certificates against Certificate Revocation Lists for Intermediate Certificate Authorities #9889
- Added: Improve System menu behavior for Certificate Manager privileges #14347
- Fixed: CA and Certificate renewal page does not properly list some SHA1 certificates as being weak #14678

Console Menu

- Fixed: PHP shell script `pfanchordrill` shows duplicate anchor content #14637

DHCP (IPv4)

- Added: Introduce Kea DHCP as an alternative DHCP server for IPv4 and IPv6 #6960

DNS Resolver

- Added: Unbound Advanced Settings entry for `sock-queue-timeout` #14731
- Changed: Update Unbound to 1.18.0_1 to address looping UDP retries when ENOBUFS is returned #14980

Dashboard

- Fixed: System Information widget does not properly form list of active hardware crypto algorithms #14417
- Fixed: Gateway widget tooltip incorrectly indicates some gateways as being default #14542

Diagnostics

- Fixed: `diag_edit.php` warning is not cleared after picking non-directory to load #7589
- Changed: Combining Interface and Rule ID state table filter fields returns no results #14399
- Fixed: Improve error handling in `status.php` #14513
- Added: Status output plugin hook for packages to include their own data #14777

Dynamic DNS

- Added: Include hostname being updated in Dynamic DNS notifications #9504
- Added: Dynamic DNS support for Porkbun #14402
- Fixed: PHP error with One.com Dynamic DNS provider #14649
- Fixed: List of Dynamic DNS types with split host+domain name is missing several providers #14783
- Fixed: Correct name of Gandi LiveDNS #14784
- Fixed: Multi-WAN Dynamic DNS does not fail over when preferred WAN loses link #14829

Gateways

- Fixed: Misleading error message when adding/editing static routes which use a gateway on a disabled interface #8846
- Fixed: Cannot select IP Alias VIP with CARP VIP parent in Virtual IP drop-down on Gateway Groups #14524
- Fixed: A default route can remain after setting the default gateway to None #14717

Hardware / Drivers

- Fixed: Unnecessary delay when querying `ixgbe(4)` interfaces with SFP ports #13911
- Added: Options to control Intel Speed Shift #14047
- Fixed: Cavium `qlnxe` / `if_qlnxe` driver is not present #14534
- Fixed: `bnxt(4)` driver errors #14569
- Added: QAT 200xx devices are not recognized as supported #14844

IGMP Proxy

- Fixed: Input validation error when saving IGMP Proxy settings #14301
- Fixed: IGMP Proxy cannot start on VirtIO (`vtnet`) interfaces #14665

IPsec

- Changed: Clarify that the IPsec keep alive check option ignores Child SA Start Action #12762
- Fixed: PHP error in `status_ipsec.php` after removing active IPsec tunnel configuration #14525
- Fixed: Multi-WAN IPsec does not fail over when preferred WAN loses link #14626
- Added: Show IPsec phase 1 authentication type in Mode column of tunnel list #14726
- Fixed: IPsec rejects certificate without any SANs #14831

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: IPv6 neighbor discovery protocol (NDP) fails in some cases #13423

Interfaces

- Fixed: GIF-based interface MTU is assigned to parent interface on boot when parent interface is a LAGG #13218
- Fixed: Cannot add a QinQ interface to a bridge #14377
- Fixed: `find_interface_ipv6_ll()` can return a VIP instead of the interface address #14392
- Fixed: Interface value is not properly validated when submitted on `interfaces_gif_edit.php` and `interfaces_gre_edit.php` #14549
- Fixed: Primary interface address is incorrectly set to the last address on the interface #14623
- Fixed: Link loss causes interfaces configured as Track Interface for IPv6 to lose their IPv4 addresses #14756
- Changed: Eliminate direct config access in `interfaces.php` #14790

Logging

- Fixed: Log rotation is not active if the configuration contains an empty <`syslog`> section or if that section is not present #14517
- Fixed: Per-log settings for file size and retention count are not honored #14545
- Added: Improve SCTP support in `filterlog` #14667

Notifications

- Added: Allow SMTP notifications from non-root processes #14337
- Fixed: PHP error when failing to write `config.cache` #14432

OpenVPN

- Fixed: OpenVPN can select the wrong interface IP address when multiple addresses are present #14646
- Changed: Prevent weak SHA1 certificates from being used with OpenVPN clients and servers #14677
- Changed: Check for deprecated OpenVPN encryption and digest options on upgrade #14686
- Changed: Update OpenVPN to 2.6.7 #14985

Operating System

- Added: Method for users to customize shell initialization behavior #14746
- Changed: Automatically configure PF states hash table size #14750
- Fixed: Panic when pfsync attempts to synchronize states between hosts with different rulesets #14804

PHP Interpreter

- Added: Option to configure a custom value for the PHP memory limit #13377

PPP Interfaces

- Fixed: PPP interface default username/password are not being populated from provider data on `interfaces.php` and `interfaces_ppps_edit.php` #14544
- Fixed: `getserviceproviders.php` does not always validate value of `$connection`, displays without encoding #14547

PPPoE Server

- Fixed: PPPoE Server address input validation is incorrectly allowing IPv6 #13903

Packet Capture

- Added: Change default match modifier from "all of" to "any of" #14650
- Fixed: `packet_capture.php` uses `count` and `length` values in command execution without validation or encoding #14809

Rules / NAT

- Added: Support interface macros in Outbound NAT rules #3288
- Fixed: Negating `<interface> net` when a VIP exists on the interface results in unintended behavior #6799
- Added: Option to wait for interface selection before displaying firewall rules #13124
- Added: Support interface groups in firewall rule source/destination fields #14448
- Fixed: "Convert interface definitions" option is not respected when bulk copying rules #14576
- Fixed: Rule separators are ordered incorrectly after removing rules in certain positions #14619
- Fixed: Rule separators are hidden when their index is greater than the number of rules #14621
- Added: Extend support for SCTP in firewall and NAT rules #14640
- Fixed: Separators get shifted when copying firewall rules between interfaces #14691
- Fixed: `ctype_digit()` returns unexpected result for values `<= 255` which can break some validation functions/usages #14702

System Logs

- Fixed: `status_logs_filter_dynamic.php` does not encode value of `interfacefilter` in raw mode #14548

Traffic Graphs

- Fixed: PHP Error when viewing Traffic Graphs in `iftop` mode #14500
- Fixed: Traffic graph filters apply incorrectly #14892

Traffic Shaper (ALTQ)

- Fixed: Kernel panic when using traffic shaping on a PPPoE interface #14497

Translations

- Fixed: Some functions fail if the Language does not exactly match an available Locale #13776

UPnP/NAT-PMP

- Fixed: Remove broken `stun.sipgate.net` from UPnP STUN server list #14673

Upgrade

- Fixed: Update check in GUI does not always honor the configured proxy settings #14609

User Manager / Privileges

- Fixed: Copy function for User Manager Groups does not work for first group in list #14695

Web Interface

- Fixed: Refactor IPsec code using config access functions #13704
- Fixed: PHP error in CSRF Magic from invalid time value #14394
- Fixed: Breadcrumb path missing on `system_register.php` #14462
- Changed: Prevent weak SHA1 certificates from being used with GUI and Captive Portal #14672
- Fixed: `status_carp.php` and `diag_dump_states.php` unresponsive with large state tables #14758
- Fixed: Logo text is partially rendered when using Compact-RED theme on CE #14807
- Fixed: GUI TCP port is not updated in the configuration when saving with the field empty to remove an existing value #14820

Wireless

- Fixed: PHP error in `handle_wireless_post()` when toggling some wireless interface options #14579

3.12 23.05.1 New Features and Changes

This is a maintenance software release including bug fixes for issues in pfSense® Plus software version 23.05.

3.12.1 General

This release follows shortly after pfSense Plus software version 23.05. See the *23.05 Release Notes* for details on changes in that release.

> **Danger:** This version includes newer ZFS features which may not be compatible with older boot loaders. These features **are not** enabled by default when upgrading to avoid potential problems with older boot loaders. Some ZFS commands run at the CLI, such as `zpool status`, may report that a pool can be upgraded, but doing so may also require manually updating the boot loader for the device to boot properly. Upgrading the ZFS pool **is not** necessary at this time. As such, the best practice is to leave it as-is. This will be handled automatically as needed in future updates.
>
> Reinstalling the OS from current installation media will result in having the most recent boot loader and ZFS feature set.

3.12.2 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: PHP error when attempting to bulk import Alias content #14412

CARP

- Fixed: Unicast CARP VIPs do not communicate using IPv6 Link Local Addresses #14385
- Fixed: CARP VIPs can become master too early at boot time #2218

Captive Portal

- Fixed: System crashes or may become unresponsive with Captive Portal #14373
- Fixed: PHP error in Captive Portal `usedmacs` handling #14446

DNS Resolver

- Fixed: Setting system DNS servers can incorrectly modify routes for interface addresses #14288
- Fixed: Discrepancy in “TTL for Host Cache Entries” Description #14358

Dashboard

- Fixed: PHP error from empty <plugins> tag in config.xml #14474

IPsec

- Fixed: Reassembled packets received on a VTI are not forwarded #14396
- Fixed: PHP error in IPsec tunnels list #14458

Interfaces

- Fixed: Panic when changing the parent of a VLAN interface used by limiters #14433

Notifications

- Fixed: Notices incorrectly set system LEDs on hardware with less than three LEDs #14482

Rules / NAT

- Fixed: Outbound NAT rule input validation error when attempting to manually specify "Other Subnet" with a valid address #14354
- Fixed: Enable IPv6 over IPv4 tunneling option results in invalid PF rule #14415

Web Interface

- Fixed: "Max Processes" value is not stored properly when saving on system_advanced_admin.php #14425

3.13 23.05 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

3.13.1 General

- This release includes support for cryptographic acceleration through the Multi-Buffer Crypto for IPsec Library (IPsec-MB, IIMB) which leverages special CPU instructions to accelerate several algorithms for multiple types of VPNs and other uses. See *Cryptographic Accelerators* for details.
- This release includes experimental support for Ethernet (Layer 2) rules. See *Ethernet (Layer 2) Rules* for details.
- As of this release, several new and recent features combined enable using the GUI alone to configure a setup compatible with the AT&T Residential Fiber Network. The same setup should work for any similar ISPs which require special handling such as Priority Code Point tagging on VLAN 0 and 802.1X authentication passthrough to a modem. Previous versions of pfSense Plus software required additional scripts (e.g. "pfatt") and/or manual changes outside the GUI.

  There is a new configuration recipe which covers using these features in the GUI to configure this use case: *WAN Connectivity with 802.1X Authentication Bridging and VLAN 0 PCP Tagging*.

- Unicast CARP support can be configured on a per-VIP basis for environments where multicast CARP cannot function. This is a step toward future enhancements in virtualization and cloud environments which are still under development, including high availability in AWS and Azure. See *VIP Configuration Options* for details.
- WireGuard is now installed by default on new installations. This **does not** affect upgrades or factory reset configurations, only fresh installations.
- Several improvements have been made to memory usage reporting and to reduce some reported cases of increased memory usage in the previous release. See *Memory Management* and *ZFS Tuning* for additional information on memory usage and tuning
- A bug in 23.01 caused some automatic dynamic gateway names to be in mixed case instead of all upper case, which may have led to loss of connectivity until the default gateway or gateway group membership was updated. This bug has been corrected, but anyone who worked around the problem by changing gateway entries will have to correct them again once they have upgraded to 23.05.

3.13.2 Security

- pfSense-SA-23_06.webgui A potential Authenticated Command Execution vulnerability from the `bridgeif` parameter on `interfaces_bridge_edit.php` in the GUI.

Note: Users of pfSense Plus software version 23.01, pfSense Plus software version 22.05.x, and pfSense CE software version 2.6.0 can obtain corrections for this issue from the Recommended Patches area of the *System Patches* package.

- pfSense-SA-23_07.kernel Denial of Service on pfSense Plus software version 23.01 due to a kernel panic from oversize IPv6 packets.

Warning: There is no patch for this issue as it is a problem in the kernel. Users must upgrade to pfSense Plus software version 23.05 or later to correct the problem.

This problem did not affect any version of pfSense Plus software prior to 23.01, nor does it affect any **released** version of pfSense CE software. Users of pfSense CE development snapshots must upgrade to a current snapshot to correct the problem.

3.13.3 Upgrade Paths

Devices running pfSense Plus software version 23.01 can upgrade directly to version 23.05.

Devices running pfSense Plus software version 22.05.1 and earlier must first upgrade to version 23.01, then they can upgrade to version 23.05.

Devices running pfSense CE software version 2.6.0 can also upgrade directly to pfSense Plus software version 23.05. Devices running pfSense CE software version 2.7.0 snapshots dated before the pfSense Plus software version 23.05 release can also upgrade directly. Snapshots after that time may still be able to upgrade, but check the forum for details.

3.13.4 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: Using PF reserved keywords for interface descriptions results in an invalid ruleset #14007
- Fixed: PHP error when attempting to bulk import Alias content #14013
- Fixed: Alias list is not sorted #14015

Authentication

- Added: Option to enable/disable console bell, enabled by default #14002

Auto Configuration Backup

- Fixed: PHP error if the configuration has an empty Auto Configuration Backup section #14076

Captive Portal

- Fixed: PHP error in Captive Portal if `usedmacs` list is empty #14172

Certificates

- Fixed: PHP errors when configuration lacks any certificates #14004
- Fixed: PHP error when exporting a CRL for an old CA #14022
- Fixed: Some blank SAN fields are not ignored when creating a certificate #14124
- Added: Ability to edit Certificate Revocation List properties #14185
- Changed: Add note to inform the user that the "Next Certificate Serial" value is ignored when the "Randomize Serial" option is enabled #14188

Console Menu

- Fixed: Console menu incorrectly shows option 99 on some ARMv7/ARM64 installations #14102
- Added: Print ZFS Boot Environment status in console menu banner #14323

Cryptographic Modules

- Added: Support for cryptographic acceleration using the Multi-Buffer Crypto for IPsec Library (IPsec-MB, IIMB) #14291

DHCP (IPv4)

- Fixed: DHCP Server generates an invalid configuration for static mappings when defining network booting and UEFI HTTPBoot URL #13573
- Fixed: Automatic DHCP failover firewall rules are not present in the ruleset when failover is active #13965
- Fixed: Multiple PHP errors in the DHCP Server when the configuration contains an empty section for an interface #13983
- Fixed: DHCP Server page does not properly select a default interface tab if neither WAN nor LAN are capable of being DHCP servers #14115

DHCP (IPv6)

- Fixed: Typo in `filter.inc` variable for DHCPv6 VLAN priority tag value #14010

DNS Forwarder

- Fixed: DNS Forwarder (`dnsmasq`) is using an invalid combination of options when "Query DNS servers sequentially" is enabled #13655

DNS Resolver

- Fixed: DNS Resolver does not generate automatic ACLs for IPv6 when Network Interfaces is set to "All" #13851

Dashboard

- Fixed: System Information Dashboard widget stops showing CPU details on aarch64 #14204
- Fixed: Changing the default IPsec widget tab removes all widgets #14053
- Fixed: Uptime displays plural seconds for multiple minutes in the System Information Dashboard widget #14176
- Added: Support for Intel PCH temperature values in thermal sensors #14255
- Fixed: PHP error in RSS widget after saving settings #14365

Diagnostics

- Added: Packet Capture GUI with granular control #13382
- Changed: Add more disk information to status output #14103

Dynamic DNS

- Changed: Improve DynDNS help text readability #14186

FreeBSD

- Fixed: Kernel panic accessing the GUI over IPsec in certain environments when using nginx `sendfile` with unmapped mbufs #13938
- Changed: Update Time Zone data to 2023c or later #14209

Gateways

- Fixed: Dynamic gateway names use mixed case instead of upper case, leading to configuration mismatches #14057
- Fixed: Gateway popup in firewall rule list does not indicate current gateway status #14327

Hardware / Drivers

- Fixed: Switch ports on 7100/1100/2100 do not have Auto MDI-X support enabled #13993
- Fixed: Undersized CESA TDMA descriptor pools can be exhausted, leading to errors #14235
- Fixed: Status LEDs on the Netgate 1100 do not function properly #14292
- Fixed: 2100/1100 PCIe bus devices are not recognized #14334
- Fixed: Intel e1000 driver (`em`, `igb`) cannot pass packets tagged with VLAN `0` #12821
- Fixed: Malicious Driver Detection event on `ixl(4)` driver #13003

IGMP Proxy

- Fixed: IGMP Proxy multicast group membership query packets have an invalid checksum #13929

IPsec

- Fixed: Deadlock in Charon VICI interface #13014
- Fixed: PHP error from upgraded IPsec tunnel containing only deprecated ciphers #14009
- Fixed: IPsec Phase 2 rekey failures with some PFS key groups #14217
- Fixed: PHP Error performing IPv6 `ip_in_subnet()` when passing a host addresses within prefix #14256

IPv6 Router Advertisements (radvd/rtsold)

- Fixed: No working IPv6 gateway if upstream RA does not contain M or O flags because rtsold does not execute script #14072

Interfaces

- Added: Priority Code Point (PCP) option on interface configuration #13511
- Fixed: SNMP logs "Device not configured" error message when queries involve built-in switch port interfaces #13976
- Fixed: PHP Error on `status_interfaces.php` with empty switch VLAN group configuration and assigned VLAN interfaces #13981
- Added: Promiscuous Mode option on interface configuration #14295
- Changed: Start `rtsold` immediately after `dhcp6c` sends a request #13492
- Fixed: DHCP client can fail permanently if an interface is down at boot #13671
- Changed: Trim blank characters from static IP address fields on the Interface configuration page #13959
- Fixed: PHP error in `gwlb.inc` when OpenVPN or IPsec instances referred to by assigned interface entries are missing #13973
- Fixed: PHP error when attempting to create a GIF interface on ARM #14035
- Fixed: Bridge interface is not properly validated when submitted on `interfaces_bridge_edit.php` #14052
- Fixed: IPv6 interface configuration race condition can lead to kernel panic #14164

Logging

- Added: Option to control log level of authentication messages in system logs ("Emergency" vs "Notice" level) #12464
- Fixed: Nothing is logged through `syslog` if the configuration contains an empty <`syslogd`> section or if that section is not present #14283

NTPD

- Fixed: PHP error in NTP widget and status with GPS data #13999
- Fixed: PHP error in NTP Server if the configuration contains a partial section of old `openntpd` settings #14033
- Fixed: PHP error when the `timeserver` section of the configuration is empty #14036

Notifications

- Fixed: Identical SMTP notifications repeat in an infinite loop under certain conditions #14031

OpenVPN

- Fixed: SSL/TLS OpenVPN Client fails with `ifconfig` error when the IPv4 Tunnel Network is defined #13350
- Fixed: OpenVPN crashes with Signal 8 with very low fragment size #13943
- Changed: Update OpenVPN Wizard to match current certificate and OpenVPN options #14183

Operating System

- Fixed: Early boot hangs on Hyper-V with Gen2 VMs #13895
- Fixed: OpenVPN and GIF interface create/destroy operations fail due to outdated `linker.hints` #13963
- Changed: Update memory graphs to account for changes in memory reporting #14011
- Fixed: FreeBSD default `cron` jobs are enabled when they should be disabled #14016
- Fixed: Kernel panic from incoming IPv6 connections #14077
- Fixed: Kernel panic when PF passes a large/fragmented ICMP6 packet #14092

PHP Interpreter

- Changed: Update PHP to 8.2.4 #14027
- Fixed: PHP error if a non-privileged shell user attempts an operation which needs to write `config.cache` #14061

PPP Interfaces

- Fixed: IPv6 does not work on secondary PPPoE WAN #13939
- Fixed: PPP interfaces do not request DNS servers when "DNS Server Override" is enabled #13962
- Fixed: PHP Error on `status_interfaces.php` from PPP interface uptime #14117

Package System

- Added: Package plugin hook for pf Ethernet rules #14293
- Added: Package plugin hook for web server configuration stanzas #13054

Rules / NAT

- Added: Support for Ethernet (L2) filtering rules #14308
- Fixed: PHP Error loading Floating rule tab with OpenVPN group rules when there are no OpenVPN instances in the configuration #13953
- Fixed: Custom default state timeouts are not respected in the ruleset #13992
- Fixed: PHP Error enabling ICMP6 using EasyRule #14037
- Fixed: The "Kill States" button does not work consistently #14091
- Changed: Match upstream changes in PF syntax to disable fragment disassembly #14098
- Fixed: PHP error when saving an ICMP firewall rule with no subtypes selected #14267
- Fixed: Associated firewall rule for NAT port forward does not inherit `nosync` property, gets synchronized #14335
- Fixed: PHP error from empty separator #14338

Services

- Fixed: Services Status page and Dashboard widget do not list the `radvd` service with certain static IPv6 configurations #14136

Setup Wizard

- Changed: Update firewall host and domain fields in the Setup Wizard to match the description and warning text from `system.php` #14250

System Logs

- Fixed: PHP error on `status_logs_settings.php` if the configuration contains an empty `syslog` section #13942
- Fixed: `syslogd` tries to bind interfaces with no IP address #14120

Traffic Graphs

- Fixed: PHP Error when viewing Traffic Graphs in `iftop` mode #14236

Traffic Shaper (Limiters)

- Fixed: Traffic shaped by limiters is dropped when routed to a GIF gateway #14055

Traffic Shaper Wizards

- Fixed: PHP errors when re-running Traffic Shaper Wizards with different settings #13915

Upgrade

- Fixed: pfSense Plus Upgrade repo data remains on the system after upgradng #14137
- Fixed: `pfSense-boot` can fail to copy the EFI bootloader #14045

User Manager / Privileges

- Fixed: "All" user group overwritten after assigning an existing user to a group #14363

Virtual IP Addresses

- Fixed: Firewall rules are not reloaded when removing a VIP, outdated rules/entries remain active #13908

Web Interface

- Changed: Replace direct config accesses for the rest of the paths in `system_advanced_admin.inc` #13701
- Changed: Replace direct config accesses in `system_advanced_sysctl` #13702
- Added: Support for `iwlwifi` wireless interfaces #14050

XMLRPC

- Fixed: PHP errors in `xmlrpc.php` during configuration synchronization if the target host has an empty XML tag for a given section #14034
- Fixed: PHP error when XMLRPC client attempts to synchronize without any synchronization settings in the configuration #14182
- Fixed: Filter/NAT rules configured with "No XMLRPC Sync" enabled are still synchronized #14316

3.14 23.01 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

3.14.1 General

- PHP has been upgraded from 7.4 to 8.1
- The base operating system has been upgraded to FreeBSD 14-CURRENT

> **Warning:** As a part of the FreeBSD upgrade this version removes several deprecated IPsec algorithms:
>
> - 3DES Encryption
> - Blowfish Encryption
> - CAST 128 Encryption
> - MD5 HMAC Authentication
>
> The best practice is to reconfigure tunnels using better encryption and test them before performing an upgrade to ensure a smoother transition.
>
> On upgrade, IPsec tunnels will be adjusted to remove any deprecated algorithms from their configuration. The upgrade process will disable tunnels if they have no valid encryption or authentication options remaining. The upgrade process will notify the user of any changes it makes.
>
> This change only affects IPsec and not other uses of these algorithms. For example, BGP can still use TCP-MD5 authentication.

- A long-standing difficult-to-reproduce crash in Unbound during reloading has been addressed. Christian McDonald tracked down the source of the Unbound SIGHUP crashes to a reference counting bug within the MaxMindDB Python module. Both a patch to MaxMind and a port revision to FreeBSD ports were submitted and accepted, and the fix is included in the 23.01 release. It is now safe again to enable DHCP registration alongside Unbound Python mode in pfBlockerNG.

- In addition to the Unbound crash, Christian also identified a memory leak with DHCP registration and Unbound Python mode (#10624). This is largely mitigated by updates to Python and related libraries, but there is additional ongoing work to resolve it further for future release.
- Due to #13507, batch copying rules between interfaces on a previous release may have created multiple rules with the same internal tracker ID. This issue has been corrected, but any rules with duplicate IDs must be corrected manually (e.g. by deleting and re-copying or re-creating the rules).
- The pfBlockerNG package has been updated to match pfBlockerNG-devel. After upgrade it is safe to uninstall pfBlockerNG-devel (keeping settings) and install pfBlockerNG instead.

Note: On systems using ZFS, the first boot post-upgrade will appear to have higher than normal memory usage due to the large volume of filesystem activity that takes place during the upgrade process. This is harmless, however. This is due to ZFS ARC memory usage, which it will yield as needed if other processes require more memory. Rebooting the firewall after the upgrade completes will return the reported memory usage to a normal level.

See also:

ZFS Disk Activity Increases Memory Usage

If an installation continues to show higher than usual memory usage after rebooting, see potentially related issues #14016 and #14011.

3.14.2 Security

pfSense Plus 23.01-RELEASE includes fixes for multiple potential vulnerabilities:

- pfSense-SA-23_01.webgui: A potential XSS vulnerability in diag_edit.php from browsing directories containing specially crafted filenames on the filesystem.
- pfSense-SA-23_02.webgui: A potential XSS vulnerability in system_camanager.php and system_certmanager.php from specially crafted descriptions when editing entries.
- pfSense-SA-23_03.webgui: A potential authenticated arbitrary file creation vulnerability from the name parameter when creating or editing URL table aliases.
- pfSense-SA-23_04.webgui: A potential authenticated arbitrary command execution vulnerability in status.php from specially crafted filenames on the filesystem.
- pfSense-SA-23_05.sshguard: Anti-brute force protection bypass for GUI authentication requests containing certain proxy headers.

Note: Users of pfSense Plus 22.05.x and pfSense CE 2.6.0 can obtain corrections for these issues from the Recommended Patches area of the *System Patches* package.

3.14.3 Errata/Known Hardware Issues

- The **Netgate 1000** does not function on FreeBSD 14 and as a consequence it is unable to upgrade to this release. Attempting to check for updates on a **Netgate 1000** device will print a notification to this effect. No other models are impacted.
- Some older installations of pfSense Plus software on **Netgate 1100**, **Netgate 2100**, and **Netgate 2100 MAX** devices contain an EFI partition which does not have sufficient space to accommodate the new EFI loader for version 23.01 and later. This primarily affects UFS-based systems **initially** installed with version 21.02-p1 or before.

  Users with affected units **must** reinstall pfSense Plus software to run version 23.01 or later.

  Read *Troubleshooting Upgrades on Netgate 1100 and Netgate 2100 Devices* for details.
- The PCI bus in the **Netgate 1100** and **Netgate 2100** models does not currently function on 23.01. This was never an advertised feature, though some users have taken advantage of it in the past. If a device relies on the PCI bus, such as an add-on Wireless card, then consider the impact of upgrading to 23.01 where that will not be available (NG 9622).
- Devices based on "ADI" or "RCC" hardware, such as the **4860**, **8860**, and potentially other similar models, may have issues with the `ichsmb0` and/or `ehci0` devices encountering an interrupt loop, leading to higher than usual CPU usage (NG 8916).

  This can typically be worked around by disabling the affected device, with some caveats.

  To disable the `ichsmb0` device, which will disable the LED status indicators, add the following *Loader Tunable*:

  ```
  hint.ichsmb.0.disabled=1
  ```

  A similar method can be used to disable `ehci0` but doing so will also disable the internal MMC drive, so that should **only** be disabled when the device is booted and running from an add-on SSD.

  This **does not** affect the **2220**, **2440** or **XG-2758**.
- There have been a small number of reports that pfSense Plus software version 23.01 installations using ZFS will not boot in Hyper-V, though it works OK for others (#13895). Test in a lab or non-production environment before attempting to deploy this version. In some cases removing the optical drive from the VM settings before upgrading has allowed it to boot successfully.
- Azure instances now use Gen2 and currently do not have a functional serial console, developers are working to address this in the next release.
- Devices using the i915 video driver require manual changes because FreeBSD moved the driver from the kernel to a package. In most cases this driver is not necessary, but it can be helpful on some platforms for HDMI hotplug support.

  To continue using the driver on 23.01, after the upgrade completes run `pkg install -y drm-510-kmod` from a shell. Then add the following *Loader Tunable*:

  ```
  kld_list="i915kms"
  ```

  Reboot the firewall after making the changes to activate the driver.
- There have been a small number of reports on non-Netgate hardware that accessing the GUI of a pfSense Plus software installation over IPsec can trigger a kernel panic. Developers have not yet been able to reproduce the crash, but there is a workaround for users encountering this problem: Create a *system tunable entry* to set `kern.ipc.mb_use_ext_pgs=0`. See #13938 for details and alternate workarounds.
- Some devices have an issue with the serial console display of password protected consoles and other aspects of the boot process, such as Boot Environment selection. The features may not render properly, but are still

functional. This is not a regression in 23.01 as it also happened on 22.05.x. This has been reported on **Netgate 4100**, **Netgate 6100**, and **Netgate 8200** models. See #13455 for more information.

- The switch ports on the **Netgate 7100** do not have Auto-MDIX enabled on 23.01-RELEASE. If a straight-through Ethernet cable is connecting two **7100** units together (e.g. back-to-back for HA), it will not link on 23.01-RELEASE. This will be addressed in a future release. Replacing the cable with a crossover Ethernet cable will allow it to link in the meantime.
- On **Netgate 3100** units, OpenVPN, GIF, and other types of virtual interfaces may not function on 23.01 until the kernel linker hints are updated by running the following command from a shell prompt:

```
kldxref /boot/kernel
```

  See #13963 for details.

- Dynamic interfaces (DHCP, PPP, etc) with mixed case descriptions may not have the same gateway name after upgrading, leading to a loss of connectivity due to the gateway name not matching the configuration. The simplest workaround is to change the interface name to all capital letters. For a patch, see #14057.

See also:

Numerous additional issues have been fixed since 23.01-RELEASE. Before reporting an error on the forum or elsewhere, first check the existing known issues on Redmine to see if the error already been reported and/or fixed.

3.14.4 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: Alias content is sometimes incomplete when an alias contains both FQDN and IP address entries #9296
- Fixed: Alias with non-resolving FQDN entry breaks underlying PF table #12708
- Fixed: Alias content is sometimes incomplete if the firewall cannot resolve an FQDN in the alias #13282
- Added: Specify CA trust store location when downloading and validating URL alias content #13367
- Fixed: Invalid alias name can still be used by code attempting to validate URL table content #13425
- Fixed: Deleting an alias marks the subsystem as unclean but also unconditionally reloads the filter configuration #13538
- Fixed: Missing descriptions for referrers to firewall aliases cause empty strings for references to be returned when deleting an in-use alias #13539

Authentication

- Fixed: Google LDAP connections fail due to lack of SNI for TLS 1.3 #11626
- Fixed: RADIUS authentication attempts no longer send RADIUS NAS IP attribute #13356
- Fixed: Unable to set web interface session timeout to `0` (i.e. never expire) #13561
- Fixed: Extra remote address information can confuse `sshguard` #13574
- Changed: Improve LDAP debugging #13718

Auto Configuration Backup

- Added: Option to list AutoConfigBackup entries in "reverse" order (newest at top) #11266
- Added: Support for international characters in the AutoConfigBackup Hint/Identifier field #13388

Backup / Restore

- Fixed: Multiple <sshdata> or <rrddata> sections in `config.xml` lead to an XML parsing error during restore #13132
- Fixed: Attempting to restore a 0 byte `config.xml` prints an error that the file cannot be read #13289
- Fixed: Configuration history restores revision no matter which option is clicked in confirmation dialog #13861
- Fixed: RRD restore process does not sanitize filenames from backup XML #13935

Build / Release

- Changed: Disable `pkg` compatibility flag which creates `txz` file extension symbolic links #12782

Captive Portal

- Fixed: Traffic passed by Captive Portal cannot use limiter queues on other rules #13148
- Fixed: Voucher CSV output has leading space before voucher code #13272
- Fixed: Error `dummynet: bad switch 21!` when using Captive Portal with Limiters #13290
- Fixed: Captive Portal breaks policy based routing for MAC address bypass clients #13323
- Fixed: Multiple Captive Portal interfaces do not properly form the list of portal IP addresses #13391
- Fixed: Custom logo or background image is created with two dots (`..`) before the file extension #13396
- Fixed: Captive Portal does not keep track of client data usage #13418
- Fixed: All Captive Portal users are given the same limiter pipe pair #13488
- Fixed: Captive Portal blocked MAC addresses are not blocked #13747
- Fixed: Rules for authenticated Captive Portal users are not removed when a zone is disabled #13756
- Fixed: Captive Portal RADIUS start/stop accounting does not reset counters at each accounting start #13838
- Fixed: Captive Portal does not apply RADIUS bandwidth limits to user pipes #13853

Certificates

- Fixed: CA path is not defined when using `curl` in the shell #12737
- Fixed: Exporting a PKCS#12 file from the certificate manager does not use the intended encryption algorithm #13257
- Fixed: Input validation is not rejecting invalid description characters when editing a CA or Certificate #13387
- Fixed: CRL expiration date with default lifetime is too long, goes past UTCTime limit #13424
- Fixed: ECDSA certificate renewal causes digest algorithm to be reset to SHA1 #13437

Configuration Backend

- Fixed: Input validation is checking RAM disk sizes when they are inactive #13479

Console Menu

- Fixed: Changing an interface IP address and gateway at the console does not save the new gateway if one already exists for the interface #12632
- Fixed: Hidden menu option `100` incorrectly handles HTTPS detection #13258

DHCP (IPv4)

- Added: Improve distinction between online and idle/offline entries in DHCP lease list #10345
- Changed: Clean up DHCP Server option language #13250
- Added: Input validation for numbered DHCP options in static mappings #13584
- Fixed: DHCP server "Disable Ping Check" option does not store value on save #13748

DHCP (IPv6)

- Fixed: `dhcp6c` is not restarted when applying settings when multiple WANs are configured for DHCP6 #13253
- Fixed: Advanced DHCP6 client settings only work for a single interface #13462
- Fixed: "Provide DNS servers to DHCPv6 clients" setting does not reflect a changed value until the page is reloaded #13594
- Fixed: DHCPv6 rules are not created for interfaces with static IPv6 #13633

DNS Forwarder

- Fixed: DNS Forwarder refuses valid retries from clients in certain cases #12901

DNS Resolver

- Fixed: Memory leak in Unbound with Python module and DHCP lease registration active #10624
- Fixed: Unbound crashes with signal 11 when reloading #11316
- Fixed: DNS Resolver is restarted during every `rc.newwanip` event even for interfaces not used in the resolver #12612
- Fixed: DNS resolver does not update its configuration or reload during link down events #13254
- Fixed: DNS Resolver responds with unexpected source address when the DNS over TLS server function is enabled #13393
- Fixed: Incorrect word in "Network Interfaces" help text on `services_unbound.php` #13453
- Changed: Update Unbound to use Python 3.11 instead of Python 3.9 #13867
- Changed: Update Unbound to 1.17.1 #13893

Dashboard

- Fixed: QAT detection on dashboard is incorrect if the driver does not attach #13674
- Fixed: APU1 hardware is not properly identified with current BIOS versions #13471

Diagnostics

- Fixed: File browser on `diag_edit.php` does not encode filenames before display #13262
- Fixed: Neighbor hostnames in the NDP Table on `diag_ndp.php` are always empty #13318
- Fixed: `status.php` uses <name> component of `/tmp/rules.packages.`<name> filenames in shell command without encoding #13426
- Changed: Add multicast group membership (`ifmcstat`) to `status.php` #13731

Dynamic DNS

- Fixed: Namecheap Dynamic DNS responses are not parsed properly #12816
- Fixed: DigitalOcean Dynamic DNS update fails with a "bad request" error #13167
- Fixed: Dynv6 Dynamic DNS client does not check the response code when updating #13298
- Fixed: DNSExit Dynamic DNS updates no longer work #13303

FilterDNS

- Fixed: Resolve interval for `filterdns` may not match the configured value #13067

FreeBSD

- Fixed: Cannot set EFI console as primary console when using both EFI and Serial #13080
- Fixed: CVE-2022-23093 / FreeBSD-SA-22:15.ping #13716

Gateway Monitoring

- Fixed: Marking a gateway as down does not affect IPsec entries using gateway groups #13076
- Fixed: Incorrect function parameters for `get_dpinger_status()` call in `gwlb.inc` #13295

Gateways

- Fixed: Recovering interface gateway may not be added back into gateway groups and rules when expected #13228

Hardware / Drivers

- Fixed: Software VLAN tagging does not work on `ixgbe(4)` interfaces #13381
- Fixed: Intel i226 network interfaces do not honor a manually selected link speed #13529
- Fixed: UDP checksum errors with `ixgbe` interfaces #13883

IPsec

- Fixed: `filterdns` does not monitor remote IPsec gateways for IPv6 address changes #12645
- Fixed: IPsec rejects certificates if any SAN is wildcard rather than rejecting when **all** SANs are wildcard #13373
- Changed: Information box on `status_ipsec.php` says "IPsec not enabled" even when a tunnel is established #13398
- Fixed: Incorrect quoting of Split DNS attribute value in `strongswan.conf` #13579
- Added: Support for ChaCha20-Poly1305 encryption with IPsec #13647
- Changed: Remove deprecated IPsec algorithms (3DES, Blowfish, and CAST 128 encryption; MD5 HMAC/Hashing) #13648

Interfaces

- Fixed: Primary interface address is not always used when VIPs are present #11545
- Added: Support for VLAN `0` #12070
- Fixed: Bridges with QinQ interfaces not properly set up at boot #13225
- Fixed: Several advanced DHCP6 client options do not inform the user when rejecting invalid input #13493
- Changed: Clean up obsolete code in `pfSense-dhclient-script` #13501
- Fixed: Assigned bridge interfaces are not configured at boot #13666
- Fixed: Code that sets IPv6 MTU can unintentionally act on IPv4 addresses #13675

OpenVPN

- Fixed: OpenVPN DCO panics with short UDP packets #13338
- Fixed: OpenVPN crashes after reaching the configured concurrent connection limit #13355
- Fixed: Traffic to OpenVPN DCO RA clients above the first available tunnel IP address is incorrectly routed #13358
- Added: Support for ChaCha20-Poly1305 and AES-128-GCM encryption with OpenVPN DCO #13649
- Fixed: GUI allows configuring OpenVPN DCO with incompatible options (TCP, compression, TAP, net30) #13664
- Fixed: OpenVPN status for multi-user VPN shows info icon to display RADIUS rules when there are none to display #13243

Operating System

- Fixed: Entries for `net.link.ifqmaxlen` duplicated in `/boot/loader.conf` #13280
- Fixed: `vmstat -m` value for `temp` is accounted for incorrectly, resulting in underflows #13316
- Fixed: Memory leak in PF when retrieving Ethernet rules #13525
- Changed: Update Python 3.9.15 to 3.9.16 in base system #13865
- Changed: Add Python 3.11.1 to base system #13866

PHP Interpreter

- Added: Upgrade PHP from 7.4 to 8.1 #13446
- Fixed: `fcgicli` fails to write packets with `nvpair` values that exceed `128` bytes #13638

PPP Interfaces

- Fixed: Services are not restarted when PPP interfaces connect #12811
- Fixed: PPP interface custom reset date/time Hour and Minute fields do not properly handle `0` value #13307

Routing

- Added: Enable `ROUTE_MPATH` multipath routing #9544

Rules / NAT

- Fixed: Rule separator positions change when deleting multiple rules #9887
- Fixed: User is forced to pick an NPt destination IPv6 prefix length even when choosing a drop-down entry which contains a defined prefix length #13240
- Fixed: The `negate_networks` table is duplicated in `rules.debug` #13308
- Fixed: Each line in the NPt destination IPv6 prefix list also contains the network of the previous line when multiple choices are present #13310
- Fixed: Using the copy (not clone) function on firewall rules unintentionally converts interface `address` to interface `net` #13364
- Fixed: PF can fail to load a new ruleset #13408
- Fixed: TCP traffic sourced from the firewall can only use the default gateway #13420
- Fixed: `easyrule` CLI script has multiple bugs and undesirable behaviors #13445
- Changed: Correct DHCP client rule descriptions in the generated firewall ruleset #13505
- Fixed: Copying multiple rules at the same time results in new rules with duplicate tracker IDs #13507
- Fixed: Toggling NAT rules using the button method does not enable/disable corresponding firewall rules #13545
- Fixed: Error creating port forward rule with port alias #13601

Traffic Shaper (ALTQ)

- Added: ALTQ GUI support for Broadcom Netextreme II (`bxe`) interfaces #13304

UPnP/NAT-PMP

- Fixed: UPnP/NAT-PMP status page does not display all port mappings #4500

User Manager / Privileges

- Fixed: RADIUS authentication not working over IPv6 #4154

Web Interface

- Fixed: Unnecessary link tag in login page #7996
- Fixed: "Dark" theme does not sufficiently distinguish between selected and deselected elements in option lists #11730
- Fixed: VGA install defaults to serial as primary console when loading/saving admin GUI settings without making changes #12960
- Changed: Spelling and typo corrections #13357
- Fixed: "Dark" theme uses the same colors for disabled and enabled input fields #13390
- Fixed: Input validation on `system_advanced_firewall.inc` uses incorrect variable references for some fields #13436
- Changed: Update external HTTPS/HTTP links #13440
- Fixed: Table row selection has poor contrast in Dark theme #13448
- Fixed: Changing the GUI port does not redirect the browser to the new port on save #13591

3.15 22.05/22.05.1 New Features and Changes

3.15.1 Version 22.05.1

pfSense Plus software version 22.05.1 is a special patch release which adds hardware support for the Netgate 8200 and newer hardware revisions of the 2100, as well as built-in dynamic repository support.

Important: The majority of pfSense Plus users will not need to run this version unless directed to do so by Netgate TAC. This limited patch release is not currently offered as an upgrade from 22.05.

3.15.2 Version 22.05

This is a regularly scheduled release of pfSense® Plus software including new features and bug fixes.

3.15.3 General

- Added: *OpenVPN Data Channel Offload (DCO)* support (Plus only)

Note: Some OpenVPN features and use cases are not compatible with DCO. See *Limitations* for a list of known DCO limitations.

- Added: *ZFS Boot Environment* (BE) snapshots support (Plus only)
- Changed: Captive Portal and Limiters now use only PF and not IPFW (Plus and CE)

3.15.4 Security

pfSense Plus 22.05-RELEASE includes a fix for the following potential vulnerability:

- pfSense-SA-22_05.webgui: A potential XSS vulnerability in firewall_aliases.php from URL table alias URLs.

Note: Users of pfSense CE 2.6.0 can obtain a correction for this issue from the Recommended Patches area of the *System Patches* package.

3.15.5 pfSense Plus

Changes in this version of pfSense Plus software.

Aliases / Tables

- Fixed: Renaming an alias does not update the alias names in static routes and OpenVPN instances #12727
- Added: Retain descriptions when exporting and importing aliases #12842

Authentication

- Added: GUI option to select the user password hashing algorithm #12855
- Fixed: LDAP setup does not display 'Global Root CA List' option unless another CA also exists #13185

Backup / Restore

- Changed: Comply with current iteration standards when encrypting and decrypting configuration files #12556
- Added: Support encrypted `config.xml` files when restoring via ECL #12685
- Added: Notify user if AutoConfigBackup is unable to successfully upload a backup #12724
- Added: Ability to sort AutoConfigBackup entries #12773
- Fixed: PHP error when upgrading from before configuration revision 21.6, `ipsec_create_vtimap()` is undefined #13097
- Added: Option to restore dashboard widget layout #13125
- Fixed: PHP error restoring DHCP lease data on fresh installation: #13157

CARP

- Changed: Reorganize CARP status page #12701
- Fixed: CARP event storm when leaving persistent CARP maintenance mode. #12961

Captive Portal

- Fixed: Allowed IP/Hostname "Direction" option is never used #12649
- Fixed: `nginx` logs an error that the port is already in use when restarting Captive Portal services #12651
- Fixed: Value of `net.inet.ip.dummynet.*` OIDs in `sysctl` are ignored #12733
- Fixed: Only TCP traffic is passed outbound though IPFW #12834
- Changed: Transition Captive Portal from IPFW to PF #13100

Certificates

- Added: Option to retain the existing serial number when renewing a CA or certificate #13010

Configuration Backend

- Added: Move command line history to a GUI option stored in `config.xml` rather than a manual flag file #12675
- Added: Eliminate duplicate shell commands from history file #12741

Configuration Upgrade

- Added: Playback script to perform a configuration upgrade on an arbitrary `config.xml` file #12973

Console Menu

- Added: Warn the user if they attempt to disable SSH from the menu while connected through SSH #13103

DHCP (IPv4)

- Fixed: Disabling DHCP Server RRD statistics does not work #12710
- Fixed: `HTTPClient` option not sent when using UEFI HTTP Boot #12892
- Fixed: `HTTPClient` option does not work for static mappings #12896
- Fixed: DHCP "Ignore denied clients" option with MAC Deny list set causes DHCP server to not start #12923
- Fixed: DHCP network boot filename can be incorrectly placed in DHCP Pool Options #12986
- Added: Relax DHCP maximum lease time input validation #13118
- Fixed: DHCP lease list displays wrong interface name in the "Leases in Use" summary if DHCP settings for a disabled interface remain in the configuration #13127

DHCP (IPv6)

- Fixed: Multiple DHCP6 WAN connections leads to multiple dhcp6c clients #6880
- Fixed: DHCPv6 server does not skip interfaces configured with invalid ranges #12527
- Fixed: RADVD can be started on both HA nodes when configured with an IPv6 link-local address #12582
- Fixed: Uninitialized array in `array_remove_duplicates()` #12749

DNS Forwarder

- Fixed: DNS Forwarder creates a loop when "Use local DNS, ignore remote DNS servers" is selected #12902
- Fixed: DNS Forwarder custom options may fail after save/restore when options are only separated by newline #13105

DNS Resolver

- Fixed: DNS Resolver does not restart during link up/down events on a static IP address interface #12613
- Added: Automatically create DNS Resolver ACLs for OpenVPN CSO entries #12636
- Fixed: DNS Resolver help text for **System Domain Local Zone Type** option refers users to `unbound.conf(5)` man page instead of pfSense docs #12781
- Fixed: DNS Resolver updates trust anchor at boot even with DNSSEC disabled which can lead to a startup delay of ~2 minutes if the firewall does not have Internet access #12985
- Fixed: DNS Resolver ACLs are not updated when OpenVPN networks change #12991
- Added: DNS Resolver option to keep probing when servers are down #13023

Dashboard

- Fixed: Firewall log widget action icon features stop working when new log entries are added dynamically #6253
- Added: Show `Inactive` for Hardware Crypto output instead of empty field on System Information dashboard widget when nothing can be accelerated #12714

Diagnostics

- Fixed: `diag_pftop.php` does not fully encode output #12915

Dynamic DNS

- Fixed: Dynamic DNS custom IPv6 service fails on 6rd tunnels #12590
- Fixed: GleSYS Dynamic DNS responses are not parsed properly #12672
- Added: IPv6 support for DNSimple Dynamic DNS #12744
- Fixed: Input validation prevents configuring wildcard Dynamic DNS records on GoDaddy #12750
- Added: Support wildcard Dynamic DNS records on DigitalOcean #12752
- Fixed: Google Domains Dynamic DNS responses are not parsed properly #12754
- Fixed: Input validation prevents configuring wildcard Dynamic DNS records on Google Domains #12761
- Fixed: Namecheap Dynamic DNS responses are not parsed properly #12816
- Fixed: Clicking Save & Force Update on a Dynamic DNS entry results in a GUI timeout #12870

Gateway Monitoring

- Fixed: Gateway monitoring should mark gateway as "offline" on PPPoE parent interface disconnect #12633
- Added: Option to disable auto-addition of static routes for `dpinger` #12687
- Changed: Update `dpinger` to 3.2 #12881

Gateways

- Fixed: `fixup_default_gateway()` should not remove a default gateway managed by a dynamic routing daemon #11692
- Fixed: IPv6 link local gateway default status not indicated in GUI #11764
- Fixed: IPv6 gateway group using link local addresses incorrectly logs a gateway change because it not including interface scope properly #12721
- Added: Retain knowledge of previous dynamic gateway IP address when interface is down #12931

Hardware / Drivers

- Added: Chelsio TOE support using the `t4_tom` module #9091
- Fixed: Hyper-V RSC support in `hn(4)` driver is enabled by default and results in very low throughput #12873

High Availability

- Added: Use consistent pf host ID and add GUI option to set a custom host ID in state synchronization settings #12702

IGMP Proxy

- Fixed: IGMP Proxy server is restarted during every `rc.newwanip` event #12609

IPsec

- Added: Option to choose default tab in IPsec status Dashboard widget #2456
- Fixed: IPsec VTI phase 2 traffic selectors default to address when defined as a network #11226
- Fixed: `filterdns` does not monitor remote IPsec gateways for IPv6 address changes #12645
- Fixed: Disallow remote gateway of `0.0.0.0` for VTI mode #12723
- Fixed: VTI gateway status stuck as "pending" after reboot #12763
- Changed: Update strongSwan #12934
- Fixed: ESP description in IPsec phase 2 proposal help text is ambiguous #12953
- Fixed: IKEv2 Mobile IPsec clients do not receive `INTERNAL_DNS_DOMAIN` (value `25`) attribute #12975
- Added: GUI option for IPsec `dns-interval` setting #13057
- Fixed: Delete function for IPsec SAD entries on `status_ipsec_sad.php` does not work #13071
- Fixed: Mobile IPsec clients cannot be manually disconnected from IPsec status screen #13131

Installer

- Fixed: Support encrypted `config.xml` files when restoring during install #12691
- Added: Recover existing SSH keys during installation #12809

Interfaces

- Added: Show SFP module details on `status_interfaces.php` #8861
- Added: Improved support for USB interfaces that may not always be present #9393
- Fixed: PPPoE WAN IP address different than expected when set static by ISP #11629
- Fixed: `devd` is not configured to act on USB interface attach/detach events #12606
- Changed: Restart services on interface changes #12619
- Fixed: Interface status "Total Interrupts" display is non-functional #12735

- Fixed: L2TP/PPTP interface assignment page loses some values after input validation error #12780
- Fixed: Link-Local IPv6 address on WAN with MAC spoofing changes if there is an IP Alias on WAN #12790
- Fixed: Link-local address does not reset after removing MAC address spoofing #12794
- Fixed: Disabled Captive Portal configuration prevents adding an interface to a bridge #12866
- Fixed: The ruleset is not regenerated after assigning an interface #12949

L2TP

- Fixed: L2TP MPD configuration is not updated when a dynamic WAN IP address changes #13066
- Fixed: L2TP stays bound to previous IP address after static IP address change #13082
- Fixed: Static routes to destinations at L2TP clients are not re-added after a client reconnects #13099

LAGG Interfaces

- Added: GUI option to configure layers for LACP hash #12819

Notifications

- Fixed: Slack notification options only allow `` -`` as a special character in channel names #13083

OpenVPN

- Fixed: OpenVPN IPv4 Tunnel Network incorrectly allows hostnames #11416
- Fixed: OpenVPN stays bound to previous IP address after interface changes #11864
- Added: OpenVPN option to limit concurrent connections per user #12267
- Fixed: OpenVPN does not clear old Cisco-AVPair anchor rules in some cases #12332
- Added: Use deferred client connections in OpenVPN #12407
- Fixed: OpenVPN re-synchronization also synchronizes override entries unnecessarily in some cases #12628
- Fixed: Automatic filter reload with OpenVPN client gateway uplink happens too soon or not at all #12771
- Fixed: PHP error when terminating OpenVPN sessions via the dashboard widget #12817
- Fixed: OpenVPN status display for TAP mode services shows peer-to-peer instead of client list in certain cases #12884
- Fixed: GUI does not reject an invalid OpenVPN tap mode configuration with an empty tunnel network "Bridge DHCP" disabled #12887
- Fixed: FQDN in network alias is omitted from OpenVPN networks list #12925
- Changed: Warn about OpenVPN shared key deprecation #12981
- Fixed: OpenVPN `remote_cert_tls` option does not behave correctly when enabled and later disabled #13056
- Fixed: Gateway events for IPv6 affect IPv4 OpenVPN instances and vice versa #13061
- Fixed: OpenVPN client `tls-client/client` configuration directive not handled properly #13116
- Changed: OpenVPN status page improvements #13129

- Fixed: OpenVPN `client-connect` file contains `topology` #13133
- Fixed: Per-user `route` files are not removed from `/tmp` when they are no longer needed #13145
- Fixed: OpenVPN override IPv4 tunnel network field changing value improperly #13274

Operating System

- Fixed: pf `hostid` value is handled inconsistently #12703
- Fixed: Some `sysctl` OIDs in `loader.conf.local` are silently removed #12862
- Fixed: Output from `pfctl -vvsr` does not include `ridentifier` value in the expected location #12868

PPP Interfaces

- Fixed: PPPoE WANs fail to reconnect after parameter negotiation failure #13092

PPPoE Server

- Fixed: PPPoE server panics with multiple client connections #13210

Package System

- Fixed: Packages are not automatically reinstalled when restoring configuration using the installer #12105
- Fixed: Packages with custom `internal_name` values do not reinstall properly when restoring a backup #12766
- Fixed: `write_rcfile()` does not create `rc_restart()` entry #13004

Packet Capture

- Added: Button to clear previous packet capture data #12968

Routing

- Fixed: Setting a default gateway of "None" does not remove the default gateway from the routing table #12536
- Fixed: Cannot remove IPv6 static routes #12728
- Fixed: Explicit PPPoE disconnect of a WAN Gateway Group member may not restore a default route. #13048

Rules / NAT

- Added: Toggle button to disable/enable multiple firewall rules #2505
- Added: Port forward NAT rules with "any" protocol #4259
- Added: Allow NPt to use dynamic IPv6 networks #4881
- Added: Button to copy rules from one interface to another #8365
- Fixed: Automatic Outbound NAT mode can create incorrect rules in some cases #11984
- Added: Utilize new `pfctl` abilities to kill states #12092

- Fixed: NAT reflection does not work for IPv6 port forwarding rules when configured for NAT+Proxy mode #12319
- Added: Allow the selection of "any" interface in floating rules #12392
- Fixed: Applying firewall rule changes does not clear dirty flag for aliases subsystem #12678
- Fixed: Automatic Outbound NAT rules do not include OpenVPN CSO entries #12792
- Fixed: Error loading ruleset due to illegal TOS value #12803
- Fixed: High latency and packet loss during a filter reload #12827
- Fixed: On startup "No routing address with matching address" might appear #12847
- Fixed: Some action buttons are always active for firewall rules, even if no rules are selected #12871
- Added: Toggle button to disable/enable multiple entries on NAT pages #12879
- Fixed: Delete button is always active for NAT rules, even if no rules are selected #12957
- Fixed: NAT Reflection generates duplicate rules when internal interface contains multiple VIPs in the same subnet #13012
- Fixed: NAT generates duplicate `no nat on` rules for port forwards with a destination of `Any` #13015
- Fixed: Input validation requires a gateway for floating `match out` rules #13027
- Fixed: Empty `negate_networks` table breaks policy routing rules #13049
- Fixed: The `negate_networks` table is not updated when an OpenVPN server is deleted #13055
- Added: Allow auto prefix with manual prefix-length in NPt #13070
- Fixed: Info icon on `firewall_nat_out.php` is incorrectly placed in manual outbound NAT mode #13164
- Fixed: Changing the redirect target for a Port Forward with an associated filter creates an incorrect firewall rule #13171
- Fixed: Incorrect usage of DSCP hex value #13178

SNMP

- Fixed: SNMP daemon is restarted during every `rc.newwanip` event #12611

Services

- Fixed: NTP service is not listed on `status_services.php` unless `config.xml` contains NTP configuration data #12775
- Fixed: Stale `sshdkeys.dirty` lock file prevents generating SSH server keys #13139

Traffic Shaper (ALTQ)

- Changed: Remove code references to unused `reset` parameter from traffic shaper pages #13042

Traffic Shaper (Limiters)

- Fixed: Incorrect ICMP reply when using limiters #9263
- Fixed: Pie and `fq_pie` are missing options and do not handle floating point number input correctly #12003
- Fixed: Utilize `dnctl(8)` to apply limiter changes without a filter reload #12579
- Fixed: Traffic routed through DUMMYNET by PF fails when IPFW is enabled #12954

Traffic Shaper Wizards

- Fixed: Traffic Shaper wizard can produce an invalid ruleset when configured with an IPv4 upstream SIP server #12937
- Fixed: Traffic shaper wizard rewrites Mbits to Kbits #13086

UPnP/NAT-PMP

- Added: uPnP fails to properly give out subsequent reservations when multiple gaming systems are playing the same game/using the same port. #7727
- Changed: Reorganize UPnP options #12624

Unknown

- Fixed: Many `exec()` functions do not use full path to executable files #11941

Upgrade

- Fixed: Upgrade does not work when using only IPv6 DNS servers #13162

User Manager / Privileges

- Fixed: Icon missing for user manager entries with a scope other than "user" #13174

Web Interface

- Fixed: Lack of DNS or Internet connectivity causes GUI to be slow #12141
- Fixed: Zero-value prefix IPv6 addresses are mishandled #12440
- Added: Option to filter state table contents by rule ID #12616
- Fixed: Changing RAM disk size does not prompt to reboot #12876
- Fixed: Input validation for IPv6 addresses allows invalid address compression in some cases #13069
- Added: Trim whitespace from MAC addresses in user input #13109

Wireless

- Fixed: Wireless interface WPA configuration fields are always visible #12998
- Fixed: Duplicate wireless interfaces are created at boot #12999

XMLRPC

- Fixed: Deleting a user on the primary node does not delete its home directory on secondary node during XMLRPC sync #12940

3.16 22.01/2.6.0 New Features and Changes

This is a regularly scheduled release of pfSense® CE and pfSense Plus software including new features, additional hardware support, and bug fixes.

> **Warning:** When upgrading to pfSense Plus 22.01 and later versions, the `pfSense-upgrade` process will forcefully reinstall all operating system packages and add-on packages to ensure a consistent state and package set. This may increase the time the upgrade will take to download and install.

3.16.1 Security

This release includes corrections for the following vulnerabilities in pfSense software:

- pfSense-SA-22_01.webgui (File overwrite in `services_ntpd_gps.php`, #12191)
- pfSense-SA-22_02.webgui (Potential vulnerabilities with route collection on `diag_routes.php` , #12257)
- pfSense-SA-22_03.webgui (Potential vulnerabilities in OpenVPN form validation, #12677)
- pfSense-SA-22_04.webgui (XSS in `pkg.php`, #12725)

3.16.2 Errata

- There is a patch available to improve NAT behavior for UPnP and multiple game consoles or clients playing the same game but the fix was discovered too late for it to be included in 22.01/2.6.0.

  For additional details and instructions on how to apply the patch, see Redmine issue #7727 note #74 and #75, the Github commit, and the forum thread for testing feedback.

3.16.3 General

- This release contains several significant changes to IPsec for stability and performance. Read the IPsec section of this document carefully.

  > **Warning:** IPsec VTI interface names have changed in this release. Configurations will be updated automatically where possible to use the new names.
  >
  > Check the interface names of assigned VTI instances under **Interfaces > Assignments** to ensure they are correct after the upgrade completes.

If any third party software configurations or other manual changes referenced the old IPsec VTI interface names directly (e.g. `ipsecNNNN`) they must be updated to the new format.

- ZFS is now the default filesystem for new installations of pfSense Plus and pfSense CE software on all platforms which support booting from ZFS.
  - It is not possible to change from UFS to ZFS in place, a reinstallation of pfSense Plus or CE is required to migrate from UFS use ZFS.
  - The ZFS pool name and datasets have also been updated and optimized. Users who were already using ZFS may want to reinstall as well to ensure they have the most optimal disk layout.
  - pfSense Plus software has a new ZFS dashboard widget to track the status of disks using ZFS.
- **Log Compression** for rotation of System Logs is now disabled by default for new ZFS installations as ZFS performs its own compression.

Tip: The best practice is to disable **Log Compression** for rotation of System Logs manually for not only existing ZFS installations, but also for any system with slower CPUs. This setting can be changed under **Status > System Logs** on the **Settings** tab.

- The default password hash format in the User Manager has been changed from bcrypt to SHA-512. New users created in the User Manager will have their password stored as a SHA-512 hash. Existing user passwords will be changed to SHA-512 next time their password is changed.

Note: User Manager passwords are only stored as a hash, thus existing users cannot be automatically changed to the new format. To convert a user password from an older hash format, change the password for the user in the User Manager.

- The firewall now *bootstraps its clock* at boot in multiple ways, one of which utilizes multiple NTP servers with static IP addresses from Google Public NTP. This avoids a chicken-and-egg problem where the firewall cannot resolve NTP servers because DNSSEC, which is enabled by default, cannot function when the clock is inaccurate. The firewall performs this sync once per boot before it starts the NTP daemon.

Note: This behavior can easily be changed or disabled. See *Changing Clock Bootstrap Behavior*.

- Several areas of the documentation have been rewritten and updated for these releases. Notably, the IPsec and OpenVPN sections have been updated significantly including all of the related configuration recipes.

3.16.4 pfSense Plus

PHP Interpreter

- Fixed: PHP exits with signal 11 on SG-3100 when calling PCRE functions #11466

3.16.5 pfSense CE

Aliases / Tables

- Fixed: Error loading rules when URL Table Ports content is empty #4893
- Fixed: Mixed use of aliases in a port range produces unloadable ruleset #11818
- Fixed: Unable to create nested URL aliases #11863
- Fixed: Creating or editing aliases fails with multiple hosts separated by spaces #12124
- Fixed: When attempting to delete an in-use alias, input validation only prints the first item using the alias in the error message #12177

Authentication

- Changed: Use SHA-512 for user password hashes #10298
- Fixed: Deny SSH access for `admin` and `root` users when the `admin` GUI account is disabled #12346

Backup / Restore

- Fixed: Restoring from AutoConfigBackup presents reboot type selection option then reboots automatically #10662
- Added: Backup and restore SSH host key(s) #11118
- Fixed: Output from reboot process is printed on Backup & Restore page when restoring a configuration file #11909
- Fixed: Custom value for AutoConfigBackup schedule Hours is not shown when loading the settings page #11946
- Added: AutoConfigBackup performance improvements #12193
- Fixed: Viewing an AutoConfigBackup entry takes approximately 60 seconds to completely load #12247
- Changed: Explicitly state where AutoConfigBackup stores encrypted backup data #12296

Build / Release

- Changed: Remove deprecated `libzmq` code and references #12060

CARP

- Fixed: Cannot enter persistent CARP maintenance mode when CARP is disabled #11727
- Fixed: When a CARP VIP VHID change is synchronized to a secondary node, the CARP VIP is removed from the interface and the old VHIDs remain active #12202
- Fixed: Changing VHID on CARP VIP does not update VHID of related IP Alias VIPs #12227
- Fixed: `rc.carpmaster` only sends notifications via SMTP #12584

Captive Portal

- Fixed: Vouchers may expire too early when using RAM disks #11894
- Fixed: Incorrect variable substitution in captive portal error page #11902
- Fixed: Clicking "logout" on portal page does not function when logout popup is disabled #12138
- Fixed: Captive Portal database and `ipfw` rules are out of sync after unclean shutdown #12355
- Fixed: Captive Portal input validation for "After authentication Redirection URL" and "Blocked MAC address redirect URL" is swapped #12388
- Fixed: Captive Portal online user statistics data is not cleared on unclean shutdown #12455

Certificates

- Fixed: Certificate Revocation tab does not list active users of CRL entries #11831
- Fixed: Certificate manager reports CA as in use by an LDAP server when LDAP is not configured for TLS #11922
- Fixed: Certificate Manager performs redundant escaping of special characters in certificate DN fields #12034
- Added: Input validation to prevent unsupported UTF-8 characters from being used in certificate subject components #12035
- Fixed: Certificate Manager shows incorrect DN for imported entries with UTF-8 encoding #12041

Console Menu

- Fixed: Cannot configure WAN IP address with `/32` CIDR mask via console menu #11581
- Changed: Suppress kernel messages when loading `dummynet` and thermal sensor modules #12454

DHCP (IPv4)

- Added: DHCPv4 client does not support `supersede` statement for option 54 #7416
- Added: Support for UEFI HTTP Boot option in DHCPv4 Server #11659
- Fixed: DHCPv4 server configuration does not include ARM TFTP filenames #11905
- Fixed: ARM 32/64 network boot options are not parsed on Static DHCP Mapping page #12216

DHCP (IPv6)

- Fixed: DHCPv6 Server should not offer configuration options for unsupported PPPoE Server interfaces #12277

DHCP Relay

- Fixed: PHP error if no DHCPv6 Relay interfaces are selected #11969

DNS Resolver

- Fixed: Unbound crashes with signal 11 when reloading #11316
- Fixed: Unbound fails to start if its configuration references a python script which does not exist #12274
- Fixed: Unbound falls back to using all outgoing network interfaces if manually selected outgoing interface(s) are unavailable #12460

Dashboard

- Fixed: System Information widget unnecessarily polls data for hidden items #12241
- Fixed: IPsec widget generates errors if no tunnels are defined #12337
- Fixed: IPsec widget treats phase 1 in "connecting" state as connected #12347
- Added: Disks dashboard widget to replace Disk Usage section of System Information widget #12349
- Fixed: Thermal Sensors Dashboard widget filter for negative values refers to invalid variable #12470

Diagnostics

- Fixed: State table content on `diag_dump_states.php` does not sort properly #11852
- Changed: Hide "Reboot and run a filesystem check" for ZFS systems #11983
- Fixed: "GoTo line #" function does not work on `diag_edit.php` #12050
- Fixed: Sanitize WireGuard private and pre-shared keys in status output #12256
- Added: Include firewall rules from packages which failed to load in status output #12269
- Added: Include firewall rules generated from OpenVPN RADIUS ACL entries in status output #12316
- Fixed: ARP table interface column empty for entries on unassigned interfaces #12698

Dynamic DNS

- Added: Option to set interval of forced Dynamic DNS updates #9092
- Added: Support DNS Made Easy authentication without a username #9341
- Fixed: RFC 2136 Dynamic DNS client uses IPv6 alias VIP instead of Track IPv6 address for AAAA records #11816
- Added: New Dynamic DNS Provider: Strato #11978
- Fixed: Dynamic DNS cache expiration time check calculation method may cause update to happen on the wrong day #12007
- Fixed: NoIP.com incorrectly encodes Dynamic DNS update credentials #12021
- Added: New Dynamic DNS Provider: deSEC #12086
- Added: Support Check IP services which return bare IP address values #12194

- Fixed: Yandex Dynamic DNS client does not set the `PddToken` value #12331
- Added: Dynamic DNS client proxy support #12342
- Fixed: Update Dynamic DNS code for one.com to use their new login process #12352
- Fixed: Dynamic DNS updates do not respect certificate authority trust store #12589
- Fixed: Dynamic DNS client updates using a private IP address when it cannot determine the public IP address #12617
- Fixed: Dynamic DNS may not use the correct interface when updating during failover #12631

FreeBSD

- Fixed: Duplicate `comconsole_port` lines in `/boot/loader.conf` #11653
- Changed: Upgrade to `pkg` 1.17.x #12171

Gateways

- Added: Support DNS server gateway selection on `system.php` for multiple gateways not assigned to interfaces #12116
- Fixed: Default IPv4 gateway may be set to IPv6 gateway value in certain cases #12282

Hardware / Drivers

- Added: Support for network interfaces using the `qlnxe` driver #11750

High Availability

- Fixed: Incorrect RADVD log message on HA event #11966

IGMP Proxy

- Added: Support `0` CIDR mask for IGMP Proxy networks #7749

IPsec

- Fixed: Disconnected IPsec phase 2 entries are not shown in IPsec status #6275
- Fixed: UDP fragments received over IPsec tunnel are not properly reassembled and forwarded #7801
- Fixed: EAP-RADIUS Mobile IPsec clients with RADIUS-assigned addresses do not get additional configuration attributes #11447
- Fixed: Incorrect phase 2 entry removed when deleting multiple items consecutively #11552
- Fixed: strongSwan configuration contains incorrect structure for mobile pool DNS records #11891
- Fixed: IPsec status tunnel descriptions are incorrect #11910
- Changed: PC/SC Smart Card Daemon `pcscd` running on all devices at all times, should be optional #11933
- Fixed: IPsec status fails when many tunnels are connected #11951

- Fixed: Mobile IPsec advanced RADIUS parameters do not allow numeric values with a decimal point #11967
- Fixed: Mobile IPsec NAT/BINAT entries missing from firewall rules #12023
- Fixed: Applying IPsec settings for many tunnels is slow or times out #12026
- Fixed: Gateway alarm always triggers IPsec restart #12039
- Changed: Improve IPsec identifier settings #12044
- Fixed: IPsec status IKE disconnect button drops all connections for the IKE ID, not a specific IKE SA ID #12052
- Fixed: Tunnels with conflicting REQID values can lead to multiple identical Child SA entries #12155
- Added: IPsec keep alive option to initiate phase 2 without using ICMP #12169
- Added: Add connect/disconnect buttons to IPsec dashboard widget #12181
- Added: GUI options to configure IKE retransmission behavior #12184
- Fixed: IPsec status shows connect buttons while tunnel is connecting #12189
- Fixed: IPsec writes CRL files when tunnel does not use certificates #12195
- Fixed: IPsec settings fail to apply when a remote gateway is set to an FQDN and there are no DNS servers available #12196
- Fixed: Mobile IPsec phase 1 should not display "Gateway duplicates" option #12197
- Fixed: Disabling an IPsec phase 1 entry does not disable related phase 2 entries #12198
- Fixed: Disabled IPsec VTI interfaces are always created #12212
- Fixed: IPsec bypass rules display help text under each entry #12236
- Fixed: IPsec phase 1 entry with `0.0.0.0` as its remote gateway does not receive correct automatic firewall rules #12262
- Changed: Update "IPsec Filter Mode" option values and help text to reflect that VTI mode also helps transport mode (e.g. GRE) #12289
- Fixed: IPsec manual initiation and termination should use a timeout value or forced actions #12298
- Fixed: IPsec tunnels using a gateway group do not get reloaded in some cases #12315
- Fixed: IPsec Phase 2 entry incorrectly orders proposals in AH mode #12323
- Fixed: Hash algorithm GUI options are disabled after switching a phase 2 entry to AH mode #12324
- Fixed: IPsec VTI interface remote endpoint is not resolved the correct way #12328
- Fixed: Incorrect label for IPsec DH group 32 #12350
- Added: Distinguish between policy-based and route-based entries on IPsec status SPD tab #12397
- Fixed: Console boot output includes `Configuring IPsec VTI interfaces` when no VTI interfaces are configured #12419
- Changed: Add IPsec phase 2 BINAT subnet size input validation #12430
- Fixed: IPsec initiates on HA backup node when a tunnel interface is set to a gateway group #12566
- Fixed: IPsec Mobile Client RADIUS Advanced parameters are not reset to default values when disabled #12575

IPv6 Router Advertisements (RADVD)

- Fixed: `radvd` only responds to the first Router Solicitation received after each multicast Router Advertisement #10304
- Fixed: "Default preferred lifetime" router advertisement validation check uses incorrect variable #12159
- Fixed: IPv6 RA DNSSL lifetime is too short, not compliant with RFC 8106 #12173
- Fixed: Default IPv6 router advertisement intervals and lifetime are too low #12280
- Fixed: "Default preferred lifetime" field for IPv6 RA does not have input validation #12439
- Fixed: IPv6 interface prefix change not reflected in RADVD configuration #12604
- Fixed: Router Advertisement DNS search domain from one interface may unintentionally be used by other interfaces #12626

Installer

- Added: Restore RRD and extra data from configuration backups when restoring during installation #12518
- Fixed: Minnowboard Turbo cannot boot a clean install #12707

Interfaces

- Fixed: GRE and GIF tunnels on dynamic IPv6 interface are not brought up during boot #6507
- Fixed: Interface column empty in list of GIF tunnels when using IP Alias on CARP VIP as Interface #11337
- Fixed: QinQ using OpenVPN `ovpn` interface as a parent is not configured at boot time #11662
- Fixed: VLAN and QinQ edit pages allows selecting incompatible OpenVPN `tun` interfaces #11675
- Fixed: Advanced DHCP client configuration "Protocol timing" help text is in the wrong location #11926
- Added: VLAN list sorting #11968
- Fixed: Boot messages contain entries about configuring LAGG/VLAN/QinQ interfaces even when no entries of those types are configured #12002
- Fixed: Input validation incorrectly rejects a second IPv4-only GRE tunnel #12049
- Fixed: Interface assignment mismatch is not detected if VLAN-only parent interface is removed #12170
- Fixed: IPv6 DNS servers from dynamic sources are not listed on `status_interfaces.php` #12252
- Fixed: IPv6 gateway for an interface is not shown on `status_interfaces.php` if the interface does not also have an IPv4 gateway #12253
- Fixed: Remove subnet overlap check on LAN interfaces when using 6rd #12371
- Fixed: "6RD Prefix" field does not have input validation #12435
- Fixed: Trying to delete an assigned PPPoE interface fails without printing an error message #12514

L2TP

- Fixed: Kernel panic during L2TP retransmit #9058
- Fixed: FQDN L2TP server address is only resolved at boot #12072

Logging

- Fixed: Logging configuration added by a package is not removed on uninstall #11846
- Fixed: Remote log server input validation allows invalid values #12000
- Added: Disable log compression on new installations when `/var/log` is a ZFS dataset with compression enabled #12011
- Changed: Improve log settings help text for file size, compression, and retention count #12012
- Added: Create a log entry when a configuration change occurs #12118
- Fixed: Rotation settings for individual log files do not take effect after saving #12366

NTPD

- Added: Poll Interval For GPS and PPS #9439
- Added: Support for NTP Peer mode #11496
- Fixed: File overwrite in `services_ntpd_gps.php` via `gpsport` parameter #12191
- Added: Support SHA-256 hash NTP authentication #12213
- Fixed: ZFS installations without an RTC battery boot with clock at BIOS/EFI default value because they do not receive initial clock value from filesystem data #12769

Notifications

- Added: Option to suppress expiration notifications for revoked certificates #12109
- Added: Support for Slack notifications #12291
- Added: Send notification for halt, reboot, and reroot events #12441
- Fixed: `rc.notify_message` only sends notifications via SMTP #12585

OpenVPN

- Added: Support aliases in OpenVPN local/remote/tunnel network fields #2668
- Changed: Set `explicit-exit-notify` option by default for new OpenVPN server instances #11684
- Fixed: OpenVPN client certificate validation with OCSP always fails #11829
- Added: Option to validate OpenVPN peer TLS certificate key usage #11865
- Added: Log external IP address of OpenVPN clients on connect and disconnect #11935
- Fixed: DNS Resolver does not add PTR record for OpenVPN clients #11938
- Fixed: OpenVPN IPv6 tunnel network is not validated properly #11999
- Fixed: OpenVPN RADIUS-based firewall rules use incorrect port ranges #12020

- Fixed: Incorrect OpenVPN Client Export help link #12022
- Fixed: OpenVPN RADIUS-based firewall rules do not use expected value for RADIUS-assigned IP addresses #12076
- Fixed: Prevent using OpenVPN "Exit Notify" option with point-to-point modes #12102
- Fixed: OpenVPN Wizard configuration missing recently added default values #12172
- Fixed: OpenVPN does not clean up previous CA and CRL files #12192
- Changed: Move "Description" option on OpenVPN server and client pages to top of the page, show internal instance ID #12218
- Fixed: Prevent using OpenVPN "Inactive" option with point-to-point modes #12219
- Fixed: Configuration files are not deleted after disabling an OpenVPN instance #12223
- Fixed: OpenVPN page allows to delete/disable instance with an assigned interface #12224
- Fixed: OpenVPN status incorrect for TAP servers without a defined tunnel network #12232
- Fixed: OpenVPN client connect/disconnect scripts are not used in Remote Access (SSL/TLS) mode #12238
- Added: Pop-up window to view firewall rules generated from RADIUS ACL entries on the OpenVPN status page #12321
- Added: Support OpenVPN `client-kill` to terminate remote clients instead of clearing their session #12416
- Fixed: Set OpenVPN Gateway Creation value to "Both" by default for new instances #12448
- Fixed: OpenVPN form validation issues #12677

Operating System

- Changed: Ensure `/usr/local/sbin/` scripts use full path to executable files #11985
- Fixed: Update NGINX to address CVE-2021-23017 #12061
- Added: Suppress kernel messages for `lo0` configuration during boot #12094
- Changed: Convert RAM disks to `tmpfs` #12145
- Changed: Improve uses of `grep` which utilize user-supplied patterns #12265
- Fixed: Update mpd5 to address vulnerabilities in < 5.9_2 #12373
- Fixed: Update python to address vulnerabilities < 3.8.12 #12374
- Fixed: Multiple cURL Vulnerabilities #12434
- Changed: Add note in log settings that disabling logging also disables `sshguard` login protection #12511
- Fixed: Kernel panic in `nd6_dad_timer()` #12548

PHP Interpreter

- Fixed: `diag_dump_states.php` no longer filters by rule ID #12605

PPP Interfaces

- Fixed: PPP interfaces lose the description field in `ifconfig` output when restarted #11959

PPPoE Server

- Added: Option to select PPPoE Server authentication protocol #12438

Package System

- Fixed: Package <plugins> and <tabs> content missing from configuration in some cases #11290
- Added: Add `librdkafka` package to the pfSense package repository #12290
- Fixed: PHP error on `pkg_mgr_install.php` when multiple instances are running #12713
- Fixed: Potential XSS in `pkg.php` via `pkg_filter` #12725

RRD Graphs

- Added: Graph for hardware temperature readings #9297

Routing

- Fixed: Static routes using aliases are not automatically updated when alias content changes #7547
- Fixed: Input validation does not prevent removing a gateway used by a DNS server #8390
- Fixed: Kernel route table entries are removed if they match disabled static route entries #10706
- Fixed: Modifying static routes results in a logged error, changes are not reflected in routing table #11599
- Added: Require user to manually apply changes after altering static route entries #11895
- Fixed: Route data collection method on `diag_routes.php` has multiple issues #12257

Rules / NAT

- Added: IPv6 support in `easyrule` CLI script #11439
- Fixed: NAT rule overlap detection is inconsistent #11734
- Fixed: Input validation not working for 1:1 NAT entries using an alias as a destination #11923
- Fixed: `easyrule` script does not function properly #12151
- Fixed: IPv6 policy routing does not work if an IPsec tunnel phase 2 remote network is configured for `::/0` #12164
- Fixed: 1:1 NAT rule with internal IP address of "Any" results in an invalid firewall rule #12168
- Fixed: Firewall rule tabs load slowly when many rules on the tab utilize gateways #12174

- Fixed: VIP network addresses are not expanded on Port Forward rules #12233
- Fixed: Duplicating a Port Forward does not copy "Filter Rule Association" values of "None" or "Pass" #12272
- Added: Display default "Reflection Timeout" value on `system_advanced_firewall.php` #12318
- Fixed: NAT rule overlap detection does not check special networks #12361
- Fixed: Input validation prevents creating 1:1 NAT rules on OpenVPN #12408
- Fixed: 1:1 NAT edit page lists incorrect entries in the Destination field #12410
- Added: Icon for traffic direction on floating rules tab #12433
- Fixed: Port forward rules are not created for special networks (pppoe, openvpn) #12452
- Fixed: Automatic outbound NAT for reflection does not support IPv6 #12500
- Fixed: Interface group name starting with a digit creates invalid XML for rule separators #12529
- Added: Change Gateway/Group name in firewall rule list to clickable link to edit page for the entry #12555
- Fixed: Automatic rule tracker IDs incorrect after multiple filter reloads #12588
- Fixed: PHP error when clicking Delete on Outbound NAT with no rules selected #12694

SNMP

- Added: IPv6 support for base system SNMP service #12325

Services

- Fixed: System attempts to stop inactive services at shutdown #12001
- Fixed: System attempts to start inactive services at boot #12038

Traffic Shaper (ALTQ)

- Added: IPv6 support in the Traffic Shaper Wizard #4769
- Fixed: Panic when using CBQ traffic shaping #11470
- Added: Allow Chelsio T6 CXGBE (`cc`) drivers to be used for ALTQ traffic shaping #12499
- Changed: Traffic shaper wizard default bandwidth type should be Mbit/s #12501

Traffic Shaper (Limiters)

- Fixed: Unable to delete limiter referenced in filter rules #12503
- Fixed: Kernel panic when using `fq_pie` limiter scheduler #12622

UPnP/NAT-PMP

- Added: UPnP/NAT-PMP STUN configuration options #10587

Upgrade

- Changed: `pfSense-upgrade` should reinstall all packages on new version upgrades #12235

User Manager / Privileges

- Added: Copy button for group entries in the User Manager #12226

Virtual IP Addresses

- Fixed: Validation when deleting a VIP does not check if the VIP is used by IPsec phase 1 entries #12356
- Fixed: Validation when deleting a VIP does not prevent deleting a CARP VIP used as a parent for an IP Aliases VIP #12362

Wake on LAN

- Added: Wake on LAN button to wake all devices #12480

Web Interface

- Changed: Update font formats to WOFF2 #11507
- Fixed: DHCP Leases page and ARP table page fail to load if DNS is not available #11512
- Fixed: Notifications page cannot be saved without configuring or disabling SMTP #12107
- Changed: Convert help shortcut links to server-side redirects #12314
- Fixed: Help text for RAM disk settings does not mention Captive Portal data #12389
- Fixed: Input validation error can unintentionally result in removal of PPP type interface settings #12498

Wireless

- Fixed: `wpa_supplicant` uses 100% of a CPU core at boot #11453
- Fixed: Interfaces page does not show Wireless EAP client options #12239

XMLRPC

- Fixed: XMLRPC sync results in an error when a failover peer IP address is specified in DHCP server settings for an unconfigured interface #10955
- Added: XMLRPC synchronization for DHCP relay settings #11957
- Changed: XMLRPC client improvements #12051
- Fixed: Changes to an existing IPsec configuration are not applied on HA secondary after XMLRPC sync #12075

3.17 21.05.2 New Features and Changes

This is a maintenance release of pfSense® Plus software. pfSense Plus software version 21.05.2 corrects an issue with the pre-installed Netgate Firmware Upgrade package on Netgate 6100 hardware devices.

In certain circumstances the pre-installed Netgate Firmware Upgrade package could have incorrectly offered to downgrade the firmware when the hardware shipped from the factory with a newer firmware version than the copy contained within the 21.05.1 installation image.

The pfSense Plus software version number was increased for all models of Netgate hardware for consistency, but there are no functional changes for other hardware platforms. Upgrading a device in the field to 21.05.2 is not necessary at this time, but users may do so if they wish.

3.18 21.05.1 New Features and Changes

This is a maintenance release including bug fixes for issues affecting pfSense® Plus software version 21.05.

3.18.1 Security

This release includes corrections for the following vulnerabilities in pfSense software:

- Additional corrections for pfSense-SA-21_02.captiveportal (XSS in Captive Portal client login page, #11843)

3.18.2 General

3.18.3 pfSense Plus

FreeBSD

- Fixed: 32-bit ARM performance regression #12200

Operating System

- Changed: Native hardware package builds for 32-bit ARM #12201

PHP Interpreter

- Changed: Disable PCRE JIT to work around PHP PCRE crashes on multi-core 32-bit ARM systems #12004

Routing

- Fixed: Static routes may not be in routing table when expected #11986

3.19 21.05 New Features and Changes

This is a regularly scheduled software release of pfSense® Plus software including new features, additional hardware support, and bug fixes.

3.19.1 Security

This release includes corrections for the following vulnerabilities in pfSense software:

- pfSense-SA-21_02.captiveportal (XSS in Captive Portal client login page, #11843)

3.19.2 General

- Added: WireGuard add-on package
- Added: OpenVPN client import add-on package
- Fixed: `ix(4)` driver fails to attach if a broken or unsupported SFP module (e.g. incompatible media type) is present at boot time [NG 1586]
- Fixed: IP Address ranges do not work in aliases on 32-bit ARM [NG 5445]

3.19.3 pfSense Plus

Aliases / Tables

- Added: PHP shell playback script to modify Alias contents #11380

Authentication

- Added: Copy button for Authentication Server entries #11390

Backup / Restore

- Added: Randomize time of scheduled AutoConfigBackup runs #10811
- Fixed: Automated corruption recovery from cached `config.xml` backup files should check multiple backups #11748

Captive Portal

- Added: Redirect Captive Portal users to login page after they logout #11264
- Fixed: Captive Portal post-auth redirect is not properly respected #11842
- Fixed: Potential XSS vulnerability in Captive Portal `redirurl` handling #11843

Certificates

- Fixed: Certificate Manager does not report Unbound as using a certificate #11678
- Fixed: PHP error on certificate list due to unreadable private key #11859
- Fixed: Export P12 icon is missing if certificate is not locally renewable #11884

Configuration Upgrade

- Fixed: PHP error in `upgrade_212_to_213()` when upgrading certain IPsec tunnels #11801

Console Menu

- Changed: Allow reroot on ZFS from console and GUI reboot menu entries #11914

DHCP (IPv6)

- Fixed: `dhcp6withoutra_script.sh` does not get executed when advanced options are set #11883

DNS Forwarder

- Fixed: Disable DNSSEC option for dnsmasq #11781
- Fixed: Update dnsmasq to 2.85 to fix CVE-2021-3448 #11866

DNS Resolver

- Fixed: Unbound Python Integration repeatedly mounts `dev` without unmounting #11456
- Fixed: Stale hostname registration data for OpenVPN clients is not deleted from the DNS Resolver configuration at boot #11704
- Changed: Temporarily move back to Unbound 1.12.x due to instability on Unbound 1.13.x #11915

Dashboard

- Fixed: Thermal sensors widget no longer shows values from certain hardware #11787
- Fixed: IPsec Dashboard widget only displays first P2 subnet when using a single traffic selector #11893
- Fixed: Editing widgets on Dashboard causes a PHP Warning #11939

Diagnostics

- Fixed: ARP Table populates hostname values using expired DHCP lease data #11510
- Fixed: Sanitize OpenVPN Client Export certificate password in status output #11767
- Fixed: Sanitize Captive Portal RADIUS MAC secret in status output #11769
- Fixed: MAC address OEM information missing from ARP table #11819

Dynamic DNS

- Added: New Dynamic DNS Provider: Mythic-Beasts #7842
- Added: New Dynamic DNS Provider: one.com #11293
- Added: New Dynamic DNS Provider: Yandex PDD #11294
- Added: New Dynamic DNS Provider: NIC.RU #11358
- Added: New Dynamic DNS Provider: Gandi LiveDNS IPv6 #11420
- Fixed: Automatic 25-day forced Dynamic DNS update removes wildcard domain #11667
- Fixed: Digital Ocean Dynamic DNS help text is incorrect #11754
- Fixed: NoIP.com Dynamic DNS update failure is not detected properly #11815
- Fixed: Dynamic DNS edit page incorrectly hides username field when switching away from Digital Ocean #11840

Gateways

- Added: Input validation to prevent setting a load balancing gateway group as default #11164

Hardware / Drivers

- Changed: Deprecate old cryptographic accelerator hardware which is not viable on modern systems #11426
- Fixed: Using SHA1 or SHA256 with AES-NI may fail if AES-NI attempts to accelerate hashing #11524

IGMP Proxy

- Fixed: IGMP Proxy restarts unnecessarily after IPv6 gateway events #11904

IPsec

- Added: GUI option to set RADIUS Timeout for EAP-RADIUS #11211
- Added: Option to switch IPsec filtering modes to choose between `enc` and `if_ipsec` filtering #11395
- Changed: Move custom IPsec NAT-T port settings to Advanced Options #11518
- Fixed: strongSwan configuration always contains user EAP/PSK values #11564
- Added: IPsec GUI option to control Child SA `start_action` #11576
- Fixed: Error when adding both IPv4 and IPv6 P2 under an IPv4 or IPv6 only IKEv1 P1 #11651
- Fixed: Cannot disable IPsec P1 when related P2s are in VTI mode and enabled #11792
- Fixed: IPsec VTI interface names are not properly formed for more than 32 interfaces #11794
- Fixed: Applying IPsec settings for more than ~30 tunnels times out PHP #11795
- Fixed: `ipsec_vti()` does not skip disabled VTI entries #11832
- Fixed: IPsec GUI allows creating multiple identical Phase 1 entries when using FQDN for remote gateway #11912

IPv6 Router Advertisements (RADVD)

- Added: Use virtual link local IP address as RA source address for HA environments #11103
- Added: Shortcut buttons for service control and logs on RADVD configuration #11911
- Fixed: RADVD breaks on SIGHUP #11913

Interfaces

- Fixed: DHCP interfaces are always treated as having a gateway, even if one is not assigned by the upstream DHCP server #5135
- Fixed: Interfaces page displays MAC Address field for interfaces which do not support L2 #11387
- Fixed: CLI interface configuration without IPv6 leaves RA enabled #11609
- Fixed: Incomplete PPPoE custom reset values lead to invalid cron entry #11698
- Fixed: Error when changing MTU if the interface is used for both IPv4 and IPv6 default routes #11855

L2TP

- Fixed: Unused L2TP VPN files are not removed when the service is disabled #11299
- Added: GUI option to set MTU for L2TP VPN server #11406

NTPD

- Fixed: NTP widget displays incorrect status #11495
- Fixed: NTP authentication input validation rejects valid keys #11850

Notifications

- Fixed: Invalid HTML encoding in modal Notices window #11765

OpenVPN

- Added: Allow the firewall to use DNS servers provided to an OpenVPN client instance #11140
- Fixed: OpenVPN Wizard does not support gateway groups #11141
- Added: Set Explicit Exit Notify to 1 by default for new OpenVPN client instances #11521
- Added: Support for Cisco AVPair `{clientipv6}` template in firewall rules returns by RADIUS #11596
- Fixed: OpenVPN does not clean up parsed `Cisco-AVPair` rules on non-graceful disconnect #11699
- Fixed: OpenVPN does not kill IPv6 client states on disconnect #11700
- Fixed: OpenVPN client starts when CARP VIP is in BACKUP status when bound to Virtual IP aliased to CARP VIP #11793
- Fixed: Certificate validation with OCSP always fails in `openvpn.tls-verify.php` #11830
- Changed: Update OpenVPN to 2.5.2 #11844
- Fixed: OpenVPN client startup error if IPv6 Tunnel Network is defined in TAP mode #11869

Operating System

- Added: Kernel modules for alternate congestion control algorithms #7092
- Added: Kernel module for RTL8153 driver #11125
- Added: Xen console support #11402
- Fixed: Unquoted variable in `dot.tcshrc` can cause proxy password to be printed #11867

Routing

- Fixed: Static route targets may still reachable via default route when the gateway they should route through is down #11296
- Fixed: IPv4 link-local (`169.254.x.x`) gateway does not function #11806

Rules / NAT

- Added: Support for IPv6 firewall entries with dynamic delegated prefix and static host address #6626
- Fixed: Disabling all interfaces associated with a floating rule causes the firewall to generate an incorrect pf rule #11688
- Fixed: Input validation prevents creating 1:1 NAT rules on IPsec #11751
- Fixed: Invalid combinations of TCP flag matching options cause `pfctl` parser error #11762
- Fixed: Error loading rules in certain cases where an interface is temporarily without an address #11861

Traffic Shaper (ALTQ)

- Fixed: Harmless error when enabling traffic shaper #11229
- Fixed: Segmentation fault when loading ALTQ traffic shaping rules using FAIRQ #11550

Traffic Shaper (Limiters)

- Fixed: Unused Limiter entries with schedules create unnecessary cron jobs #11636
- Fixed: Error when setting queue limit on CODELQ limiter #11725

Upgrade

- Fixed: Language presented to user during upgrade is misleading #11897

Web Interface

- Added: Replace HTTP links with HTTPS in the GUI #11228
- Fixed: Ambiguous text in help and input validation error for system domain name #11658
- Fixed: PHP error if `PHP_error.log` file is too large #11685
- Fixed: RAM Disk Settings shows Kernel Memory at `0` Kb and does not allow the user to create RAM disks #11702
- Fixed: HTTP Referer error message text is incorrect #11873
- Fixed: Missing `/0` subnet when cloning repeatable CIDR mask controls #11880

WireGuard

- Fixed: Ignore WireGuard configurations under `<installedpackages></installedpackages>` #11808

Wireless

- Added: GUI options for WPA Enterprise with identity/password #2400

XMLRPC

- Fixed: XMLRPC synchronization restarts all OpenVPN instances on the secondary node when making any change on the primary node #11082
- Fixed: XMLRPC Client does not honor its default timeout value #11718

3.20 21.02.2/2.5.1 New Features and Changes

pfSense® Plus software version 21.02.2 and pfSense CE software version 2.5.1 are maintenance releases to address recently identified issues.

Warning: WireGuard was removed from the base system in releases after pfSense Plus 21.02-p1 and pfSense CE 2.5.0, when it was removed from FreeBSD.

If upgrading from a version that has WireGuard active, the upgrade will abort until all WireGuard tunnels are removed. For more details, see the *Release Notes*

WireGuard is available as an add-on package on pfSense Plus 21.05, pfSense CE 2.5.2, and later versions. The settings for the WireGuard add-on package are not compatible with the older base system configuration.

Note: The WireGuard package is still under active development. Follow the development progress on the developer's YouTube channel

Tip: To remove WireGuard tunnels, navigate to **VPN > WireGuard** and click the delete button for each tunnel. When the page displays **No WireGuard tunnels have been configured.**, the upgrade can proceed.

Note: This pfSense Plus software version contains all of the items noted below for pfSense CE as well.

Tip: For those who have not yet updated to 2.4.5-p1, consult the *previous release notes* and blog posts for those releases to read all important information and warnings before proceeding.

3.20.1 Known Issues / Errata

- There is an issue in this release with port forwarding on pfSense CE software installations with multiple WANs, see #11805 for details.
- There is an issue with AES-NI hash acceleration for SHA1 and SHA-256. If the AES-NI driver detects a system capable of accelerating SHA1 or SHA-256 and the firewall attempts to utilize one of those hashes, the affected operation may fail. This affects IPsec and OpenVPN, among other uses. pfSense Plus users can change to QAT acceleration on supported hardware instead. In cases where QAT is unavailable, change to AES-GCM, change to a different unaccelerated hash (e.g. SHA-512), or disable AES-NI. See #11524 for details.
- There is a similar issue which affects SafeXcel SHA1 and SHA2 hash acceleration on SG-1100 and SG-2100. On that hardware, change to an AEAD cipher such as AES-GCM or switch to an unaccelerated hash. This issue is being tracked internally on NG #6005
- The FRR package on pfSense Plus 21.02 and pfSense CE 2.5.0 and later no longer exchanges routes with BGP peers by default without being explicitly allowed to do so. This is more secure behavior but requires a manual change. To replicate the previous behavior, use **ONE** of the following workarounds:
  - Navigate to **Services > FRR BGP** on the **Advanced** tab and check *Disable eBGP Require Policy*, then **Save**.
  - Instead of disabling the policy check, create route maps which match and allow expected incoming and outgoing routes explicitly. This is the most secure method. See *Peer Filtering* and *BGP Example Configuration* for more information.
  - Manually create a route map to permit all routes (Name: `allow-all`, Action: *Permit*, Sequence: `100`), then set that route map on BGP neighbors for inbound and outbound peer filtering. This can be used as a placeholder for later migration to more secure route map filtering.

3.20.2 Security

This release includes corrections for the following vulnerabilities in pfSense software:

- pfSense-SA-21_01.webgui (XSS in Wake on LAN, #11616)

3.20.3 General

3.20.4 pfSense Plus

Certificates

- Fixed: CA and certificate validity end dates after 2038 are not handled properly on 32-bit ARM #11504

Interfaces

- Added: Interface Status page information for switch uplinks may be replaced by switch port data when media state monitoring is set #10804

Rules / NAT

- Fixed: State matching problem with reponses to packets arriving on non-default WANs #11436

Upgrade

- Fixed: LEDs do not indicate available upgrade status #11689

3.20.5 pfSense CE

Aliases / Tables

- Fixed: Alias name change is not reflected in firewall rules #11568

Authentication

- Fixed: Unreachable LDAP server for SSH auth causes boot process to stop at 'Synchronizing user settings' and no user can login over SSH #11644

Certificates

- Fixed: Invalid certificate data can cause a PHP error #11489
- Fixed: Renewing a self-signed CA or certificate does not update the serial number #11514
- Fixed: Unable to renew a certificate without a SAN #11652
- Fixed: Certificates with escaped x509 characters display the escaped version when renewing #11654
- Fixed: Creating a certificate while creating a user does not fully configure the certificate properly #11705
- Fixed: Renewing a certificate without a `type` value assumes a server certificate #11706

DNS Resolver

- Fixed: DNS Resolver does not add a `local-zone` type for `ip6.arpa` domain override #11403
- Fixed: DNS Resolver does not bind to an interface when it recovers from a down state #11547

Dashboard

- Fixed: CPU details are incorrect in the System Information widget after resetting log files #11428
- Fixed: Disabling 'State Table Size' in the System Information widget prevents other data from being displayed #11443

Gateway Monitoring

- Fixed: Automatic default gateway mode does not select expected entries #11729

Gateways

- Fixed: Gateways with "Use non-local gateway" set are not added to routing table #11433

IPsec

- Fixed: IPsec status incorrect for entries using expanded IKE connection numbers #11435
- Fixed: Distinguished Name (FQDN) IPsec peer identifier type is not formatted properly in `swanctl.conf` secrets #11442
- Fixed: Mobile IPsec DNS server input validation does not reject unsupported IPv4-mapped IPv6 addresses #11446
- Fixed: Broken help link on IPsec Advanced Settings tab #11474
- Fixed: Connect and disconnect buttons on the IPsec status page do not work for all tunnels #11486
- Fixed: IPsec tunnels using expanded IKE connection numbers do not have proper child SA names in `swanctl.conf` #11487
- Fixed: IPsec tunnel definitions have `pools =` entry in `swanctl.conf` with no value #11488
- Fixed: Mobile IPsec broken when using strict certificate revocation list checking #11526
- Fixed: IPsec VTI tunnel between IPv6 peers may not configure correctly #11537
- Fixed: IPsec peer ID of "Any" does not generate a proper remote definition or related secrets #11555
- Fixed: IPsec tunnel does not function when configured on a 6RD interface #11643

IPv6 Router Advertisements (RADVD)

- Fixed: IPv6 RA RDNSS lifetime is too short, not compliant with RFC 8106 #11105

Installer

- Fixed: Installer does not add required module to loader.conf when using ZFS #11483

Interfaces

- Fixed: IPv4 MSS value is incorrectly applied to IPv6 packets #11409
- Fixed: Gateway value for DHCP6 interfaces missing after RA events triggered script without gateway information #11454
- Fixed: Delayed packet transmission in cxgbe driver can lead to latency and reduced performance #11602
- Fixed: DHCP6 interfaces are reconfigured multiple times at boot when more than one interface is set to Track #11633

Logging

- Fixed: Entries from rotated log files may be displayed out of order when log display includes contents from multiple files #11639

Notifications

- Fixed: Telegram and Pushover notification API calls do not respect proxy configuration #11476

OpenVPN

- Fixed: OpenVPN authentication and certificate validation fail due to size of data passed through `fcgicli` #4521
- Added: Display negotiated data encryption algorithm in OpenVPN connection status #7077
- Fixed: OpenVPN does not start with several authentication sources selected #11104
- Fixed: OpenVPN client configuration page displays Shared Key option when set for SSL/TLS #11382
- Fixed: Incorrect order of `route-nopull` option in OpenVPN client-specific override configuration #11448
- Fixed: OpenVPN using the wrong OpenSSL command to list digest algorithms #11500
- Fixed: Selected Data Encryption Algorithms list items reset when an input validation error occurs #11554
- Fixed: OpenVPN does not start with a long list of Data Encryption Algorithms #11559
- Fixed: ACLs generated from RADIUS reply attributes do not parse `{clientip}` macro #11561
- Fixed: ACLs generated from RADIUS reply attributes have incorrect syntax #11569
- Fixed: OpenVPN binds to all interfaces when configured on a 6RD interface #11674

Operating System

- Fixed: Unexpected Operator error on console at boot with ZFS and RAM Disks #11617
- Changed: Upgrade OpenSSL to 1.1.1k #11755

Routing

- Fixed: Disabled static route entries trigger 'route delete' error at boot #3709
- Fixed: Route tables with many entries can lead to PHP errors and timeouts when looking up routes #11475
- Fixed: Error when removing automatic DNS server route #11578
- Fixed: IPv6 routes with a prefix length of 128 result in an invalid route table entry #11594
- Fixed: Error when deleting IPv6 link-local routes #11713

Rules / NAT

- Fixed: Saved state timeout values not loaded into GUI fields on system_advanced_firewall.php #11565
- Fixed: Firewall rule schedule cannot be changed #11747

Upgrade

- Fixed: pfSense Proxy Authentication not working #11383

Wake on LAN

- Fixed: Potential stored XSS vulnerability in services_wol.php #11616

Web Interface

- Fixed: Requests to `ews.netgate.com` do not honor proxy configuration #11464

XMLRPC

- Fixed: XMLRPC error with Captive Portal and CARP failover when GUI is on non-standard port #11425
- Fixed: Incorrect DHCP failover IP address configured on peer after XMLRPC sync #11519
- Fixed: PHP error in logs from XMLRPC if no sections are selected to sync #11638

3.21 2.7.0 New Features and Changes

This pfSense® CE software release includes new features and bug fixes.

3.21.1 Upgrade Notes

> **Warning:** Due to major changes in PHP and base OS versions, there is a higher than usual chance that packages will interfere with the upgrade process.
>
> To give an upgrade the best possible chance of going smoothly, uninstall **all** packages before starting the upgrade.

3.21.2 General

- PHP has been upgraded from 7.4.x to 8.2.6
- The base operating system has been upgraded to FreeBSD 14-CURRENT

> **Warning:** As a part of the FreeBSD upgrade this version removes several deprecated IPsec algorithms:
>
> – 3DES Encryption
>
> – Blowfish Encryption

  - CAST 128 Encryption
  - MD5 HMAC Authentication

The best practice is to reconfigure tunnels using better encryption and test them before performing an upgrade to ensure a smoother transition.

On upgrade, IPsec tunnels will be adjusted to remove any deprecated algorithms from their configuration. The upgrade process will disable tunnels if they have no valid encryption or authentication options remaining. The upgrade process will notify the user of any changes it makes.

This change only affects IPsec and not other uses of these algorithms. For example, BGP can still use TCP-MD5 authentication.

- Added support for ChaCha20-Poly1305 encryption with IPsec
- Captive Portal has been migrated from IPFW to PF
- A long-standing difficult-to-reproduce crash in Unbound during reloading has been addressed. Christian McDonald tracked down the source of the Unbound SIGHUP crashes to a reference counting bug within the MaxMindDB Python module. Both a patch to MaxMind and a port revision to FreeBSD ports were submitted and accepted, and the fix is included in the 2.7.0 release. It is now safe again to enable DHCP registration alongside Unbound Python mode in pfBlockerNG.
- In addition to the Unbound crash, Christian also identified a memory leak with DHCP registration and Unbound Python mode (#10624). This is largely mitigated by updates to Python and related libraries, but there is additional ongoing work to resolve it further for future release.
- Fix for UPnP and multiple game systems
- New gateway state killing options for smoother failover
- Firewall/NAT rule usability improvements such as buttons to toggle multiple rules and copy rules to other interfaces
- OpenVPN upgraded to 2.6.4
- OpenVPN Shared Key Tunnels Deprecated – They still work, but will trigger warnings in the logs and GUI.
- New Packet Capture GUI
- UDP Broadcast Relay Package

Danger: This version includes newer ZFS features which may not be compatible with older boot loaders. These features **are not** enabled by default when upgrading to avoid potential problems with older boot loaders. Some ZFS commands run at the CLI, such as `zpool status`, may report that a pool can be upgraded, but doing so may also require manually updating the boot loader for the device to boot properly. Upgrading the ZFS pool **is not** necessary at this time. As such, the best practice is to leave it as-is. This will be handled automatically as needed in future updates.

Reinstalling the OS from current installation media will result in having the most recent boot loader and ZFS feature set.

3.21.3 Security

pfSense CE 2.7.0-RELEASE includes fixes for the following potential vulnerabilities:

- pfSense-SA-22_05.webgui: A potential XSS vulnerability in firewall_aliases.php from URL table alias URLs.
- pfSense-SA-23_01.webgui: A potential XSS vulnerability in diag_edit.php from browsing directories containing specially crafted filenames on the filesystem.
- pfSense-SA-23_02.webgui: A potential XSS vulnerability in system_camanager.php and system_certmanager.php from specially crafted descriptions when editing entries.
- pfSense-SA-23_03.webgui: A potential authenticated arbitrary file creation vulnerability from the name parameter when creating or editing URL table aliases.
- pfSense-SA-23_04.webgui: A potential authenticated arbitrary command execution vulnerability in status.php from specially crafted filenames on the filesystem.
- pfSense-SA-23_05.sshguard: Anti-brute force protection bypass for GUI authentication requests containing certain proxy headers.
- pfSense-SA-23_06.webgui A potential Authenticated Command Execution vulnerability from the `bridgeif` parameter on `interfaces_bridge_edit.php` in the GUI.

3.21.4 pfSense CE

Changes in this version of pfSense CE software.

Aliases / Tables

- Fixed: Alias content is sometimes incomplete when an alias contains both FQDN and IP address entries #9296
- Fixed: Alias with non-resolving FQDN entry breaks underlying PF table #12708
- Fixed: Renaming an alias does not update the alias names in static routes and OpenVPN instances #12727
- Added: Retain descriptions when exporting and importing aliases #12842
- Fixed: Potential XSS from URL and URL Table alias URLs #13060
- Fixed: Alias content is sometimes incomplete if the firewall cannot resolve an FQDN in the alias #13282
- Added: Specify CA trust store location when downloading and validating URL alias content #13367
- Fixed: Invalid alias name can still be used by code attempting to validate URL table content #13425
- Fixed: Deleting an alias marks the subsystem as unclean but also unconditionally reloads the filter configuration #13538
- Fixed: Missing descriptions for referrers to firewall aliases cause empty strings for references to be returned when deleting an in-use alias #13539
- Fixed: Using PF reserved keywords for interface descriptions results in an invalid ruleset #14007
- Fixed: Alias list is not sorted #14015

Authentication

- Fixed: User password hashes pseudo-random number generator may return insecure salt value #12801
- Added: GUI option to select the user password hashing algorithm #12855
- Fixed: LDAP setup does not display 'Global Root CA List' option unless another CA also exists #13185
- Fixed: Unable to set web interface session timeout to `0` (i.e. never expire) #13561
- Fixed: Extra remote address information can confuse `sshguard` #13574
- Changed: Improve LDAP debugging #13718
- Added: Option to enable/disable console bell, enabled by default #14002

Auto Configuration Backup

- Added: Option to list AutoConfigBackup entries in "reverse" order (newest at top) #11266
- Added: Support for international characters in the AutoConfigBackup Hint/Identifier field #13388
- Fixed: Auto Config Backup prints a confusing decryption error when using the wrong key #14060

Backup / Restore

- Changed: Comply with current iteration standards when encrypting and decrypting configuration files #12556
- Added: Support encrypted `config.xml` files when restoring via ECL #12685
- Added: Notify user if AutoConfigBackup is unable to successfully upload a backup #12724
- Added: Ability to sort AutoConfigBackup entries #12773
- Fixed: Sanitize SHA-512 user password hashes in `status.php` output #12810
- Added: Option to restore dashboard widget layout #13125
- Fixed: PHP error restoring DHCP lease data on fresh installation: #13157
- Fixed: Attempting to restore a 0 byte `config.xml` prints an error that the file cannot be read #13289
- Fixed: Configuration history restores revision no matter which option is clicked in confirmation dialog #13861
- Fixed: RRD restore process does not sanitize filenames from backup XML #13935

Build / Release

- Changed: Disable `pkg` compatibility flag which creates `txz` file extension symbolic links #12782

CARP

- Fixed: CARP VIPs can become master too early at boot time #2218
- Changed: Reorganize CARP status page #12701
- Fixed: CARP event storm when leaving persistent CARP maintenance mode #12961

Captive Portal

- Fixed: Allowed IP/Hostname "Direction" option is never used #12649
- Fixed: `nginx` logs an error that the port is already in use when restarting Captive Portal services #12651
- Fixed: Value of `net.inet.ip.dummynet.*` OIDs in `sysctl` are ignored #12733
- Fixed: Only TCP traffic is passed outbound through IPFW #12834
- Changed: Transition Captive Portal from IPFW to PF #13100
- Fixed: Voucher CSV output has leading space before voucher code #13272
- Fixed: Captive Portal breaks policy based routing for MAC address bypass clients #13323
- Fixed: Multiple Captive Portal interfaces do not properly form the list of portal IP addresses #13391
- Fixed: Custom logo or background image is created with two dots (`..`) before the file extension #13396
- Fixed: Captive Portal does not keep track of client data usage #13418
- Fixed: All Captive Portal users are given the same limiter pipe pair #13488
- Fixed: Captive Portal RADIUS start/stop accounting does not reset counters at each accounting start #13838
- Fixed: Captive Portal does not apply RADIUS bandwidth limits to user pipes #13853

Certificates

- Fixed: CA path is not defined when using `curl` in the shell #12737
- Added: Option to retain the existing serial number when renewing a CA or certificate #13010
- Fixed: Exporting a PKCS#12 file from the certificate manager does not use the intended encryption algorithm #13257
- Fixed: Input validation is not rejecting invalid description characters when editing a CA or Certificate #13387
- Fixed: CRL expiration date with default lifetime is too long, goes past UTCTime limit #13424
- Fixed: ECDSA certificate renewal causes digest algorithm to be reset to SHA1 #13437
- Fixed: Some blank SAN fields are not ignored when creating a certificate #14124
- Added: Ability to edit Certificate Revocation List properties #14185
- Changed: Add note to inform the user that the "Next Certificate Serial" value is ignored when the "Randomize Serial" option is enabled #14188

Configuration Backend

- Added: Move command line history to a GUI option stored in `config.xml` rather than a manual flag file #12675
- Added: Eliminate duplicate shell commands from history file #12741
- Fixed: Input validation is checking RAM disk sizes when they are inactive #13479

Configuration Upgrade

- Added: Playback script to perform a configuration upgrade on an arbitrary `config.xml` file #12973
- Fixed: PHP Error in `upgrade216_ipsec_create_vtimap()` #14400

Console Menu

- Fixed: Changing an interface IP address and gateway at the console does not save the new gateway if one already exists for the interface #12632
- Added: Warn the user if they attempt to disable SSH from the menu while connected through SSH #13103
- Fixed: Hidden menu option `100` incorrectly handles HTTPS detection #13258

DHCP (IPv4)

- Added: Improve distinction between online and idle/offline entries in DHCP lease list #10345
- Fixed: Disabling DHCP Server RRD statistics does not work #12710
- Fixed: `HTTPClient` option not sent when using UEFI HTTP Boot #12892
- Fixed: `HTTPClient` option does not work for static mappings #12896
- Fixed: DHCP "Ignore denied clients" option with MAC Deny list set causes DHCP server to not start #12923
- Added: Relax DHCP maximum lease time input validation #13118
- Fixed: DHCP lease list displays wrong interface name in the "Leases in Use" summary if DHCP settings for a disabled interface remain in the configuration #13127
- Changed: Clean up DHCP Server option language #13250
- Fixed: DHCP Server generates an invalid configuration for static mappings when defining network booting and UEFI HTTPBoot URL #13573
- Added: Input validation for numbered DHCP options in static mappings #13584
- Fixed: DHCP Server page does not properly select a default interface tab if neither WAN nor LAN are capable of being DHCP servers #14115

DHCP (IPv6)

- Fixed: Multiple DHCP6 WAN connections leads to multiple dhcp6c clients #6880
- Fixed: DHCPv6 server does not skip interfaces configured with invalid ranges #12527
- Fixed: RADVD can be started on both HA nodes when configured with an IPv6 link-local address #12582
- Fixed: Uninitialized array in `array_remove_duplicates()` #12749
- Fixed: Advanced DHCP6 client settings only work for a single interface #13462
- Fixed: "Provide DNS servers to DHCPv6 clients" setting does not reflect a changed value until the page is reloaded #13594
- Fixed: DHCPv6 rules are not created for interfaces with static IPv6 #13633

DNS Forwarder

- Fixed: DNS Forwarder refuses valid retries from clients in certain cases #12901
- Fixed: DNS Forwarder creates a loop when "Use local DNS, ignore remote DNS servers" is selected #12902
- Fixed: DNS Forwarder custom options may fail after save/restore when options are only separated by newline #13105
- Fixed: DNS Forwarder (`dnsmasq`) is using an invalid combination of options when "Query DNS servers sequentially" is enabled #13655

DNS Resolver

- Fixed: Memory leak in Unbound with Python module and DHCP lease registration active #10624
- Fixed: Unbound crashes with signal 11 when reloading #11316
- Fixed: DNS Resolver is restarted during every `rc.newwanip` event even for interfaces not used in the resolver #12612
- Fixed: DNS Resolver does not restart during link up/down events on a static IP address interface #12613
- Added: Automatically create DNS Resolver ACLs for OpenVPN CSO entries #12636
- Fixed: DNS Resolver help text for **System Domain Local Zone Type** option refers users to `unbound.conf(5)` man page instead of pfSense docs #12781
- Fixed: DNS Resolver updates trust anchor at boot even with DNSSEC disabled which can lead to a startup delay of ~2 minutes if the firewall does not have Internet access #12985
- Fixed: DNS Resolver ACLs are not updated when OpenVPN networks change #12991
- Added: DNS Resolver option to keep probing when servers are down #13023
- Fixed: DNS resolver does not update its configuration or reload during link down events #13254
- Fixed: DNS Resolver responds with unexpected source address when the DNS over TLS server function is enabled #13393
- Fixed: Incorrect word in "Network Interfaces" help text on `services_unbound.php` #13453
- Fixed: DNS Resolver does not generate automatic ACLs for IPv6 when Network Interfaces is set to "All" #13851
- Changed: Update Unbound to use Python 3.11 instead of Python 3.9 #13867
- Changed: Update Unbound to 1.17.1 #13893

- Fixed: DNS Resolver experiences intermittent resolution failures with SSL over TLS due to ASLR #14056
- Fixed: Setting system DNS servers can incorrectly modify routes for interface addresses #14288
- Fixed: Discrepancy in "TTL for Host Cache Entries" Description #14358

Dashboard

- Fixed: Firewall log widget action icon features stop working when new log entries are added dynamically #6253
- Added: Show `Inactive` for Hardware Crypto output instead of empty field on System Information dashboard widget when nothing can be accelerated #12714
- Fixed: Uptime displays plural seconds for multiple minutes in the System Information Dashboard widget #14176
- Added: Support for Intel PCH temperature values in thermal sensors #14255

Diagnostics

- Fixed: `diag_pftop.php` does not fully encode output #12915
- Fixed: File browser on `diag_edit.php` does not encode filenames before display #13262
- Fixed: Neighbor hostnames in the NDP Table on `diag_ndp.php` are always empty #13318
- Fixed: `status.php` uses `<name>` component of `/tmp/rules.packages.<name>` filenames in shell command without encoding #13426
- Changed: Add multicast group membership (`ifmcstat`) to `status.php` #13731
- Changed: Add more disk information to status output #14103

Dynamic DNS

- Fixed: Dynamic DNS custom IPv6 service fails on 6rd tunnels #12590
- Fixed: GleSYS Dynamic DNS responses are not parsed properly #12672
- Added: IPv6 support for DNSimple Dynamic DNS #12744
- Fixed: Input validation prevents configuring wildcard Dynamic DNS records on GoDaddy #12750
- Added: Support wildcard Dynamic DNS records on DigitalOcean #12752
- Fixed: Google Domains Dynamic DNS responses are not parsed properly #12754
- Fixed: Input validation prevents configuring wildcard Dynamic DNS records on Google Domains #12761
- Fixed: Namecheap Dynamic DNS responses are not parsed properly #12816
- Fixed: Clicking Save & Force Update on a Dynamic DNS entry results in a GUI timeout #12870
- Fixed: DigitalOcean Dynamic DNS update fails with a "bad request" error #13167
- Fixed: Dynv6 Dynamic DNS client does not check the response code when updating #13298
- Fixed: DNSExit Dynamic DNS updates no longer work #13303
- Changed: Improve DynDNS help text readability #14186

FilterDNS

- Fixed: Resolve interval for `filterdns` may not match the configured value #13067

FreeBSD

- Fixed: Cannot set EFI console as primary console when using both EFI and Serial #13080
- Fixed: CVE-2022-23093 / FreeBSD-SA-22:15.ping #13716
- Changed: Update Time Zone data to 2023c or later #14209

Gateway Monitoring

- Fixed: Gateway monitoring should mark gateway as "offline" on PPPoE parent interface disconnect #12633
- Added: Option to disable auto-addition of static routes for `dpinger` #12687
- Changed: Update `dpinger` to 3.2 #12881
- Fixed: Marking a gateway as down does not affect IPsec entries using gateway groups #13076
- Fixed: Incorrect function parameters for `get_dpinger_status()` call in `gwlb.inc` #13295

Gateways

- Fixed: `fixup_default_gateway()` should not remove a default gateway managed by a dynamic routing daemon #11692
- Fixed: IPv6 link local gateway default status not indicated in GUI #11764
- Fixed: IPv6 gateway group using link local addresses incorrectly logs a gateway change because it not including interface scope properly #12721
- Added: Retain knowledge of previous dynamic gateway IP address when interface is down #12931
- Fixed: Recovering interface gateway may not be added back into gateway groups and rules when expected #13228
- Fixed: Gateway popup in firewall rule list does not indicate current gateway status #14327

Hardware / Drivers

- Added: Chelsio TOE support using the `t4_tom` module #9091
- Fixed: Intel e1000 driver (`em`, `igb`) cannot pass packets tagged with VLAN `0` #12821
- Fixed: Hyper-V RSC support in `hn(4)` driver is enabled by default and results in very low throughput #12873
- Fixed: Malicious Driver Detection event on `ixl(4)` driver #13003
- Fixed: UDP checksum errors with `ixgbe` interfaces #13883

High Availability

- Added: Use consistent pf host ID and add GUI option to set a custom host ID in state synchronization settings #12702

IGMP Proxy

- Fixed: IGMP Proxy server is restarted during every `rc.newwanip` event #12609

IPsec

- Added: Option to choose default tab in IPsec status Dashboard widget #2456
- Fixed: IPsec VTI phase 2 traffic selectors default to address when defined as a network #11226
- Fixed: `filterdns` does not monitor remote IPsec gateways for IPv6 address changes #12645
- Fixed: Disallow remote gateway of `0.0.0.0` for VTI mode #12723
- Fixed: VTI gateway status stuck as "pending" after reboot #12763
- Fixed: ESP description in IPsec phase 2 proposal help text is ambiguous #12953
- Fixed: IKEv2 Mobile IPsec clients do not receive `INTERNAL_DNS_DOMAIN` (value `25`) attribute #12975
- Fixed: Deadlock in Charon VICI interface #13014
- Added: GUI option for IPsec `dns-interval` setting #13057
- Fixed: Delete function for IPsec SAD entries on `status_ipsec_sad.php` does not work #13071
- Fixed: Mobile IPsec clients cannot be manually disconnected from IPsec status screen #13131
- Fixed: IPsec rejects certificates if any SAN is wildcard rather than rejecting when **all** SANs are wildcard #13373
- Changed: Information box on `status_ipsec.php` says "IPsec not enabled" even when a tunnel is established #13398
- Fixed: Incorrect quoting of Split DNS attribute value in `strongswan.conf` #13579
- Added: Support for ChaCha20-Poly1305 encryption with IPsec #13647
- Changed: Remove deprecated IPsec algorithms (3DES, Blowfish, and CAST 128 encryption; MD5 HMAC/Hashing) #13648
- Fixed: Reassembled packets received on a VTI are not forwarded #14396

Installer

- Fixed: Support encrypted `config.xml` files when restoring during install #12691
- Added: Recover existing SSH keys during installation #12809

Interfaces

- Added: Show SFP module details on `status_interfaces.php` #8861
- Added: Improved support for USB interfaces that may not always be present #9393
- Fixed: Primary interface address is not always used when VIPs are present #11545
- Fixed: PPPoE WAN IP address different than expected when set static by ISP #11629
- Added: Support for VLAN `0` #12070
- Fixed: `devd` is not configured to act on USB interface attach/detach events #12606
- Changed: Restart services on interface changes #12619
- Fixed: Interface status "Total Interrupts" display is non-functional #12735
- Fixed: L2TP/PPTP interface assignment page loses some values after input validation error #12780
- Fixed: Link-Local IPv6 address on WAN with MAC spoofing changes if there is an IP Alias on WAN #12790
- Fixed: Link-local address does not reset after removing MAC address spoofing #12794
- Fixed: Disabled Captive Portal configuration prevents adding an interface to a bridge #12866
- Fixed: The ruleset is not regenerated after assigning an interface #12949
- Fixed: Bridges with QinQ interfaces not properly set up at boot #13225
- Changed: Start `rtsold` immediately after `dhcp6c` sends a request #13492
- Fixed: Several advanced DHCP6 client options do not inform the user when rejecting invalid input #13493
- Changed: Clean up obsolete code in `pfSense-dhclient-script` #13501
- Fixed: DHCP client can fail permanently if an interface is down at boot #13671
- Fixed: Code that sets IPv6 MTU can unintentionally act on IPv4 addresses #13675
- Changed: Trim blank characters from static IP address fields on the Interface configuration page #13959
- Fixed: Bridge interface is not properly validated when submitted on `interfaces_bridge_edit.php` #14052

L2TP

- Fixed: L2TP MPD configuration is not updated when a dynamic WAN IP address changes #13066
- Fixed: L2TP stays bound to previous IP address after static IP address change #13082
- Fixed: Static routes to destinations at L2TP clients are not re-added after a client reconnects #13099

LAGG Interfaces

- Added: GUI option to configure layers for LACP hash #12819

Logging

- Added: Option to control log level of authentication messages in system logs (“Emergency” vs “Notice” level) #12464

Notifications

- Fixed: Slack notification options only allow - as a special character in channel names #13083
- Fixed: Identical SMTP notifications repeat in an infinite loop under certain conditions #14031
- Fixed: Notices incorrectly set system LEDs on hardware with less than three LEDs #14482

OpenVPN

- Fixed: OpenVPN IPv4 Tunnel Network incorrectly allows hostnames #11416
- Fixed: OpenVPN stays bound to previous IP address after interface changes #11864
- Added: OpenVPN option to limit concurrent connections per user #12267
- Fixed: OpenVPN does not clear old Cisco-AVPair anchor rules in some cases #12332
- Added: Use deferred client connections in OpenVPN #12407
- Fixed: OpenVPN re-synchronization also synchronizes override entries unnecessarily in some cases #12628
- Fixed: Automatic filter reload with OpenVPN client gateway uplink happens too soon or not at all #12771
- Fixed: PHP error when terminating OpenVPN sessions via the dashboard widget #12817
- Fixed: OpenVPN status display for TAP mode services shows peer-to-peer instead of client list in certain cases #12884
- Fixed: GUI does not reject an invalid OpenVPN tap mode configuration with an empty tunnel network “Bridge DHCP” disabled #12887
- Fixed: FQDN in network alias is omitted from OpenVPN networks list #12925
- Changed: Warn about OpenVPN shared key deprecation #12981
- Fixed: OpenVPN `remote_cert_tls` option does not behave correctly when enabled and later disabled #13056
- Fixed: Gateway events for IPv6 affect IPv4 OpenVPN instances and vice versa #13061
- Fixed: OpenVPN Client Overrides: properly hide/show form fields #13088
- Fixed: OpenVPN client `tls-client/client` configuration directive not handled properly #13116
- Changed: OpenVPN status page improvements #13129
- Fixed: OpenVPN `client-connect` file contains `topology` #13133
- Fixed: Per-user `route` files are not removed from `/tmp` when they are no longer needed #13145
- Fixed: OpenVPN status for multi-user VPN shows info icon to display RADIUS rules when there are none to display #13243
- Fixed: OpenVPN override IPv4 tunnel network field changing value improperly #13274
- Changed: Update OpenVPN Wizard to match current certificate and OpenVPN options #14183
- Changed: Remove deprecated NCP enable/disable toggle from OpenVPN #14201

Operating System

- Fixed: pf `hostid` value is handled inconsistently #12703
- Fixed: Some `sysctl` OIDs in `loader.conf.local` are silently removed #12862
- Fixed: Output from `pfctl -vvsr` does not include `ridentifier` value in the expected location #12868
- Changed: Update memory graphs to account for changes in memory reporting #14011
- Fixed: Netlink debug messages from IPsec #14370
- Added: wpa_supplicant: add VLAN 0 support #14457

PHP Interpreter

- Added: Upgrade PHP from 7.4 to 8.1 #13446
- Fixed: `fcgicli` fails to write packets with `nvpair` values that exceed `128` bytes #13638
- Changed: Update PHP to 8.2.6 #14027

PPP Interfaces

- Fixed: Services are not restarted when PPP interfaces connect #12811
- Fixed: PPPoE WANs fail to reconnect after parameter negotiation failure #13092
- Fixed: PPP interface custom reset date/time Hour and Minute fields do not properly handle `0` value #13307
- Fixed: IPv6 does not work on secondary PPPoE WAN #13939

PPPoE Server

- Fixed: PPPoE server panics with multiple client connections #13210

Package System

- Fixed: Packages are not automatically reinstalled when restoring configuration using the installer #12105
- Fixed: Packages with custom `internal_name` values do not reinstall properly when restoring a backup #12766
- Fixed: `write_rcfile()` does not create `rc_restart()` entry #13004
- Added: Package plugin hook for web server configuration stanzas #13054

Packet Capture

- Added: Button to clear previous packet capture data #12968
- Added: Packet Capture GUI with granular control #13382

Routing

- Added: Enable `ROUTE_MPATH` multipath routing #9544
- Fixed: Setting a default gateway of "None" does not remove the default gateway from the routing table #12536
- Fixed: Cannot remove IPv6 static routes #12728
- Fixed: Explicit PPPoE disconnect of a WAN Gateway Group member may not restore a default route #13048

Rules / NAT

- Added: Toggle button to disable/enable multiple firewall rules #2505
- Added: Port forward NAT rules with "any" protocol #4259
- Added: Allow NPt to use dynamic IPv6 networks #4881
- Added: Button to copy rules from one interface to another #8365
- Fixed: Rule separator positions change when deleting multiple rules #9887
- Fixed: Automatic Outbound NAT mode can create incorrect rules in some cases #11984
- Added: Utilize new `pfctl` abilities to kill states #12092
- Fixed: NAT reflection does not work for IPv6 port forwarding rules when configured for NAT+Proxy mode #12319
- Added: Allow the selection of "any" interface in floating rules #12392
- Fixed: Applying firewall rule changes does not clear dirty flag for aliases subsystem #12678
- Fixed: Automatic Outbound NAT rules do not include OpenVPN CSO entries #12792
- Fixed: Error loading ruleset due to illegal TOS value #12803
- Fixed: High latency and packet loss during a filter reload #12827
- Fixed: On startup "No routing address with matching address" might appear #12847
- Added: Toggle button to disable/enable multiple entries on NAT pages #12879
- Fixed: Delete button is always active for NAT rules, even if no rules are selected #12957
- Fixed: NAT Reflection generates duplicate rules when internal interface contains multiple VIPs in the same subnet #13012
- Fixed: NAT generates duplicate `no nat on` rules for port forwards with a destination of `Any` #13015
- Fixed: Input validation requires a gateway for floating `match out` rules #13027
- Fixed: Empty `negate_networks` table breaks policy routing rules #13049
- Fixed: The `negate_networks` table is not updated when an OpenVPN server is deleted #13055
- Added: Allow auto prefix with manual prefix-length in NPt #13070
- Fixed: Info icon on `firewall_nat_out.php` is incorrectly placed in manual outbound NAT mode #13164
- Fixed: Changing the redirect target for a Port Forward with an associated filter creates an incorrect firewall rule #13171
- Fixed: Incorrect usage of DSCP hex value #13178
- Fixed: TCP traffic sourced from the firewall can only use the default gateway #13420
- Fixed: `easyrule` CLI script has multiple bugs and undesirable behaviors #13445

- Changed: Correct DHCP client rule descriptions in the generated firewall ruleset #13505
- Fixed: Toggling NAT rules using the button method does not enable/disable corresponding firewall rules #13545
- Fixed: The "Kill States" button does not work consistently #14091
- Changed: Match upstream changes in PF syntax to disable fragment disassembly #14098
- Fixed: Associated firewall rule for NAT port forward does not inherit `nosync` property, gets synchronized #14335
- Fixed: Default tab on `firewall_rules.php` is not selected if the configuration has no WAN interface #14345
- Fixed: Outbound NAT rule input validation error when attempting to manually specify "Other Subnet" with a valid address #14354
- Fixed: Enable IPv6 over IPv4 tunneling option results in invalid PF rule #14415

SNMP

- Fixed: SNMP daemon is restarted during every `rc.newwanip` event #12611

Services

- Fixed: NTP service is not listed on `status_services.php` unless `config.xml` contains NTP configuration data #12775

Setup Wizard

- Changed: Update firewall host and domain fields in the Setup Wizard to match the description and warning text from `system.php` #14250

System Logs

- Fixed: Firewall log parser does not handle SCTP log entries #13940

Traffic Shaper (ALTQ)

- Changed: Remove code references to unused `reset` parameter from traffic shaper pages #13042
- Added: ALTQ GUI support for Broadcom Netextreme II (`bxe`) interfaces #13304
- Added: Include `ixv` in ALTQ capable NIC list #14408

Traffic Shaper (Limiters)

- Fixed: Incorrect ICMP reply when using limiters #9263
- Fixed: Pie and `fq_pie` are missing options and do not handle floating point number input correctly #12003
- Fixed: Utilize `dnctl(8)` to apply limiter changes without a filter reload #12579
- Fixed: Traffic routed through DUMMYNET by PF fails when IPFW is enabled #12954
- Fixed: Traffic shaped by limiters is dropped when routed to a GIF gateway #14055

Traffic Shaper Wizards

- Fixed: Traffic Shaper wizard can produce an invalid ruleset when configured with an IPv4 upstream SIP server #12937

Translations

- Fixed: Polish translation contains an invalid `sprintf()` format in the text for `firewall_nat_out_edit.php` #13946

UPnP/NAT-PMP

- Fixed: UPnP/NAT-PMP status page does not display all port mappings #4500
- Added: uPnP fails to properly give out subsequent reservations when multiple gaming systems are playing the same game/using the same port #7727
- Changed: Reorganize UPnP options #12624
- Changed: Update miniupnpd to 2.3.3 #14307

Unknown

- Fixed: Many `exec()` functions do not use full path to executable files #11941
- Fixed: URL scheme is not properly validated in some cases #14356

Upgrade

- Fixed: Upgrade does not work when using only IPv6 DNS servers #13162
- Fixed: `pfSense-boot` can fail to copy the EFI bootloader #14045

User Manager / Privileges

- Added: Support for RADIUS authentication over IPv6 #4154
- Fixed: Icon missing for user manager entries with a scope other than "user" #13174

Virtual IP Addresses

- Fixed: Firewall rules are not reloaded when removing a VIP, outdated rules/entries remain active #13908

Web Interface

- Fixed: Unnecessary link tag in login page #7996
- Fixed: "Dark" theme does not sufficiently distinguish between selected and deselected elements in option lists #11730
- Fixed: Lack of DNS or Internet connectivity causes GUI to be slow #12141
- Changed: GUI pages should use `POST` for AJAX calls, not `GET` #12431
- Fixed: Zero-value prefix IPv6 addresses are mishandled #12440
- Added: Option to filter state table contents by rule ID #12616
- Fixed: Changing RAM disk size does not prompt to reboot #12876
- Fixed: VGA install defaults to serial as primary console when loading/saving admin GUI settings without making changes #12960
- Fixed: Input validation for IPv6 addresses allows invalid address compression in some cases #13069
- Added: Trim whitespace from MAC addresses in user input #13109
- Changed: Spelling and typo corrections #13357
- Fixed: "Dark" theme uses the same colors for disabled and enabled input fields #13390
- Fixed: Input validation on `system_advanced_firewall.inc` uses incorrect variable references for some fields #13436
- Changed: Update external HTTPS/HTTP links #13440
- Fixed: Table row selection has poor contrast in Dark theme #13448
- Added: Support for `iwlwifi` wireless interfaces #14050

Wireless

- Fixed: Wireless interface WPA configuration fields are always visible #12998
- Fixed: Duplicate wireless interfaces are created at boot #12999

XMLRPC

- Fixed: Deleting a user on the primary node does not delete its home directory on secondary node during XMLRPC sync #12940
- Fixed: Filter/NAT rules configured with "No XMLRPC Sync" enabled are still synchronized #14316

3.22 2.5.2 New Features and Changes

This is a regularly scheduled software release including new features and bug fixes.

3.22.1 Known Issues / Errata

- Dynamic DNS incorrectly encodes NoIP update credentials #12021

3.22.2 Security

This release includes corrections for the following vulnerabilities in pfSense® software:

- pfSense-SA-21_02.captiveportal (XSS in Captive Portal client login page, #11843)

3.22.3 General

- Added: WireGuard add-on package

3.22.4 pfSense CE

Aliases / Tables

- Added: PHP shell playback script to modify Alias contents #11380

Authentication

- Added: Copy button for Authentication Server entries #11390

Backup / Restore

- Added: Randomize time of scheduled AutoConfigBackup runs #10811
- Fixed: Automated corruption recovery from cached `config.xml` backup files should check multiple backups #11748
- Fixed: AutoConfigBackup schedule custom hour value lost on page load #11946

Captive Portal

- Added: Redirect Captive Portal users to login page after they logout #11264
- Fixed: Captive Portal post-auth redirect is not properly respected #11842
- Fixed: Potential XSS vulnerability in Captive Portal `redirurl` handling #11843

Certificates

- Fixed: Certificate Manager does not report Unbound as using a certificate #11678
- Fixed: PHP error on certificate list due to unreadable private key #11859
- Fixed: Export P12 icon is missing if certificate is not locally renewable #11884

Configuration Upgrade

- Fixed: PHP error in `upgrade_212_to_213()` when upgrading certain IPsec tunnels #11801

Console Menu

- Changed: Allow reroot on ZFS from console and GUI reboot menu entries #11914

DHCP (IPv6)

- Fixed: `dhcp6withoutra_script.sh` does not get executed when advanced options are set #11883

DNS Forwarder

- Fixed: Disable DNSSEC option for dnsmasq #11781
- Fixed: Update dnsmasq to 2.85 to fix CVE-2021-3448 #11866

DNS Resolver

- Fixed: Unbound Python Integration repeatedly mounts `dev` without unmounting #11456
- Fixed: Stale hostname registration data for OpenVPN clients is not deleted from the DNS Resolver configuration at boot #11704
- Changed: Temporarily move back to Unbound 1.12.x due to instability on Unbound 1.13.x #11915

Dashboard

- Fixed: Thermal sensors widget no longer shows values from certain hardware #11787
- Fixed: IPsec Dashboard widget only displays first P2 subnet when using a single traffic selector #11893
- Fixed: Editing widgets on Dashboard causes a PHP Warning #11939

Diagnostics

- Fixed: ARP Table populates hostname values using expired DHCP lease data #11510
- Fixed: Sanitize OpenVPN Client Export certificate password in status output #11767
- Fixed: Sanitize Captive Portal RADIUS MAC secret in status output #11769
- Fixed: MAC address OEM information missing from ARP table #11819
- Fixed: State table content on `diag_dump_states.php` does not sort properly #11852

Dynamic DNS

- Added: New Dynamic DNS Provider: Mythic-Beasts #7842
- Added: New Dynamic DNS Provider: one.com #11293
- Added: New Dynamic DNS Provider: Yandex PDD #11294
- Added: New Dynamic DNS Provider: NIC.RU #11358
- Added: New Dynamic DNS Provider: Gandi LiveDNS IPv6 #11420
- Fixed: Automatic 25-day forced Dynamic DNS update removes wildcard domain #11667
- Fixed: Digital Ocean Dynamic DNS help text is incorrect #11754
- Fixed: NoIP.com Dynamic DNS update failure is not detected properly #11815
- Fixed: Dynamic DNS edit page incorrectly hides username field when switching away from Digital Ocean #11840

Gateways

- Added: Input validation to prevent setting a load balancing gateway group as default #11164

Hardware / Drivers

- Changed: Deprecate old cryptographic accelerator hardware which is not viable on modern systems #11426
- Fixed: Using SHA1 or SHA256 with AES-NI may fail if AES-NI attempts to accelerate hashing #11524

High Availability

- Fixed: Incorrect RADVD log message on HA event #11966

IGMP Proxy

- Fixed: IGMP Proxy restarts unnecessarily after IPv6 gateway events #11904

IPsec

- Added: GUI option to set RADIUS Timeout for EAP-RADIUS #11211
- Added: Option to switch IPsec filtering modes to choose between `enc` and `if_ipsec` filtering #11395
- Changed: Move custom IPsec NAT-T port settings to Advanced Options #11518
- Fixed: strongSwan configuration always contains user EAP/PSK values #11564
- Added: IPsec GUI option to control Child SA `start_action` #11576
- Fixed: Error when adding both IPv4 and IPv6 P2 under an IPv4 or IPv6 only IKEv1 P1 #11651
- Fixed: Cannot disable IPsec P1 when related P2s are in VTI mode and enabled #11792
- Fixed: IPsec VTI interface names are not properly formed for more than 32 interfaces #11794
- Fixed: Applying IPsec settings for more than ~30 tunnels times out PHP #11795

- Fixed: `ipsec_vti()` does not skip disabled VTI entries #11832
- Fixed: IPsec GUI allows creating multiple identical Phase 1 entries when using FQDN for remote gateway #11912
- Fixed: Mobile IPsec advanced RADIUS parameters do not allow numeric values with a decimal point #11967

IPv6 Router Advertisements (RADVD)

- Added: Use virtual link local IP address as RA source address for HA environments #11103
- Added: Shortcut buttons for service control and logs on RADVD configuration #11911
- Fixed: RADVD breaks on SIGHUP #11913

Interfaces

- Fixed: DHCP interfaces are always treated as having a gateway, even if one is not assigned by the upstream DHCP server #5135
- Fixed: Interfaces page displays MAC Address field for interfaces which do not support L2 #11387
- Fixed: CLI interface configuration without IPv6 leaves RA enabled #11609
- Fixed: Incomplete PPPoE custom reset values lead to invalid cron entry #11698
- Fixed: Error when changing MTU if the interface is used for both IPv4 and IPv6 default routes #11855
- Added: VLAN list sorting #11968

L2TP

- Fixed: Unused L2TP VPN files are not removed when the service is disabled #11299
- Added: GUI option to set MTU for L2TP VPN server #11406

NTPD

- Fixed: NTP widget displays incorrect status #11495
- Fixed: NTP authentication input validation rejects valid keys #11850

Notifications

- Fixed: Invalid HTML encoding in modal Notices window #11765

OpenVPN

- Added: Allow the firewall to use DNS servers provided to an OpenVPN client instance #11140
- Fixed: OpenVPN Wizard does not support gateway groups #11141
- Added: Set Explicit Exit Notify to 1 by default for new OpenVPN client instances #11521
- Added: Support for Cisco AVPair `{clientipv6}` template in firewall rules returns by RADIUS #11596
- Changed: Set `explicit-exit-notify` option by default for new OpenVPN server instances #11684
- Fixed: OpenVPN does not clean up parsed `Cisco-AVPair` rules on non-graceful disconnect #11699
- Fixed: OpenVPN does not kill IPv6 client states on disconnect #11700
- Fixed: OpenVPN client starts when CARP VIP is in BACKUP status when bound to Virtual IP aliased to CARP VIP #11793
- Fixed: Certificate validation with OCSP always fails in `openvpn.tls-verify.php` #11830
- Changed: Update OpenVPN to 2.5.2 #11844
- Fixed: OpenVPN client startup error if IPv6 Tunnel Network is defined in TAP mode #11869

Operating System

- Added: Kernel modules for alternate congestion control algorithms #7092
- Added: Kernel module for RTL8153 driver #11125
- Added: Xen console support #11402
- Fixed: Unquoted variable in `dot.tcshrc` can cause proxy password to be printed #11867

Routing

- Fixed: IPv4 link-local (`169.254.x.x`) gateway does not function #11806

Rules / NAT

- Added: Support for IPv6 firewall entries with dynamic delegated prefix and static host address #6626
- Fixed: Disabling all interfaces associated with a floating rule causes the firewall to generate an incorrect pf rule #11688
- Fixed: Input validation prevents creating 1:1 NAT rules on IPsec #11751
- Fixed: Invalid combinations of TCP flag matching options cause `pfctl` parser error #11762
- Fixed: Port forward rules only function through the default gateway interface, `reply-to` does not work for Multi-WAN (CE Only) #11805
- Fixed: Error loading rules in certain cases where an interface is temporarily without an address #11861
- Fixed: NAT 1:1 fail to validate aliases #11923

Traffic Shaper (ALTQ)

- Fixed: Harmless error when enabling traffic shaper #11229
- Fixed: Segmentation fault when loading ALTQ traffic shaping rules using FAIRQ #11550

Traffic Shaper (Limiters)

- Fixed: Unused Limiter entries with schedules create unnecessary cron jobs #11636
- Fixed: Error when setting queue limit on CODELQ limiter #11725

Upgrade

- Fixed: Language presented to user during upgrade is misleading #11897

Web Interface

- Added: Replace HTTP links with HTTPS in the GUI #11228
- Fixed: Ambiguous text in help and input validation error for system domain name #11658
- Fixed: PHP error if `PHP_error.log` file is too large #11685
- Fixed: RAM Disk Settings shows Kernel Memory at `0` Kb and does not allow the user to create RAM disks #11702
- Fixed: HTTP Referer error message text is incorrect #11873
- Fixed: Missing `/0` subnet when cloning repeatable CIDR mask controls #11880
- Fixed: Update NGINX to address CVE-2021-23017 #12061

WireGuard

- Fixed: Ignore WireGuard configurations under `<installedpackages></installedpackages>` #11808

Wireless

- Added: GUI options for WPA Enterprise with identity/password #2400
- Fixed: `wpa_supplicant` uses 100% of a CPU core at boot #11453

XMLRPC

- Fixed: XMLRPC synchronization restarts all OpenVPN instances on the secondary node when making any change on the primary node #11082
- Fixed: XMLRPC Client does not honor its default timeout value #11718

3.23 21.02/21.02-p1/2.5.0 New Features and Changes

pfSense® Plus software version 21.02 and pfSense Community Edition (CE) software version 2.5.0 include a major OS version upgrade, a kernel WireGuard implementation, OpenSSL upgrades, VPN and related security improvements, plus numerous other bug fixes and new features.

Warning: The original plan was to include a RESTCONF API in pfSense® Plus software version 21.02 and pfSense software version 2.5.0, which for security reasons would have required hardware AES-NI or equivalent cryptographic accelerator support. Plans have since changed, and these versions do not contain the planned RESTCONF API, thus **pfSense® Plus software version 21.02 and pfSense Community Edition (CE) software version 2.5.0 DO NOT require AES-NI**.

Tip: For those who have not yet updated to 2.4.5-p1, consult the *previous release notes* and blog posts for those releases to read all important information and warnings before proceeding.

3.23.1 pfSense Plus

Version 21.02 is the first release of pfSense Plus software, formerly known as Factory Edition. For more details about the distinctions between pfSense Plus and pfSense CE, read the pfSense Plus Announcement. Customers running the Factory Edition of pfSense software version 2.4.5-p1 and older can upgrade in-place automatically to pfSense Plus software version 21.02 as with any other previous upgrade.

In this version, the changes in pfSense Plus software and pfSense CE software are roughly the same, with a few notable exceptions which are only available in pfSense Plus software:

- Support for Intel® QuickAssist Technology, also known as QAT.
  - QAT accelerates cryptographic and hashing operations on supported hardware, and can be used to accelerate IPsec, OpenVPN, and other OpenCrypto Framework-aware software.
  - Supported hardware includes many Intel-based systems sold by Netgate (e.g. XG-7100, SG-5100) and add-on cards.
  - From the FreeBSD man page:
    - The `qat` driver supports the QAT devices integrated with Atom C2000 and C3000 and Xeon C620 and D-1500 chipsets, and the Intel QAT Adapter 8950.
    - It can accelerate AES in CBC, CTR, XTS (except for the C2000) and GCM modes, and can perform authenticated encryption combining the CBC, CTR and XTS modes with SHA1-HMAC and SHA2-HMAC. The qat driver can also compute SHA1 and SHA2 digests.
- Improved SafeXcel cryptographic accelerator support for SG-2100 and SG-1100 which can improve IPsec performance.
  - From the FreeBSD man page:
    - The driver can accelerate the following AES modes: AES-CBC, AES-CTR, AES-XTS, AES-GCM, AES-CCM
    - The driver also implements SHA1 and SHA2 transforms, and can combine AES-CBC and AES-CTR with SHA1-HMAC and SHA2-HMAC for encrypt-then-authenticate operations.
- Updated *IPsec profile export*

  - Exports Apple profiles compatible with current iOS and macOS versions
  - New export function for Windows clients to configure tunnels using PowerShell

Version 21.02-p1

pfSense Plus software version 21.02-p1 is a special patch release to address a kernel problem affecting the SG-3100 which caused system instability (#11444). No additional fixes are present in the 21.02-p1 release.

See the detailed bug analysis blog post for more details.

3.23.2 Operating System / Architecture changes

- Base OS upgraded to FreeBSD **12.2-STABLE**
- OpenSSL upgraded to **1.1.1i-freebsd**
- PHP upgraded to **7.4** #9365 #10659
- Python upgraded to **3.7** #9360

3.23.3 Known Issues / Errata

- Deprecated the built-in relayd Load Balancer #9386
  - `relayd` does not function with OpenSSL 1.1.x
  - The `relayd` FreeBSD port has been changed to require libressl – There is no apparent sign of work to make it compatible with OpenSSL 1.1.x
  - The HAProxy package may be used in its place; It is a much more robust and more feature-complete load balancer and reverse proxy
  - For more information on implementing HAProxy, see *HAProxy package* and the Hangout
- There is an issue in this release with port forwarding on pfSense Plus software installations with multiple WANs, which has been resolved in the 21.02.2 patch release, see #11436 for details.
- There is an issue with AES-NI hash acceleration for SHA1 and SHA-256. If the AES-NI driver detects a system capable of accelerating SHA1 or SHA-256 and the firewall attempts to utilize one of those hashes, the affected operation may fail. This affects IPsec and OpenVPN, among other uses. pfSense Plus users can change to QAT acceleration on supported hardware instead. In cases where QAT is unavailable, change to AES-GCM, change to a different unaccelerated hash (e.g. SHA-512), or disable AES-NI. See #11524 for details.
- There is a similar issue which affects SafeXcel SHA1 and SHA2 hash acceleration on SG-1100 and SG-2100. On that hardware, change to an AEAD cipher such as AES-GCM or switch to an unaccelerated hash. This issue is being tracked internally on NG #6005
- The FRR package on pfSense Plus 21.02 and pfSense CE 2.5.0 and later no longer exchanges routes with BGP peers by default without being explicitly allowed to do so. This is more secure behavior but requires a manual change. To replicate the previous behavior, use **ONE** of the following workarounds:
  - Navigate to **Services > FRR BGP** on the **Advanced** tab and check *Disable eBGP Require Policy*, then **Save**.
  - Instead of disabling the policy check, create route maps which match and allow expected incoming and outgoing routes explicitly. This is the most secure method. See *Peer Filtering* and *BGP Example Configuration* for more information.

- Manually create a route map to permit all routes (Name: `allow-all`, Action: *Permit*, Sequence: `100`), then set that route map on BGP neighbors for inbound and outbound peer filtering. This can be used as a placeholder for later migration to more secure route map filtering.

Warning: See the FreeBSD 12.0 Release Notes for information on deprecated hardware drivers that may impact firewalls upgrading to pfSense software version 2.5.0. Some of these were renamed or folded into other drivers, others have been removed, and more are slated for removal in FreeBSD 13 in the future.

3.23.4 Aliases/Tables

- Fixed aliases to allow IPv6 prefix entries which end in IPv4 addresses (e.g. `x:x:x:x:x:x:d.d.d.d` from RFC 4291 section 2.2.2) #10694
- Fixed a PHP error processing aliases when the configuration contains no aliases section #9936
- Fixed URL-based Alias only storing last-most entry in the configuration #9074
- Fixed an issue with PF tables remaining active after they had been deleted #9790
- Added Internationalized domain names support for aliases #7255
- Added the ability to copy an existing alias when creating a new entry #6908
- Fixed handling of URL-based aliases containing multiple URLs #11256

3.23.5 Authentication

- Added RADIUS authentication for SSH users #10545
- Added LDAP authentication for SSH users #8698
- Added option to control behavior of unauthenticated LDAP binds #9909
- Converted LDAP TLS setup from environment variables to LDAP_OPT_X_TLS_* options #9417
- Set RADIUS NAS Identifier to include `webConfigurator` and the firewall hostname when logging in the GUI #9209
- Added LDAP extended query for groups in RFC2307 containers #9527
- Fixed errors when using RADIUS for GUI authentication while the WAN is down #11109

3.23.6 Backup/Restore

- Changed `crypt_data()` to use stronger key derivation #9421
- Updated `crypt_data()` syntax for OpenSSL 1.1.x #9420 #10178
- Disabled AutoConfigBackup manual backups when AutoConfigBackup is disabled #9785
- Improved error handling when attempting to restore encrypted and otherwise invalid configurations which result in errors (e.g. wrong encryption passphrase, malformed XML) #10179
- Added option to include the DHCP v4/v6 leases database in config.xml backups #10910
- Added option to include the Captive Portal database in config.xml backups #10868
- Added option to include the Captive Portal used MACs database in config.xml backups #10856

- Added option to prevent all extra data from being added to config.xml backups #10914
- Added password confirmation when encrypting a config.xml backup #10301
- Added support for GPT partitioned drives to the External Configuration Locator #9097
- Added support for Limiters to the Traffic Shaper backup and restore area option #4763
- Added option to backup Dynamic DNS area #3559
- Fixed restoration of active voucher data from backup #3128

3.23.7 Captive Portal

- Improved XMLRPC sync of Captive Portal database information #97
- Changed Captive Portal vouchers to use `phpseclib` so it can generate keys natively in PHP, and to work around OpenSSL deprecating key sizes needed for vouchers #9443
- Added `trim()` to the submitted username, so that spaces before/after in input do not cause authentication errors #9274
- Optimized Captive Portal authentication attempts when using multiple authentication servers #9255
- Fixed Captive Portal session timeout values for RADIUS users who do not have a timeout returned from the server #9208
- Changed Captive Portal so that users no longer get disconnected when changes are made to Captive Portal settings #8616
- Added an option so that Captive Portals may choose to remove or retain logins across reboot #5644
- Fixed deletion of related files when removing a Captive Portal zone #10891
- Fixed XMLRPC sync of Captive Portal used MACs database #10857
- Added validation of Captive Portal zone names to prevent using reserved words #10798
- Added support for IDN hostnames to Captive Portal Allowed Hostnames tab #10747
- Improved Captive Portal Allowed Hostnames so it supports multiple DNS records in responses #10724
- Fixed retention of automatic pass-through MAC entries when using Captive Portal Vouchers #9933
- Fixed Captive Portal Bandwidth per-user bandwidth limit values being applied when disabled #9437 #9311
- Changed handling of voucher logins with Concurrent Login option so that new logins are prevented rather than removing old sessions #9432 #2146
- Changed XMLRPC behavior to not remove zones from secondary node when disabling Captive Portal #9303
- Fixed XMLRPC sync failing to propagate voucher roll option changes to the secondary node #8809
- Fixed XMLRPC sync failing to create Captive Portal voucher files on secondary node #8807
- Fixed Captive Portal + Bridge interface validation #6528
- Added support for masking of Captive Portal pass-thru MACs #2424
- Added support for pre-filling voucher codes via URL parameters, so they can be used via QR code #1984

3.23.8 Certificates

- Fixed OCSP stapling detection for OpenSSL 1.1.x #9408
- Fixed GUI detection of revoked status for certificates issued and revoked by an intermediate CA #9924
- Removed PKCS#12 export links for entries which cannot be exported in that format (e.g. no private key) #10284
- Added an option to globally trust local CA manager entries #4068
- Added support for randomized certificate serial numbers when creating or signing certificates with local internal CAs #9883
- Added validation for CA/CRL serial numbers #9883 #9869
- Added support for importing ECDSA keys in certificates and when completing signing requests #9745
- Added support for creating and signing certificates using ECDSA keys #9843 #10658
- Added detailed certificate information block to the CA list, using code shared with the Certificate list #9856
- Added Certificate Lifetime to certificate information block #7332
- Added CA validity checks when attempting to pre-fill certificate fields from a CA #3956
- Added a daily certificate expiration check and notice, with settings to control its behavior and notifications (Default: 27 days) #7332
- Added functionality to import certificates without private keys (e.g. PKCS#11) #9834
- Added functionality to upload a PKCS#12 file to import a certificate #8645
- Added CA/Certificate renewal functionality #9842
  - This allows a CA or certificate to be renewed using its current settings (or a more secure profile), replacing the entry with a fresh one, and optionally retaining the existing key.
- **Added an "Edit" screen for Certificate entries**
  - This view allows editing the Certificate **Descriptive name** field #7861
  - This view also adds a (not stored) password field and buttons for exporting encrypted private keys and PKCS#12 archives #1192
- **Improved default GUI certificate strength and handling of weak values #9825**
  - Reduced the default GUI web server certificate lifetime to 398 days to prevent errors on Apple platforms #9825
  - Added notes on CA/Cert pages about using potentially insecure parameter choices
  - Added visible warnings on CA/Cert pages if parameters are known to be insecure or not recommended
- **Revamped CRL management to be easier to use and more capable**
  - Added the ability to revoke certificates by serial number #9869
  - Added the ability to revoke multiple entries at a time #3258
  - Decluttered the main CRL list screen
  - Moved to a single CRL create control to the bottom under the list rather than multiple buttons
- **Optimized CA/Cert/CRL code in various ways, including:**
  - Actions are now performed by `refid` rather than array index, which is more accurate and not as prone to being affected by parallel changes
  - Improved configuration change descriptions as shown in the GUI and configuration history/backups

  - Miscellaneous style and code re-use improvements
  - Changed CA/Cert date calculations to use a more accurate method, which ensures accuracy on ARM past the 2038 date barrier #9899

3.23.9 Configuration Backend

- Changed error handling on boot error 'XML configuration file not found' so the user is given an opportunity to fix the problem manually #10556

3.23.10 Configuration Upgrade

- Retired m0n0wall configuration upgrade support #10997

3.23.11 Console Menu

- Fixed `rc.initial` execution of `rc.local.running` #10978
- Fixed `rc.initial` handling of `-c` commands with arguments #10603
- Fixed console menu display of subnet masks for DHCP interfaces #10740

3.23.12 Dashboard

- Added PPP uptime to the Dashboard Interfaces Widget #9426
- Improved long description truncation behavior in the services status widget #10795
- Fixed Dashboard traffic graph widget display of bandwidth units (b/s vs. B/s) #9072
- Added adaptive state timeout indication to the state table usage meter #7016
- Fixed Thermal Sensors dashboard widget showing invalid sensors #10963
- Added default route indicator to Gateways widget #11057
- Added hardware interface name as a tooltip on Interfaces widget entries #11041

3.23.13 DHCP (IPv4)

- Fixed handling of spaces in DHCP lease hostnames by `dhcpleases` #9758
- Fixed DHCP leases hostname parsing problems which prevented some hostnames from being displayed in the GUI #3500
- Added OMAPI settings to the DHCP Server #7304
- Increased number of NTP servers sent via DHCP to 3 #9661
- Added an option to prevent known DHCP clients from obtaining addresses on any interface (e.g. known clients may only obtain an address from the interface where the entry is defined) #1605
- Added count of static mappings to list when editing DHCP settings for an interface #9282
- Fixed handling of client identifiers on static mappings containing double quotes #10295
- Added ARM32/64 network booting support to the DHCP Server #10374

- Increased the number of NTP servers for DHCP Static Mappings #10333
- Fix DHCP Dynamic DNS handling of per-host zone and key options from static mappings #10224
- Added per-host custom BOOTP/DHCP Options to static mappings #8990
- Added a button to clear all DHCP leases #7406
- Fixed ARPA zone declaration formatting in DHCP server configuration file #11224

3.23.14 DHCP (IPv6)

- Added options to disable pushing IPv6 DNS servers to clients via DHCP6 #9302
- Fixed DHCPv6 domain search list #10200
- Fixed validation to allow omission of DHCPv6 range for use with stateless DHCP #9596
- Fixed issues creating IPv6 Static Mappings #7443
- Fixed DHCPv6 merging an IPv6 prefix with the input submitted in DNS servers field when using Track Interface #7384
- Fixed prefix delegation not being requested if no interfaces were set to track6 #11005
- Fixed DHCPv6 Dynamic DNS domain key name validation #10844
- Fixed line formatting issues in the DHCPv6 configuration file #10675
- Fixed prefix not being included in the DNS entry registered by DHCPv6 #8156
- Fixed DHCPv6 static mapping changes requiring a restart of the DNS resolver to activate #10882
- Fixed issues running DHCPv6 on certain types of tracked interfaces (e.g. bridges, VLANs) #3965
- Fixed issues with WAN not renewing IPv6 address after an upstream failure #10966

3.23.15 DHCP Relay

- Fixed DHCP Relay validation to allow OpenVPN TAP interfaces #10711
- Fixed inconsistent validation behavior for DHCP relay and bridges #7778

3.23.16 Diagnostics

- Added Reroot and Reboot with Filesystem Check options to GUI Reboot page #9771
- Added option to control wait time between ICMP echo request (ping) packets `diag_ping.php` #9862
- Improved data sanitization in status.php #10946 #10944 Sanitize MaxMind GeoIP key #10797 #10569 #10794
- Added config history list to status.php #10696
- Added DNS Resolver configuration to status.php #10635
- Added L2TP VPN configuration to status.php #10583
- Changed pftop page to hide filtering controls for views which do not support filtering #10625
- Added support for IDN hostnames to DNS Lookup, Ping, and Traceroute #10538
- Fixed diag_dns.php link to Ping passing incorrect parameters #10537
- Added a button to clear the NDP cache #10975

- Added a button to clear the ARP cache #4038
- Fixed hostname being ignored when DNS Lookup calculates response time #11018
- Fixed **Kill States** button on `diag_dump_states.php` when used with CIDR-masked subnets #9270

3.23.17 DNS Forwarder

- Updated `dnsmasq` to 2.84 #11278

3.23.18 DNS Resolver

- Added IPv6 OpenVPN client addresses resolution to the DNS Resolver #8624
- Added DNS64 options to the DNS Resolver #10274
- Added support for multiple IP addresses in a DNS Resolver Host Override entry #10896
- Fixed DNS Resolver restart commands to work around potential environment issues #10781
- Fixed saving DNS Resolver ACL entries when using a non-English translation #10742
- Added support for IDN symbols in DNS Resolver ACL entries #10730
- Added Aggressive NSEC option to the DNS Resolver #10449
- Fixed DNS Resolver unintentionally retaining DHCP registration entries after disabling that feature #8981
- Fixed DNS Resolver restarting on every OpenVPN client connection when registering clients in DNS #11129
- Fixed issues with the DNS Resolver not starting when bound to disabled interfaces or interfaces without carrier #11087
- Fixed DNS Resolver custom TLS listen port being ignored #11051
- Improved formatting and ordering of items in the DNS Resolver access list configuration file #11309

3.23.19 Dynamic DNS

- Fixed Dynamic DNS Dashboard Widget address parsing for entries with split hostname/domain (e.g. Namecheap) #9564
- Added support for new CloudFlare Dynamic DNS API tokens #9639
- Added IPv6 support to No-IP Dynamic DNS #10256
- Fixed issues with Hover Dynamic DNS #10241
- Updated Cloudflare Dynamic DNS to query Zone ID with token #10992
- Added support for IPv6 to easyDNS Dynamic DNS #10972
- Added support for Domeneshop Dynamic DNS #10826
- Added Zone option to RFC 2136 Dynamic DNS #10684
- Updated FreeDNS Dynamic DNS to use their v2 API #10617
- Fixed DigitalOcean Dynamic DNS processing of zones with multiple pages of records #10592
- Improved Dynamic DNS Logging #10459
- Added support for dynv6.com Dynamic DNS #9642

- Fixed handling of Dynamic DNS AAAA records on 6rd tunnel interfaces bound to PPPoE interfaces #9641
- Added a button to duplicate Dynamic DNS entries #8952
- Fixed Dynamic DNS update for HE.net Tunnelbroker always setting IP address of the default WAN interface #11024
- Updated HE.net Tunnelbroker Dynamic DNS to use their current API #11037
- Added support for Wildcard A records for Gandi Dynamic DNS #11159
- Updated No-IP Dynamic DNS to use a newer API #6638
- Fixed Namecheap Dynamic DNS error code checking #5308
- Improved color blind accessibility of Dynamic DNS status #3229

3.23.20 Gateways

- Added support for obtaining a gateway via DHCP which is outside of the interface subnet #7380
- Added validation to prevent using descriptions on interfaces which would cause gateway names to exceeded the maximum allowed length #9401
- Added tooltip text to icons on the Gateways #10719
- Fixed issues with dpinger failing to update IPv6 gateway address on DHCPv6 WAN interfaces #8136

3.23.21 Hardware / Drivers

- Added `bnxt` driver for Broadcom NetXtreme interfaces #9155
- Added iOS/Android/Generic USB tethering driver #7467

3.23.22 IGMP Proxy

- Added input validation for IGMP Proxy settings #7163

3.23.23 Installer

- Created separate **Auto (UFS) UEFI** and **Auto (UFS) BIOS** installation options to avoid problems on hardware which boots differently on USB and non-USB disks #8638
- Fixed reinstalling with UFS on a ZFS formatted drive #10690
- Fixed platform detection for MBT-4220 and MBT-2220 on newer BIOS revisions #9242
- Fixed an issue with shutting down instead of rebooting after installing using ZFS #7307

3.23.24 Interfaces

- Added support for using IPv4 and IPv6 addresses on GRE interfaces at the same time #10392
- Added a check to disable Hardware Checksum Offloading in environments with interfaces which do not support it (e.g. `vtnet`, `ena`) #10723
- Changed the way interface VLAN support is detected so it does not rely on the VLANMTU flag #9548
- Added a PHP shell playback script `restartallwan` which restarts all WAN-type interfaces #9688
- Changed assignment of the `fe80::1:1` default IPv6 link-local LAN address so it does not remove existing entries, which could cause problems such as Unbound failing to start #9998
- Added automatic MTU adjustment for GRE interfaces using IPsec as a transport #10222
- Fixed SLAAC interface selection when using IPv6 on a link which also uses PPP #9324
- Added GUI interface descriptions to Operating System interfaces #1557
- Added the ability to assign virtual type interfaces (IPsec, OpenVPN, GIF, GRE, etc) during console interface assignment #10947
- Fixed TSO not being disabled in some cases #10836
- Fixed group name length input validation #10835
- Improved interface caching for environments with many interfaces #10680
- Fixed `fe80::1:1` being added to interfaces without track6 #10661
- Added a check to prevent stf (6RD/6to4) interfaces from being used as parent interfaces #10626
- Fixed redundant disabling of static ARP at boot before it could be enabled #10589
- Fixed initialization of bridges which include a GIF interface at boot #10524
- Fixed problems with post-install interface changes not being retained if the user did not complete the wizard #10383
- Fixed inefficiencies when applying settings to a VLAN parent interface #9154
- Fixed interface MTU setting not being applied to all IPv6 routes #6868
- Fixed handling of MTU setting for 6rd and 6to4 interfaces #6377
- Fixed IPv6 IP Alias preventing Track Interface from working with DHCPv6 and RA #5999
- Changed DHCP interface renewal behavior to not restart services if the IP address did not change #11142
- Fixed an error when changing bridge STP settings #11122
- Added a binary package with updated Realtek interface drivers #11079
- Improved link state visibility on Status > Interfaces #11045
- Removed VTI interfaces from Interface Group selection since they do not currently function in this manner #11134
- Fixed issues with IPv6 on top of IPv4 PPPoE placing default route on incorrect interface #9324

3.23.25 IPsec

- Added 25519 curve-based IPsec DH and PFS groups 31 and 32 #9531
- Enabled the strongSwan PKCS#11 plugin #6775
- Added support for ECDSA certificates to IPsec for IKE #4991
- Renamed IPsec "RSA" options to "Certificate" since both RSA and ECDSA certificates are now supported, and it is also easier for users to recognize #9903
- Converted IPsec configuration code from `ipsec.conf` `ipsec/stroke` style to `swanctl.conf` `swanctl/vici` style #9603
  - Split up much of the single large IPsec configuration function into multiple functions as appropriate.
  - Optimized code along the way, including reducing code duplication and finding ways to generalize functions to support future expansion.
  - For IKEv1 and IKEv2 with Split Connections enabled, P2 settings are properly respected for each individual P2, such as separate encryption algorithms #6263
    - **N.B.:** In rare cases this may expose a previous misconfiguration which allowed a Phase 2 SA to connect with improper settings, for example if a required encryption algorithm was enabled on one P2 but not another.
  - New GUI option under **VPN > IPsec**, **Mobile Clients** tab to enable RADIUS Accounting which was previously on by default. This is now disabled by default as RADIUS accounting data will be sent for every tunnel, not only mobile clients, and if the accounting data fails to reach the RADIUS server, tunnels may be disconnected.
  - Additional developer & advanced user notes:
    - For those who may have scripts which touched files in `/var/etc/ipsec`, note that the structure of this directory has changed to the new swanctl layout.
    - Any usage of `/usr/local/sbin/ipsec` or the stroke plugin must also be changed to `/usr/local/sbin/swanctl` and VICI. Note that some commands have no direct equivalents, but the same or better information is available in other ways.
    - IPsec start/stop/reload functions now use `/usr/local/sbin/strongswanrc`
    - IPsec-related functions were converged into `ipsec.inc`, removed from `vpn.inc`, and renamed from `vpn_ipsec_<name>` to `ipsec_<name>`
  - Reworked how reauthentication and rekey behavior functions, giving more control to the user compared to previous options #9983
- Reformatted `status_ipsec.php` to include more available information (rekey timer, encryption key size, IKE SPIs, ports) #9979
- Added support for PKCS#11 authentication (e.g. hardware tokens such as Yubikey) for IPsec #9878
- Fixed usage of Hash Algorithm on child ESP/AH proposals using AEAD ciphers #9726
- Added support for IPsec remote gateway entries using FQDNs which resolve to IPv6 addresses #9405
- Added manual selection of Pseudo-Random Function (PRF) for use with AEAD ciphers #9309
- Added support for using per-user addresses from RADIUS and falling back to a local pool otherwise #8160
- Added an option which allows multiple tunnels to use the same remote peer in certain situations (read warnings on the option before use) #10214
- Improved visible distinction of online/offline mobile IPsec users in the IPsec status and dashboard widget #10340

- Added options to change the IPsec NAT-T ports (local and remote) #10870
- Improved boot-time initialization of IPsec VTI interfaces #10842
- Added support for limiting IPsec VPN access by RADIUS user group #10748
- Changed IPsec to share the same RADIUS Cisco-AVPair parser code as OpenVPN for Xauth users #10469
- Fixed handling of IPsec VTI interfaces in environments with large numbers of IPsec tunnels #9592
- Added IPsec Advanced option to control maximum allowed Parallel P2 Rekey exchanges #9331
- Fixed issues with bringing up new Phase 2 entries on IPsec tunnels with "Split connections" enabled #8472
- Fixed issues where, in rare cases, IPsec tunnels would not reconnect until the firewall was rebooted #8015
- Improved the Remote Gateway field description for IPsec Phase 1 entries to indicate that `0.0.0.0` is allowed #7095
- Fixed issues with IKEv2 IPsec tunnels with multiple phase 2 entries combining traffic selectors in unexpected ways (set "Split Connections" to isolate them) #6324
- Added options to create IPsec bypass rules which prevent specific source and destination network pairs from entering policy-based IPsec tunnels #3329
- Documented settings which work around SA duplication issues experienced by users in certain cases #10176
- Improved IPsec GUI options for P1/P2 SA expiration and replacement to help prevent SA duplication #11219
- Fixed a PHP error in mobile IPsec input validation #11212
- Added validation to prevent unsupported wildcard certificates from being selected for use with IPsec #11297

3.23.26 IPv6 Router Advertisements (RADVD)

- Fixed Router Advertisement configuration missing information in Unmanaged mode #9710
- Fixed Router Advertisement lifetime input validation #10709

3.23.27 L2TP

- Fixed L2TP secret using an empty value after removing it from the GUI #10710
- Fixed L2TP input validation to allow leaving the remote address field blank when assigning addresses from RADIUS #7562
- Fixed inefficiencies in the initial L2TP reconfiguration process #7558
- Fixed L2TP Server and Client both using `l2tpX` for interface names #11006
- Fixed static routes on L2TP interfaces not being reapplied when reconnecting #10407
- Fixed L2TP server being restarted when making user account changes #11059

3.23.28 LAGG Interfaces

- Improved Interface Status and Widget information for LAGG #9187
- Fixed route for GIF/GRE peer when using VLAN on LAGG #10623
- Added option to toggle LACP PDU transmission fast timeout #10504
- Fixed LAGG member interface events causing filter reloads #10365
- Fixed issues with LAGG interface MTU being incorrectly applied to VLAN subinterfaces #8585
- Added option to control the master interface for LAGG in Failover mode #1019

3.23.29 Logging

- Changed system logging to use plain text logging and log rotation, the old binary clog format has been deprecated #8350
- Updated default log size (512k + rotated copies), default lines to display (500, was 50), and max line limits (200k, up from 2k) #9734
- Added log tabs for nginx, userlog, utx/lastlog, and some other previously hidden logs #9714
- Relocated Package Logs into a tab under System Logs and standardized display/filtering of package logs #9714
- Added GUI options to control log rotation #9711
- Added code for packages to set their own log rotation parameters #9712
- Removed the redundant `nginx-error.log` file #7198
- Fixed some instances where logs were mixed into the wrong log files/tabs (Captive Portal/DHCP/squid/php/others) #1375
- Reorganized/restructured several log tabs #9714
- Added a dedicated authentication log #9754
- Added an option for RFC 5424 format log messages which have RFC 3339 timestamps #9808
- Fixed an issue where a firewall log entry for loopback source/destination occasionally reported `127.0.0.1` as `127.0.01` #10776
- Fixed issues with `syslogd` using an old IP address after an interface IP address change #9660
- Added `watchfrr` to routing log #11207

3.23.30 Multi-WAN

- Fixed Gateways being removed from routing groups based on low alert thresholds #10546
- Fixed a possible race condition in gateway group fail-over causing unexpected behavior #9450
- Fixed a load balancing failure when one gateway had a weight of 1 and another gateway had a weight >1 #6025

3.23.31 NAT Reflection

- Fixed port forwards where the destination is a network alias creating invalid refection rules if multiple subnets are in that alias #7614

3.23.32 Notifications

- Deprecated & Removed Growl Notifications #8821
- Added a daily certificate expiration notification with settings to control its behavior #7332
- Fixed input validation of SMTP notification settings #8522
- Added support for sending notifications via Pushover API #10495
- Added support for sending notifications via Telegram #10354
- Fixed a PHP error when SMTP notifications fail #11063

3.23.33 NTPD

- Added GUI options for NTP sync/poll intervals #6787
- Added validation to prevent using `noselect` and `noserve` with pools #9830
- Added feature to automatically detect GPS baud rate #7284
- Fixed status and widget display of long hostnames and stratum #10307
- Fixed handling of the checkbox options on NTP servers #10276
- Updated GPS initialization commands for Garmin devices #10327
- Added an option to limit NTP pool server usage #10323
- Added option to force IPv4/IPv6 DNS resolution for NTP servers #10322
- Added support for NTP server authentication #8794
- Added an option to disable NTP #3567
- Added units to the NTP status page #2850

3.23.34 OpenVPN

- Updated OpenVPN to 2.5.0 #11020
  - The default compression behavior has changed for security reasons. Incoming packets will be decompressed, outgoing packets will not be compressed. There is a GUI control to alter this behavior.
  - Data cipher negotiation (Formerly known as Negotiable Cryptographic Parameters, or NCP) is now compulsory. Disabling negotiation has been deprecated. The option is still present in the GUI, but negotiation will be unilaterally enabled on upgrade. The upgrade process will attempt to use the expected data encryption algorithms before and after the upgrade completes, but in some cases more secure algorithms may be enabled as well. #10919

    We strongly encourage using AEAD ciphers such as AES-GCM, future versions of OpenVPN will require them and will not have configurable cipher lists.
- Added connection count to OpenVPN status and widget #9788

- Enabled the OpenVPN x509-alt-username build option #9884
- Restructured the OpenVPN settings directory layout
  - Changed from `/var/etc/openvpn[-csc]/<mode><id>.<file>` to `/var/etc/openvpn/<mode><id>/<x>`
    - This keeps all settings for each client and server in a clean structure
- Moved to `CApath` style CA structure for OpenVPN CA/CRL usage #9915
- Added support for OCSP verification of client certificates #7767
- Fixed a potential race condition in OpenVPN client ACLs obtained via RADIUS #9206
- Added support for more protocols (IP, ICMP), ports, and a template variable (`{clientip}`) in OpenVPN client ACLs obtained via RADIUS #9206
- Added the ability to register OpenVPN Remote Access (User Auth) clients in the DNS Resolver #10999
- Fixed an issue where duplicating an OpenVPN instance did not copy the password #10703
- Fixed issues with OpenVPN TCP clients failing to start #10650
- Added support for IPv6 OpenVPN ACLs obtained via RADIUS #10454
- Fixed validation to enforce OpenVPN client password usage when setting a username, to prevent a missing password from interrupting the boot process #10409
- Enabled asynchronous push in OpenVPN binary #10273
- Added OpenVPN client-specific override option to ignore routes pushed by the server ("push-reset") #9702
- Clarified behavior of OpenVPN server option for Duplicate Connections #10363

3.23.35 Operating System

- Fixed a network performance regression in the fast forwarding path with IP redirects enabled NG4965
- Fixed double ZFS entries in `loader.conf` #10375
- Added a method to enable persistent command history in the shell #11029
- Changed the default domain name of the firewall from `.localdomain` to `.home.arpa` #10533

3.23.36 Package System

- Disabled spell checking on package upgrade progress textarea #10637
- Fixed issues with package upgrade or reinstall hanging indefinitely #10610
- Fixed description used for buttons when editing packages #11208
- Deprecated the following packages: OpenBGPd, Quagga_OSPF, routed, blinkled, and gwled

3.23.37 PPP Interfaces

- Fixed issues with PPPoE over a VLAN failing to reconnect #9148
- Enabled selection of QinQ interfaces for use with PPP #9472
- Added option to set `Host-Uniq` value for PPPoE #10597
- Fixed incorrect interface assignment after switching from PPPoE #10240
- Fixed IPv6 not being disabled in mpd.conf when the IPv6 GUI option is set to 'disabled' #7386
- Fixed PPPoE interface errors due to MTU settings #11035

3.23.38 PPPoE Server

- Fixed PPPoE server ignoring secondary RADIUS Server #10926
- Fixed PPPoE server Accounting updates option #10869
- Removed unnecessary restarts of the PPPoE server when adding/modifying users #10318
- Added input validation to prevent enabling the PPPoE server on a PPPoE client interface #4510

3.23.39 Routing

- Fixed automatic static routes set for DNS gateway bindings not being removed when no longer necessary #8922
- Fixed missing tooltip text for icons on the Static Routes Page #10889

3.23.40 RRD Graphs

- Fixed RRD graph handling of NTP graph data with negative `freq` values #6503
- Fixed RRD graph creation for interfaces using CODELQ #6277

3.23.41 Rules / NAT

- Added the ability to configure negated tagging, to match packets which do not contain a given tag #10186
- Added support for IPv6 Port Forwards #10984
- Fixed handling of IPv6 NPt rules on 6rd WAN interfaces #10757
- Fixed 1:1 NAT issue when internal interface has VIPs #10752
- Fixed policy routing rules not being written correctly for a down gateway #10716
- Added EoIP to firewall rule Protocol list #10698
- Fixed separator bars on floating rules not covering the full table width #10667
- Fixed 1:1 NAT for IPv6 applying wrong subnet mask to "Single Host" #7742
- Added validation to prevent accidentally overlapping NPt networks and interface networks #7741
- Added support for dynamic interface addresses in 1:1 NAT rules #7705
- Added default values of TCP and UDP timeouts to the GUI #7362
- Fixed handling of IPv6 floating rules on 6rd interfaces #7142

- Fixed firewall rules for "PPPoE clients" only including the first PPPoE server instance #6598
- Fixed duplicated tracker IDs on block private networks rules #6030
- Fixed `reply-to` on rules for PPPoE WANs with IPv6 SLAAC #5258
- Added gateway/group IP addresses to mouseover on rules #885
- Fixed formatting of floating rules with large numbers interfaces #10892
- Fixed form rendering issues with Port Forward Address Fields in Safari #10674
- Fixed firewall ruleset failing to load at boot when new ruleset would be invalid #6028
- Fixed an issue adding or deleting separator bars when no rules are present #10827

3.23.42 S.M.A.R.T.

- Updated S.M.A.R.T. Page with new capabilities #9367

3.23.43 SNMP

- Fixed SNMP reporting incorrect speed for switch uplink interface on Netgate SG-3100 #10793
- Fixed SNMP input validation to require the Host Resources module when the PF module is also enabled #10471

3.23.44 Traffic Graphs

- Changed the Traffic Graph page from `rate` to `iftop` which brings IPv6 support and various other improvements #3334

3.23.45 Traffic Shaper (ALTQ)

- Changed default ALTQ queue bandwidth type to Mbit/s #10988
- Updated traffic shaper wizard settings for XBox and Wii ports #10837
- Added Broadcom NetXtreme to ALTQ-capable list #10762
- Added ALTQ support to the `ix(4)` driver #7378
- Fixed deletion of associated shaper queues when deleting an interface #3488
- Fixed ALTQ root queue bandwidth calculation #3381
- Fixed input validation for amount of queues supported by ALTQ schedulers #1353
- Added Google Stadia port range to the traffic shaper wizard #10743
- Fixed PHP errors in the traffic shaper wizard #10660
- Fixed ALTQ on `hn(4)` interfaces #8954

3.23.46 Traffic Shaper (Limiters)

- Fixed issues with `net.inet.ip.dummynet.*` tunables being ignored #10780
- Fixed issues with renaming limiters removing them from firewall rules #3924
- Fixed mask options not applying to sched limiter #10838
- Changed default Limiter queue bandwidth type to Mbit/s #10727

3.23.47 Translations

- Added Italian translation #9716

3.23.48 Upgrade

- Fixed issues with checking for updates from the GUI behind a proxy with authentication #9478
- Changed phrasing of message indicating the firewall is rebooting to upgrade #10387
- Fixed issues with the GUI incorrectly reporting "The system is on the latest version" #8870

3.23.49 UPnP

- Improved handling of UPnP with multiple gaming systems #7727

3.23.50 User Manager / Privileges

- Added menu entry for User Password Manager if the user does not have permission to reach the User Manager #9428
- Improved consistency of SSL/TLS references in LDAP authentication servers #10172
- Fixed irrelevant output being printed to users with `ssh_tunnel_shell` #9260
- Fixed theme not being applied to LDAP test results modal #7912
- Changed to more secure default values for certificates created through the user manager #11167
- Changed SSL/TLS LDAP authentication implementation to improve handling of multiple secure LDAP (SSL/TLS or STARTTLS) servers used at the same time #10704

3.23.51 Virtual IP Addresses

- Fixed a problem with PID file handling for the proxy ARP daemon #7379
- Fixed IP Alias VIPs on PPPoE interfaces #7132

3.23.52 Web Interface

- Updated JQuery to address multiple issues #10676
- Updated Bootstrap to 3.4.1 #9892
- Updated Font-Awesome to v5 #9052
- Increased the number of colors available for the login screen #9706
- Added TLS 1.3 to GUI and Captive Portal web server configuration, and removed older versions (TLS 1.0 removed from Captive Portal, TLS 1.1 removed from GUI) #9607
- Fixed empty lines in various forms throughout the GUI #9449
- Improved validation of FQDNs #9023
- Added `CHACHA20-POLY1305` to `nginx` cipher list #9896
- Fixed Setup Wizard input validation to allow Primary/Secondary DNS Server field to remain empty #10982
- Fixed Setup Wizard input validation for IPv6 DNS Servers #10720
- Added an option to omit DNS Servers from resolv.conf #10931
- Fixed the icon area within buttons not being clickable #10846
- Fixed visibility issues with multiple selection form control in the `pfsense-BETA-dark` theme #10705
- Updated documentation links in the GUI #10481
- Fixed netmask/prefix form control incorrectly resetting to `128/32` #10433
- Updated Help shortcut links #10135
- Improved handling of multiple login form submissions to avoid a potential CSRF error #9855
- Fixed reboot message when changing the Hardware Checksum Offloading setting #3031
- Added support for new site icons requested by current versions of Safari #11068
- Added descriptions to all `write_config()` calls #204

3.23.53 WireGuard

- Added kernel-based *WireGuard* VPN implementation #8786

3.23.54 Wireless

- Added support for the `athp(4)` wireless interface driver #9538 #9600
- Added support for the `ral(4)` wireless interface driver to arm64 #10934
- Added support for the `rtwn(4)` wireless interface driver #10639
- Added support for selecting 802.11n channel width (HT) #10678

3.23.55 Development

- Added a "periodic" style framework to allow for daily/weekly/monthly tasks from the base system or packages by way of plugin calls #7332
- Added a central file download function for internal use throughout the GUI
- Added TCP_RFC7413 in kernel, required for the BIND package #7293

3.23.56 XMLRPC

- Fixed XMLRPC synchronization of admin authorized keys for the admin user #9539
- Added option to synchronize changes for the account used for XMLRPC sync #9622
- Fixed XMLRPC synchronization for firewall rule descriptions with special characters #1478
- Fixed Incorrect synchronize IP address value causing XMLRPC errors #11017

3.24 2.4.5-p1 New Features and Changes

pfSense® software version 2.4.5-p1 addresses performance, security, and other miscellaneous issues found in *2.4.5*.

Warning: Proceed with caution when upgrading pfSense software while COVID-19 travel restrictions are in effect.

During this time of travel limitations, remote upgrades of pfSense software should be carefully considered, and avoided where possible. Travel restrictions may complicate any repair of any issue, including hardware-related issues that render the system unreachable. Should these issues require onsite physical access to remedy, repair of the issue may not be possible while travel restrictions related to COVID-19 are in effect.

Tip: For those who have not yet updated to 2.4.5-p1, consult the *previous release notes* and blog posts for those releases to read all important information and warnings before proceeding.

Note: Upgrading to pfSense software version 2.4.5-p1 requires `pfSense-upgrade` version 0.70 or later. Most installations will automatically pick up the new version and upgrade normally. If this does not happen automatically and the upgrade to version 2.4.5-p1 is not offered, use the following procedure:

- Navigate to **System > Updates**
- Set **Branch** to *Previous stable version*
- Wait a few moments for the upgrade check to complete
- Optional: Confirm that the latest version of *pfSense-upgrade* is present (version >= 0.70) using `pkg-static info -x pfSense-upgrade`.

  If the correct version is not present, wait a bit longer and check again as that package may be updating in the background.
- Set **Branch** to *Latest stable version*
- Wait a few moments for the upgrade check to complete

At this point, the upgrade check should see 2.4.5-p1 and the upgrade can proceed.

Note: pfSense software version 2.4.5-p1 includes `pkg` version 1.13.x which introduces a new metadata version. Most installations will automatically pick up the new version and upgrade normally. In certain cases, especially coming from much older versions, the `pkg` utility may require a manual update before it can correctly process the new metadata.

The `pkg` utility can be upgraded manually with the following command run from an ssh or console shell:

```
# pkg-static bootstrap -f
```

See *Repository Metadata Version Errors* for more details.

3.24.1 Security / Errata

- Addressed an issue with large PF tables causing system instability and high CPU usage during filter reload events #10414
- Fixed an issue with `sshguard` which could prevent it from protecting against brute force logins #10488
- Updated `unbound` to address CVE-2020-12662 and CVE-2020-12663 #10576
- Updated `json-c` to address CVE-2020-12762 #10609
- Addressed FreeBSD Security Advisories & Errata Notices including:
  - FreeBSD-SA-20:10.ipfw
  - FreeBSD-SA-20:12.libalias
  - FreeBSD-SA-20:13.libalias
  - FreeBSD-SA-20:15.cryptodev

3.24.2 Aliases / Tables

- Fixed handling of URL/URL table aliases with IDN hostnames #10321

3.24.3 Authentication

- Fixed handling of misconfigured groups which prevented the `admin` user from making configuration changes #10492
- Fixed a potential temporary privilege downgrade when deleting an account #9259

3.24.4 Backup / Restore

- Fixed handling of redundant/extraneous RRD tags when making configuration backups #10508

3.24.5 CARP

- Fixed handling of IPv6 CARP VIPs with non-significant zeros during XMLRPC sync #6579

3.24.6 Certificates

- Fixed a bug which prevented the user from removing a CA private key when editing #10509

3.24.7 Configuration Upgrade

- Fixed a PHP error during upgrade from <2.4.3 with empty tags in the IPsec configuration #10458

3.24.8 Console Menu

- Changed the naming convention of gateways created at the console to be the same as those created in the GUI #10264

3.24.9 DHCP (IPv6)

- Added default value placeholders to some DHCPv6 RA configuration options #10448
- Fixed DHCPv6 service Dynamic DNS errors #10346
- Fixed `rc.newwanipv6` being called for Request messages which `dhcp6c` should have discarded #9634
- Added dashed DUID support to DHCPv6 static mappings #2568

3.24.10 DHCP Relay

- Fixed DHCP Relay handling of scenarios where a target server may be on the same interface as some clients #10416
- Excluded unsupported interface types from DHCP Relay #10341

3.24.11 DHCP Server

- Fixed DHCPv6 static entries not being updated on external Dynamic DNS servers #10412
- Fixed DHCPv6 `domain-search` list not being sent to clients #10200
- Fixed DHCP Server not accepting IPv6 addresses for Dynamic DNS servers #6600

3.24.12 Diagnostics

- Several improvements and items added to status.php diagnostic output #10455 #10424 #10423 #10350 #10349 #10568
- Fixed Require State Filter setting on `diag_states.php` breaking filter rule link to associated states #10359

3.24.13 DNS Resolver

- Fixed IPsec and OpenVPN IPv6 tunnel network/pool prefixes not being added to automatic DNS Resolver ACLs #10460
- Fixed EDNS buffer size values to prepare for 2020 DNS flag day #10293
- Fixed DNS Resolver handling of entries from DHCP server which contain a trailing dot in domain names #8054

3.24.14 Dynamic DNS

- Fixed DigitalOcean Dynamic DNS client handling of IPv6 addresses #10390
- Fixed DNSExit update URL #9632

3.24.15 Hardware / Drivers

- Added support for `iwm` devices #7725

  Note: This device only supports Station mode. It does not support acting as an access point.

- Added `ng_etf` module to `armv6` and `aarch64` kernels #10463
- Added QLogic 10G driver (`qlxgb/qla80xx`) #9891
- Added `virtio_console` to the kernel #9985

3.24.16 IPsec

- Fixed selection of IPsec VTI Phase 2 local network address/mask values #10418
- Fixed saving IPsec connection breaking FRR BGP on VTI interfaces #10351
- Updated DH group warnings to say that group 5 is also weak #10221
- Fixed disabling IPsec Phase 1 with a VTI Phase 2 #10190
- Fixed disabled IPsec Phase 2 entries being unintentionally included in `vpn_networks` table #7622

3.24.17 L2TP

- Changed L2TP `mpd.secret` handling so that the server is not restarted after adding/modifying L2TP users #4866
- Fixed handling of L2TP usernames containing a realm separator (`@`) #9828
- Fixed Shared Secret handling in L2TP #10531 #10527

3.24.18 Limiters

- Fixed input validation of limiters with ECN #10211
- Fixed bogus extra warning dialog on when deleting limiters #9334

3.24.19 Notifications

- Fixed SMTP notification SSL validation to respect the user-selected behavior #10317

3.24.20 NTPD

- Added `localhost` to NTP Interface selection options #10348

3.24.21 OpenVPN

- Fixed OpenVPN `remote` statement protocol handling #10368
- Added option to configure OpenVPN username as common name behavior #8289

3.24.22 Operating System

- Fixed handling of RAM disk sizes not accounting for existing disk usage when calculating available kernel memory, which could prevent saving #10420
- Updated `pkg` to 1.13.x #10564
- Fixed problems preventing the Netgate Coreboot Package from updating Coreboot properly #10573

3.24.23 Packages

- Fixed handling of FreeRADIUS passwords containing non-XML-safe characters #4497
- Fixed handling of Squid LDAP search filters containing an accent #7654
- Fixed issues preventing FRR from working on certain platforms such as SG-1100 (arm64/aarch64) #10444
- Fixed issues preventing Suricata from working on certain platforms such as SG-1100 (arm64/aarch64) #10228

3.24.24 Rules / NAT

- Fixed Duplicate Outbound NAT entries from L2TP server addresses #10247
- Fixed Outbound NAT rules for mobile IPsec users with per-user addresses defined #9320
- Fixed IPv6 IP Alias VIPs not being added to Interface Network macros #8256
- Fixed Destination port range "Any" in Port Forward rules #7704
- Fixed display of interfaces on the Floating rules list #4629
- Fixed rule description validation to reject \ #10542
- Fixed setting NAT reflection timeout values #10591

3.24.25 Translations

- Fixed language selection for Chinese (Taiwan) / HK Translations #10525

3.24.26 Services

- Fixed `is_process_running()` handling of empty process, which could lead to an error when using the CLI to query the status of a service which does not exist #10540

3.24.27 Web Interface

- Fixed dark theme auto-complete popup field having dark text on dark background #10499
- Fixed using special characters in Schedule descriptions #10305
- Fixed WebGUI main page loading very slowly when there is no Internet connectivity #8987

3.25 2.4.5 New Features and Changes

pfSense® software version 2.4.5 contains a variety of bug fixes and maintenance updates.

Warning: Proceed with caution when upgrading pfSense software while COVID-19 travel restrictions are in effect.

During this time of travel limitations, remote upgrades of pfSense software should be carefully considered, and avoided where possible. Travel restrictions may complicate any repair of any issue, including hardware-related issues that render the system unreachable. Should these issues require onsite physical access to remedy, repair of the issue may not be possible while travel restrictions related to COVID-19 are in effect.

Tip: For those who have not yet updated to 2.4.5-p1, consult the *previous release notes* and blog posts for those releases to read all important information and warnings before proceeding.

3.25.1 Operating System / Architecture changes

- Base OS upgraded to FreeBSD 11.3-STABLE@r357046
- PHP upgraded to 7.2.29

3.25.2 Security / Errata

- Fixed dependency issues with `pfSense-upgrade` which may have caused it not to update itself properly #10303

Tip: If the update check fails, or the update does not complete, run `pkg install -y pfSense-upgrade` to ensure that `pfSense-upgrade` is present.

- Added encoding to the hostname in `services_acb.php` #9584
- Added encoding to error output in `services_captiveportal_mac.php` #9609
- Improved Picture Widget input validation #9610 #9731 #9804
- Added a `fsck` run with `-z` for `UFS` filesystems on upgrade to address FreeBSD-SA-19:10.ufs #9612
- Fixed format of XMLRPC auth error to match GUI auth error #9782
- Added a custom CSRF Error page with warnings and confirmation prompts before resubmitting potentially harmful data #9799
- Fixed Status_Monitoring `rrd_fetch_json.php` error encoding #9601
- Fixed encoding of the user full name on `system_usermanager_addprivs.php` #10324
- Fixed input validation and output encoding of host on `diag_ping.php` #10355
- Addressed FreeBSD Security Advisories & Errata Notices
  - FreeBSD-SA-20:05.if_oce_ioctl
  - FreeBSD-SA-20:04.tcp
  - FreeBSD-SA-19:24.mqueuefs
  - FreeBSD-SA-19:23.midi
  - FreeBSD-SA-19:22.mbuf
  - FreeBSD-SA-19:21.bhyve
  - FreeBSD-SA-19:20.bsnmp
  - FreeBSD-SA-19:19.mldv2
  - FreeBSD-SA-19:18.bzip2
  - FreeBSD-SA-19:17.fd
  - FreeBSD-SA-19:16.bhyve
  - FreeBSD-SA-19:15.mqueuefs
  - FreeBSD-SA-19:14.freebsd32
  - FreeBSD-SA-19:13.pts
  - FreeBSD-SA-19:12.telnet
  - FreeBSD-SA-19:11.cd_ioctl

- FreeBSD-SA-19:10.ufs
- FreeBSD-SA-19:09.iconv
- FreeBSD-SA-19:08.rack
- FreeBSD-EN-20:06.ipv6
- FreeBSD-EN-20:04.pfctl
- FreeBSD-EN-19:18.tzdata
- FreeBSD-EN-19:17.ipfw
- FreeBSD-EN-19:16.bhyve
- FreeBSD-EN-19:15.libunwind
- FreeBSD-EN-19:14.epoch
- FreeBSD-EN-19:13.mds
- FreeBSD-EN-19:12.tzdata
- FreeBSD-EN-19:11.net

3.25.3 Aliases/Tables

- Fixed an issue when resolving FQDN entries in aliases where some entries could be missing #9296
- Improved URL Table aliases to support FQDNs which return muliple entries #8531
- Added a function to download the contents of an individual alias #9816

3.25.4 Authentication

- Added exception handling to authentication attempts #9150

3.25.5 Backup/Restore

- Added a special string (`NoReMoToBaCkUp`) that when used in `write_config()` descriptions will prevent a remote backup #9693
- Removed legacy AutoConfigBackup options (there were no more active accounts using the retired legacy service) #9687 #9785
- Added CDATA protection to the `encryption_password` XML tag, which allows international characters to be used in that field #7186
- Added CDATA escape to more auth-related fields #9327
- Ensured that `kern.cam.boot_delay` is set for new installations and upgrades so that USB devices are properly initialized in time for configuration restore in the installer and ECL to function #9533

3.25.6 Captive Portal

- Fixed Captive Portal vouchers shortcut links #9722
- Changed Captive Portal redirect page selection order #9819
- Fixed a rare and intermittent issue where users could encounter an `nginx` error when restarting Captive Portal instances #10159

3.25.7 Certificates

- Added sorting and search/filtering to Certificate Authority & Certificate manager #9412
- Corrected wording of CA/Cert CN input validation #9234
- Fixed certificate Descriptive Name field behavior when adding a user certificate #9719
- Added `clientAuth` EKU to Server type certificates #9868
- Reduced the default GUI web server certificate lifetime to 398 days to prevent errors on Apple platforms #9825

3.25.8 Dashboard

- Added option to disable PTI display in System Information widget #9323

3.25.9 DHCP

- Fixed incorrect expansion of Dynamic DNS advanced options on the DHCPv6 Server page #9448
- Changed DHCP relay backend code to determine and specify separate upstream and downstream interface lists #9466
- Prevented OpenVPN interfaces from being used by DHCP relay, since that type of interface is not compatible #8443
- Added an option to disable ping check in dhcpd #9285
- Fixed **Show all configured leases** so it is persistent after deleting a DHCP lease #9133
- Added search/filter to DHCP/DHCPv6 leases #9791
- Improved DHCP client handling of timeout conditions and script failures #9267

3.25.10 Diagnostics

- Fixed a PHP warning in diag_dump_states.php #9780
- Fixed reverse lookup of IPv6 addresses on diag_dns.php #9543
- Fixed diag_system_activity.php to use batch mode for top so it displays process list w/o terminal, and increased amount of output displayed #9522
- Added search/filter ARP table and NDP status #9791

3.25.11 DNS

- Added `127.0.0.0/8` to the DNS Resolver `private-address` list for DNS rebinding protection #9708
- Fixed CIDR selection issues with /32 entries in DNS Resolver Access List entries #9586
- Fixed an issue saving DNS over TLS hostnames on systems with only one gateway #9898
- Fixed an issue where manually configured DNS servers may not have been active if "allow override" was disabled and they were also assigned dynamically #9963
- Added DNS Resolver (Unbound) Python Integration #9251

3.25.12 Dynamic DNS

- Fixed Dynamic DNS class constructor name #9779
- Fixed errors in DNSimple Dynamic DNS #9580
- Fixed handling of wildcard (*) hostname entries in Cloudflare Dynamic DNS #9361
- Added support for AAAA records to Digital Ocean Dynamic DNS #9280
- Fixed issues with Digital Ocean Dynamic DNS handling of empty hostnames #9602
- Cleaned up whitespace issues in Azure Dynamic DNS backend code #9271
- Added support for Linode Dynamic DNS #9268
- Fixed issues with IPv6 on Azure Dynamic DNS #9248
- Fixed handling of wildcards in Route53 Dynamic DNS #9053
- Fixed handling of wildcards in Loopia Dynamic DNS #8014
- Fixed CloudFlare Dynamic DNS processing when `proxied` is enabled #9362
- Fixed CloudFlare Dynamic DNS "Invalid TTL" error due to CloudFlare API update #10196
- Changed hostname to optional for DNS-O-Matic Dynamic DNS #7601
- Added support for Gandi LiveDNS Dynamic DNS #9452

3.25.13 Gateways

- Corrected PHP errors when marking gateways down in certain edge cases #9851

3.25.14 Interfaces

- Added more prefix delegation size entries to selection list on interfaces.php #9590
- Added initialization to the VLAN array in console setup #9582
- Fixed issues with Netgate & hardware model detection which caused problems with default interface mappings #8051
- Fixed issues with display of previously-entered IP address values on interfaces_ppps_edit.php #9741
- Added a confirmation prompt to disconnect/release actions on status_interfaces.php #9911
- Added drivers for Mellanox `mlx4` and `mlx5` network interface cards #7537

3.25.15 IPsec

- Fixed IPsec VTI interface creation logic #9781
- Added GUI option for IPsec P2/Child SA close action #9767
- Added IPsec DH and PFS groups 25, 26, and 27 #9757
- Added 25519 curve-based IPsec DH and PFS group 31 #9531
- Enabled NAT-T controls for IKEv2 #9695
- Improved handling of IPsec restarts breaking VTI routing #9668
- Fixed input validation that incorrectly prevented deleting IPsec P2 entries in some cases with VTI #9258
- Fixed IPsec `keyid` identifier handling #9243
- Fixed IPsec VTI MTU boot-time configuration #9111
- Escape Windows domain backslash in IPsec widget #9747
- Fixed VTI IPv6 address handling #9801
- Fixed Child SA button JS hide on status_ipsec.php, along with other cosmetic improvements #8847
- Added **Connect Children** button to status_ipsec.php to connect when IKE (Phase 1) is up but Child SAs (Phase 2 entries) are not #9954
- Fixed IPsec Phase 2 Remote Network field show/hide when changing between Phase 2 modes #9720
- Fixed IPsec configuration generation so that encryption options for every P2 on a given P1 are not duplicated on each P2 #6263
- Fixed a PHP error in IPsec package plugin hook processing #10217

3.25.16 Load Balancer

- Fixed a PHP when processing services when the configuration does not contain Load Balancer entries #10308

3.25.17 Logging

- Moved `igmpproxy` logs to `routing.log` #10139
- Moved `igmpproxy` verbose logging option to `services_igmpproxy.php` (formerly at `status_logs_settings.php`) #10139
- Updated `sshguard` and fixed a log processing regression #9971
- Fixed PHP errors in filter log processing when entries contain an invalid port #10255

3.25.18 Monitoring

- Fixed custom view titles being forced to lower case #9681
- Fixed packet graph scaling #9807
- Fixed a PHP error in RRD processing of ALTQ data #10248

3.25.19 Notifications

- Fixed SMTP notification password being unintentionally changed when testing SMTP settings #9684
- Reduced frequency of GEOM rebuild notifications #9256

3.25.20 NTPD

- Added validation to ensure NTP values are treated as numbers before use #9558
- Changed the default NTP pool server to `2.<domain>` so that it can use IPv6 #9931
- Improved handling of errors on the NTP status page to work/fail gracefully with custom ACLs for localhost in place #9829

3.25.21 OpenVPN

- Fixed JavaScript issue when selecting multiple OpenVPN NCP algorithms #9756
- Fixed OpenVPN wizard so it does not show DH parameter lengths that are not available #9748
- Fixed issues with OpenVPN resynchronizing when running on a gateway group #9595
- Added an option to set the OpenVPN TLS Key Direction #9030
- Added GUI options to configure OpenVPN keepalive parameters #3473
- Fixed instances of hidden invalid OpenVPN options affecting save operations #9674
- Added a copy action to OpenVPN pages #5851
- Improved sorting of bytes sent/receives on OpenVPN status page #7359
- Fixed visibility of the OpenVPN 'interface' option when multihome is selected #7840
- Reduced the OpenVPN server certificate lifetime to 398 days in the wizard to prevent errors on Apple platforms #9825
- Added input validation to prevent OpenVPN tunnel network reuse #3244
- Added Exit Notify to OpenVPN servers/client options #9078

3.25.22 Operating System

- Fixed serial console terminal size issues #9569
- Added the `strings` binary to base builds for troubleshooting #7791
- Changed UFS filesystem defaults to `noatime` on new installations #9483
- Fixed an issue where the IP header checksum was incorrect when reassembling packet fragments to a link with a different MTU #10189

3.25.23 Packet Capture

- Changed Packet Capture GUI to allow multiple TCP/UDP ports to be specified #9766
- Added start time to Packet Capture display #9831
- Added OSPF/OSPFv3 to Packet Capture protocols #9905
- Fixed Packet Capture to match both IPv4+IPv6 CARP when that protocol is selected #9867
- Fixed Packet Capture for the `pfsync` protocol #10183

3.25.24 Routing

- Fixed `(Default)` designation on routes to match the default route in the OS #9292
- Fixed static routes remaining in routing table after removal #9969

3.25.25 Rules / NAT

- Fixed state kill ordering in rc.newwanip #4674
- Added the ability to search firewall logs by tracking ID #8703
- Added GUI option to disable default blocking of APIPA networks #9966
- Added more common ports to the firewall rule drop-down list #10166
- Added input validation to prevent selecting `!*` (“not any”) in source or destination #10168
- Fixed invalid rules generated when using NAT reflection with a negated destination #10246

3.25.26 S.M.A.R.T.

- Updated the SMART page with new capabilities #9367

3.25.27 SNMP

- Fixed SNMP sysDescr contents to include hostname and patch version #9218

3.25.28 Traffic Shaping / Limiters

- Added input validation for Limiter delay values #9921
- Fixed the queue statistics parser to handle large values #9938

3.25.29 Translations

- Fixed an issue with international characters in configuration descriptions, which led to failures in certain cases, such as failing to set Manual Outbound NAT when the Language was set to pt_BR #6195
- Fixed a PHP error on `system_advanced_admin.php` when the language was set to French #10331

3.25.30 Upgrade / Installation

- Revised update check to provide a more consistent version string in JSON format #9778
- Disabled serial console on VGA memstick images #9488
- Fixed a PHP error when upgrading older configurations from revision 14.4 to 14.5 #9840

3.25.31 UPnP

- Fixed display of active UPnP sessions when configured with an alternate external address #9961

3.25.32 User Manager / Privileges

- Added input validation to prevent changing the authentication server name #9692
- Added privilege to manage integrated switches #9620
- Fixed privilege matching to handle JS anchor links #9550
- Removed wildcards incorrectly used in `isAllowedPage()` #9541
  - This issue could prevent a user in the admins group from reaching certain pages such as the User Manager.
- Improved Deny Config Write privilege handling in the User & Group Manager #9259
- Fixed input validation of group name sizes to allow longer remote groups #3792
- Fixed handling of L2TP and PPPoE user passwords containing invalid characters #10275

3.25.33 Web Interface

- Corrected input validation for firewall rule VLAN priority/set #9763
- Restricted Thoth tests to arm64 in status.php NG 2569
- Added kernel memory usage to status.php output #9705
- Redacted several additional fields in status.php output #9784 #9729 #9728 #9727 #9694 #9736 #9764
- Fixed a potential source of PHP errors when saving per-log settings #9540
- Added GUI components for MDS mitigation #9532
- Fixed integrated switch LAGG member editing on switch_ports.php #9447
- Fixed wizard.php selection option size attribute handling #8907
- Fixed platform detection for certain C2558/C2758 systems #6846
- Set `autocomplete=new-password` for forms containing authentication fields to help prevent browser auto-fill from completing irrelevant fields #9864
- Fixed processing of shortcuts for XML-based packages #9770
- Updated jQuery #9407
- Improved consistency of SSL/TLS references throughout the GUI #10172
- Updated various help references and links to use the pfSense book instead of external resources #10135 #10184

3.25.34 XMLRPC

- Fixed removal of the last ALTQ traffic shaping entry from the target system when performing an XMLRPC sync #9469
- Fixed removal of the last limiter entry from the target system when performing an XMLRPC sync #9468

3.26 2.4.4-p3 New Features and Changes

pfSense® software version 2.4.4-p3 addresses security and other issues found in *2.4.4-p2*.

Tip: For those who have not yet updated to 2.4.5-p1, consult the *previous release notes* and blog posts for those releases to read all important information and warnings before proceeding.

Warning: The upcoming pfSense release version 2.5.0 deprecates the built-in load balancer, and all related code has been removed as it is not compatible with FreeBSD 12. Plan migrations to alternate solutions such as the HAProxy package now.

See the *2.5.0 release notes* for more information.

3.26.1 Security / Errata

- Changed `sshguard` to block both ssh and the GUI using a single table, and removed the unnecessary manual scheduled table expiration pfSense-SA-19_02.sshguard #9223
- Fixed potential XSS vectors
  - pfSense-SA-19_01.webgui : Fixed potential XSS vectors in system_advanced_admin.php, interfaces_assign.php, firewall_rules_edit.php, firewall_shaper.php, services_igmpproxy_edit.php, services_ntpd_gps.php and diag_traceroute.php #9294
  - pfSense-SA-19_03.webgui : Fixed potential XSS vector in status_filter_reload.php #9499
  - pfSense-SA-19_04.webgui : Fixed potential XSS vector in the WOL widget #9507
  - pfSense-SA-19_05.webgui : Fixed potential XSS vector in services_acb.php #9508
- Fixed privilege issues
  - pfSense-SA-19_06.webgui : Restrict edit access to OpenVPN-related advanced settings, and added new privilege to delegate edit permissions #9511
  - pfSense-SA-19_07.webgui : Strengthen widget privilege matching to avoid a potential privilege bypass for users granted access to widgets #9512
  - pfSense-SA-19_08.webgui : Strengthen path privilege check to avoid a potential directory-traversal-like bypass method #9513
  - Added privileges for Auto Config Backup pages #9519
  - Updated privileges: Added misc missing pages, removed obsolete pages
- Addressed FreeBSD Security Advisories:
  - FreeBSD-SA-19:03.wpa
  - FreeBSD-SA-19:04.ntp
  - FreeBSD-SA-19:05.pf
  - FreeBSD-SA-19:06.pf
  - FreeBSD-SA-19:07.mds
  - FreeBSD-EN-19:08.tzdata
- Added DNS over TLS host verification #8602
  - Configure hostnames for DNS over TLS servers under **System > General**
- sqlite updates #9205

3.26.2 Backup / Restore

- Fixed issues with output buffering causing configuration backup download failures #9390
- Fixed automatic package reinstallation after restoring config.xml from the installer #9214
- Force <`enableserial`> when restoring a backup on a device with serial only console

3.26.3 Certificates

- Added missing countries from CA list on certificate pages #9308
- Fixed an error when adding a new user and choosing to generate a certificate #9317

3.26.4 DNS

- Fixed input validation on diag_dns.php to allow a trailing dot on hostnames #9276
- Removed non-functional tools links from diag_dns.php #9275
- Fixed rewriting of the DNS Resolver file `remotecontrol.conf` if it is present but empty #9470

3.26.5 Firewall Rules / NAT / Aliases

- Fixed intermittent pf errors when NAT reflection is enabled #9446
- Fixed reserved pf keyword matching when creating and editing aliases #9231
- Fixed duplicate entries showing on diag_tables.php from lockout tables #9359
- Fixed a PHP error deleting an imported NAT rule with no firewall rules present #9193
- Do not show scheduler icon when scheduler tag is empty

3.26.6 Gateways / Routing

- Fixed issues with the default IPv4 gateway set to a group failing after restart #9004

3.26.7 Interfaces

- Fixed PHP error from interface groups when editing QinQ entries

3.26.8 IPsec

- Fixed IPsec Phase 1 entries on upgrade to have their `protocol` field populated properly #9207

3.26.9 Operating System

- Fixed support for ZFS encrypted+mirrored swap #9281
- Fixed problems saving crash dumps when `/var` is a RAM disk #9409

3.26.10 Traffic Shaping

- Fixed a PHP error when loading a limiter that does not exist #9313
- Fixed limiter selection validation
- Fixed Queues menu items ending with ":" in certain languages #8970

3.26.11 WebGUI

- Numerous optimizations and improvements for status.php diagnostics output #9290
- Fixed a PHP error on system_advanced_network.php when disabling "IPv6 over IPv4 Tunneling" #9264
- Improved handling of large captures on diag_packet_capture.php and disabled viewing of captures larger than 50MiB. #9239
- Added hostname to login page title if the user has enabled **Show hostname on login banner** #9096
- Centralized the list of country codes used by multiple areas #9308
- Updated translation files

3.26.12 XMLRPC

- Clarified conditions for synchronizing certificates in HA Sync options #9283

3.27 2.4.4-p2 New Features and Changes

pfSense® software version 2.4.4-p2 adds support for new Netgate hardware and corrects issues found with *2.4.4-p1*.

Warning: For those who have not yet updated to 2.4.4-p1 or 2.4.4, consult the release notes and blog posts for those releases to read all important information and warnings before proceeding.

3.27.1 Miscellaneous

- Hardware support/improvements for Netgate products
- Fixed swap slice labeling in MBR mode and changed the way swap is located at boot time to detect and work around incorrect `fstab` swap labels created by the installer #9182
- Fixed handling of IPv6 name servers with nginx when using a certificate that requires OCSP stapling #9160
- Fixed handling of NPt rules using a /128 prefix #9163
- Fixed a PHP error in the Setup Wizard when dealing with static gateways #9170
- Updated Dynamic DNS to accommodate recent changes in the Digital Ocean API #9171
- Fixed OpenVPN RADIUS authentication use of `calling_station_id` #9178
- Fixed input validation that rejected certain valid hash algorithms when signing a CSR #9180
- Removed obsolete and unused OLSRD code #9117

3.28 2.4.4-p1 New Features and Changes

pfSense® software version 2.4.4-p1 corrects issues found with 2.4.4-RELEASE.

3.28.1 Security / Errata

- FreeBSD Errata Notice FreeBSD-EN-18:09.ip: IP fragment remediation causes IPv6 fragment reassembly failure #8934
- FreeBSD Errata Notice FreeBSD-EN-18:10.syscall NULL pointer dereference in freebsd4_getfsstat system call (CVE-2018-17154)
- FreeBSD Errata Notice FreeBSD-EN-18:11.listen Denial of service in listen syscall over IPv6 socket (CVE-2018-6925)
- FreeBSD Errata Notice FreeBSD-EN-18:12.mem Small kernel memory disclosures in two system calls (CVE-2018-17155)
- Fixed a potential authenticated command injection issue with PowerD settings pfSense-SA-18_09.webgui #9061
- Fixed handling of privileges on the **All** group that were previously ignored #9051

> **Warning:** Check the privileges on the **All** group before upgrading to avoid unintended privileges for accounts being respected that were not honored before

3.28.2 Certificates

- Fixed CRL lifetime errors due to 2038 rollover on 32-bit ARM platforms #9098
- Fixed date display of CA/Certificate validity ending dates after 2038 rollover on 32-bit ARM platforms #9100
- Fixed PHP errors when creating certificate entries #9099

3.28.3 DNS

- Updated Unbound to 1.8.1 to address issues with memory leaks, especially in DNS over TLS support #9059
- Fixed issues with the DNS search domain for the firewall being omitted from `resolv.conf` in certain cases #9056
- Fixed PHP errors in the DNS Forwarder #8967

3.28.4 Dynamic DNS

- Fixed an issue with FreeDNS Dynamic DNS sending an IP address with an update #8924
- Fixed issues with Custom (v6) Dynamic DNS logging a hostname error #8977

3.28.5 DHCP Server

- Fixed issues with DHCPv6 network boot settings #8949

3.28.6 Routing/Gateways

- Reduced the logging output of gateway change events #8914
- Fixed an issue with `dpinger` PID files causing it to get stuck in Pending status #8921
- Fixed display of a configured gateway monitor IP address when gateway monitoring is disabled #8953
- Fixed issues with double quotes in gateway descriptions causing a blank gateway drop-down on firewall rules #8962
- Fixed an issue where the default gateway was lost in certain cases with HA after a CARP VIP status transition #8465

3.28.7 IPsec

- Updated strongSwan to 5.7.1 #8898
- Added `0.0.0.0/0` to both sides of an IPsec VTI P2 to allow connections with third-party routed IPsec implementations that require its presence #8859
- Fixed boot-time handling of IPsec VTI static routes #9116
- Fixed IKEv2 EAP Identity/Client ID matching so that it is strictly performed, to avoid users getting incorrect per-user settings #9055
- Fixed handling of RADIUS server names containing a `.` in the IPsec configuration with strongSwan 5.7.1 #9106
- Updated AWS IPsec wizard to use EC2 instance profiles and security groups, and switched the wizard from OpenBGPD to FRR

3.28.8 Interfaces/VIPs

- Fixed issues with DHCP client MTU causing interface configure loops when advanced options are present #8507
- Fixed issues with the Hyper-V `hn(4)` driver and ALTQ #8954
- Fixed issues with Hyper-V `hn(4)` interfaces dropping UDP6 traffic when transmit checksums were enabled #9019
- Fixed an issue with IGMP proxy failing to start on PPPoE interfaces #8935
- Fixed an issue with IPv6 Transmit checksums not being disabled when hardware checksums were set to be disabled #8980
- Updated mpd to 5.8_8 to address issues with Orange MTU #8995
- Fixed PPPoE service name checks to allow `:` and other alphanumeric characters #9002
- Fixed PHP errors when creating QinQ entries #9109
- Fixed the MAC address shown when editing a LAGG entry to always show the hardware MAC for each NIC and not the currently active address, which is no longer accurate for LAGG members #8937
- Fixed a PHP error when setting an interface address to act as a DHCP server from the console, when no other DHCP servers are already configured #9144

- Fixed a situation where editing a VLAN interface caused all other VLAN interfaces with the same parent to be reconfigured, which led to several other issues #9115

> **Warning:** Editing a VLAN parent interface can still cause problems. If this becomes an issue on a firewall, consider moving from using the untagged parent to having that traffic be tagged so that the parent interface is not assigned or in use. #9154
>
> Known issues include:
>
> – PPPoE instances on VLANs will not reconnect after the interface is reconfigured #9148
>
> – VLAN interfaces that use IPv6 tracking may lose their addresses #9136

3.28.9 Hardware/Platform

- Fixed handling of EFI console when a device boots from UEFI, where `vidconsole` is not valid #8978
- Fixed PHP errors in switch configuration on platforms including integrated switches
- Added support for SG-5100 hardware watchdog

Note: Enable the Watchdog daemon under **System > Advanced** on the **Miscellaneous** tab, and then reboot and enable it in the BIOS with a timeout longer than the timeout configured in the GUI.

3.28.10 User Management / Authentication

- Fixed handling of privileges on the **All** group that were previously ignored #9051

> **Warning:** Check the privileges on the **All** group before upgrading to avoid unintended privileges for accounts being respected that were not honored before

- Added GUI options to control `sshguard` sensitivity and whitelisting to allow users to fine-tune the behavior of the brute force login protection #8864
- Added an option to enable SSH agent forwarding (disabled by default) #8590
- Fixed inconsistencies with ssh settings in the configuration #8974
- Fixed PHP errors with ssh settings #8606
- Added support for LDAP client certificates on authentication servers (Factory only) #9007
- Fixed an issue with **Local Database** authentication when using non-English languages in certain cases, such as with Captive Portal #9086

3.28.11 Captive Portal

- Fixed Captive Portal RADIUS NAS Identifier default values to include the zone name #8998
- Restored the ability to set a custom NAS Identifier on Captive Portal RADIUS settings #8998
- Fixed issues with Captive Portal logout popup #9010
- Fixed handling of the login page displayed when RADIUS MAC Authentication fails #9032
- Fixed username sent in RADIUS accounting with MAC-based authentication #9131
- Fixed an issue with the blocked MAC address redirect URL #9114

3.28.12 WebGUI / Dashboard

- Fixed `nginx` restart handling when toggling GUI web server options under **System > Advanced**, **Admin Access** tab
- Fixed empty crash reports after upgrade #8915
- Added CDATA protection to common name fields so they can safely contain international characters #9006

3.28.13 Firewall Rules / Aliases / NAT

- The `filterdns` daemon has been rewritten, solving a number of issues with the old implementation, including:
  - Fixes `filterdns` triggering every 16 seconds even when DNS records have not changed #7143
  - Fixes invalid FQDN entries in aliases causing an alias table to fail silently #8001
  - Fixes `filterdns` failing on a regular basis #8758
- Fixed `/etc/rc.kill_states` not correctly parsing `pfctl` output #8554
- Fixed formatting of alias names to still wrap but not replace underscores #8893
- Fixed PHP errors from `filter_rules_sort()` when a configuration contains no rules #8993
- Fixed PHP errors when creating schedules #9009
- Fixed PHP errors when creating entries on NAT pages #9080
- Fixed PHP errors from `easyrule` when no aliases are present #9119
- Fixed "Drag to reorder" description in rule list when rule drag-and-drop is disabled #9128

3.28.14 Traffic Shaping (ALTQ/Limiters)

- Fixed issues with Limiter queue display on upgraded configurations #8956
- Fixed the default limiter scheduler to match previous version (WF2Q+) #8973
- Added scheduler information to the limiter information page #8973

3.28.15 Packages

- Fixed issues with package installation causing problems when crossing major PHP versions #8938
- Fixed PHP errors when installing packages #9067

3.28.16 Backup/Restore

- Added schedule (cron) support to AutoConfigBackup #8947
- Fixed issues with AutoConfigBackup restoring a configuration from a different host #8901
- Fixed the AutoConfigBackup menu from the deprecated package still showing when the package is no longer present #8959
- Fixed an issue with **Reinstall Packages** hanging when run from **Diagnostics > Backup & Restore** #8933
- Fixed issues with multiple <rrddata> tags in `config.xml` #8994
- Fixed a race condition in package operations after a configuration restore that could lead to no packages being reinstalled #9045
- Fixed issues with the External Config Locator not finding a `config.xml` in `/config` #9066
- Fixed an issue where packages may not be reinstalled during a configuration restore performed immediately after a fresh install #9071
- Fixed a `stream_select()` error when restoring packages #9102

3.28.17 Wake on LAN

- Fixed issues with ordering of entries in Wake on LAN #8926
- Added top control buttons to Wake on LAN for **Add** and **Wake all Devices** when there are more than 25 entries #8943

3.28.18 NTP

- Fixed issues with NTP status when using `noquery` in the default permissions along with a specific ACL for localhost #7609

3.28.19 Logging / Notifications

- Fixed an issue with log file sizes >= `2^32/2` #9081
- Fixed PHP errors when saving log settings #9095
- Added a checkbox to disable TLS certificate verification for SMTP notifications #9001

3.28.20 Install/Upgrade

- Added a FAT partition to the installer memstick to make it easier to restore a `config.xml` file during the install process. Also includes a copy of the license and a README. #9104
- Fixed PHP errors in upgrade code for IPsec #9083

3.28.21 Miscellaneous

- Fixed HTTPS proxy authentication support for connections on the firewall itself #9029
- Clarified wording of **Kernel PTI** options on **System > Advanced**, **Miscellaneous** tab #9026
- Added a Save button to **Status > Traffic Graphs** to store default settings to use when loading the page #8976
- Added support for `nvme` controllers to the S.M.A.R.T. diagnostics page #9042

3.29 2.4.4 New Features and Changes

3.29.1 Significant Changes

OS Upgrade
Base Operating System upgraded to FreeBSD 11.2-RELEASE-p3. As a part of moving to FreeBSD 11.2, new hardware support is included for C3000-based hardware.

PHP 7.2
PHP upgraded to 7.2, which required numerous changes to syntax throughout the source code and packages.

Routed IPsec (VTI)
Routed IPsec is now possible using FreeBSD `if_ipsec(4)` Virtual Tunnel Interfaces (VTI). #8544 (See also: *Routed IPsec (VTI)*)

IPsec Speed Improvements
The new **Asynchronous Cryptography** option under the IPsec **Advanced Settings** tab can dramatically improve IPsec performance on multi-core hardware #8772

Default Gateway Group
The default gateway may now be configured using a Gateway Group setup for failover (each gateway on a different tier), which replaces Default Gateway Switching. #8187

Limiter AQM/Queue Schedulers
Limiters now include support for several Active Queue Management (AQM) methods and Queue Scheduler configurations such as FQ_CODEL. #6620 (See also: pfSense PR #3941)

Certificate Subject Requirements
The Certificate Manager and OpenVPN wizard now only require the **Common Name** to be set, and all other fields are optional. #8381

AutoConfigBackup is free!
AutoConfigBackup now integrated and free for all to use. (See also: *Automatic Configuration Backup Service*)

DNS over TLS
The DNS Resolver now includes support for DNS over TLS as both a client and a server, including for domain overrides. #8388 #8030 #8431

Captive Portal Authentication
Captive Portal authentication is now integrated with the User Manager system. Captive Portal instances may now use RADIUS, LDAP, or Local Authentication like other integrated services. The firewall will migrate existing Captive Portal RADIUS settings to the User Manager automatically on upgrade.

Captive Portal HTML Design and Usability
The default Captive Portal page has been redesigned. Controls have also been added which allow for the logo and background images and Terms of Service text to be customized without editing and uploading custom HTML code. #8793

Integrated Switch Improvements
Netgate devices with integrated switches such as the SG-3100 and XG-7100 can now configure per-port speed and duplex settings, discrete port configuration interfaces can now be tied to switch ports for up/down status, and LAGG support is also now available (Load Balance mode only)

3.29.2 Security

- FreeBSD SA for CVE-2018-6922: Resource exhaustion in TCP reassembly FreeBSD-SA-18:08.tcp
- FreeBSD SA for CVE-2018-3620, CVE-2018-3646: L1 Terminal Fault (L1TF) Kernel Information Disclosure FreeBSD-SA-18:09.l1tf
- FreeBSD SA for CVE-2018-6923: Resource exhaustion in IP fragment reassembly FreeBSD-SA-18:10.ip
- FreeBSD SA for CVE-2018-14526: Unauthenticated EAPOL-Key Decryption Vulnerability FreeBSD-SA-18:11.hostapd
- FreeBSD SA for CVE-2018-6924: Improper ELF header parsing FreeBSD-SA-18:12.elf
- FreeBSD errata notice for LazyFPU remediation causing potential data corruption FreeBSD-EN-18:08.lazyfpu
- Fixed a potential XSS vulnerability via GUI rule separators pfSense-SA-18_06.webgui #8654
- Fixed a potential XSS via custom GUI/dashboard settings pfSense-SA-18_07.webgui #8726
- Fixed a potential authenticated ACE vulnerability pfSense-SA-18_08.webgui #8843
- Upgraded strongSwan to 5.6.3 to address a buffer underflow leading to denial of service (CVE-2018-5388) #8746
- Updated default cryptographic settings for OpenVPN, IPsec, and Certificates #8594
- Changed the included DH groups to those defined in RFC 7919 #8582
- Added stronger IPsec Pre-Shared Key usage warnings, and a button to generate a secure PSK #8667
- Disabled OpenVPN compression by default on new instances for security reasons due to VORACLE – Users should strongly consider disabling compression on OpenVPN instances if they pass unencrypted data such as HTTP to arbitrary Internet sites #8788
- Patched OpenSSH for CVE-2018-15473, username enumeration/disclosure through malformed packets.
- Changed from `sshlockout_pf` to `sshguard` for monitoring failed logins and locking out offenders, this allows the lockout to work on IPv4 and IPv6 and also terminates states when adding offenders to the block list #7694 #7695

3.29.3 Errata

Warning: Third party packages from **alternate repositories** are causing problems for users with the upgrade process and also with post-upgrade behavior. These packages have never been supported, and had to be manually added by users outside of the GUI.

Due to the major changes required for FreeBSD 11.2 and PHP 7.2, third party packages from alternate repositories cannot be present during the upgrade. There is no way to predict if a third party package is compatible with the new version or will cause the upgrade itself to fail.

The upgrade process will automatically remove `pfSense-pkg-*` packages installed from alternate repositories. After the upgrade completes, the user can reinstall these packages. Packages from **alternate repositories** will not appear in the **Installed Packages** list in the GUI, and must be entirely managed in the command line.

This change does not affect packages installed from the official pfSense® package repository.

- Removed options for the deprecated FEC LAGG Protocol #8734

3.29.4 Certificates

- Changed the Certificate Manager and OpenVPN wizard to only require the **Common Name** for the CA/Cert subject #8381
- Updated default cryptographic settings Certificates #8594
- Added support for OCSP Must-Staple certificates in the GUI (and ACME package) #8418
- Changed CRL support from using an abandoned PHP OpenSSL module patch to a pure PHP implementation compatible with PHP 7.2 #8762
- Fixed issues with several areas not properly parsing CA fields properly when they were not in the expected order #8801
- Changed the default CA and Certificate create action from "Import…" to "Create an internal…" #8851

3.29.5 DNS

- Added DNS over TLS for upstream forwarders to the DNS Resolver #8388
- Added DNS over TLS server support to the DNS Resolver #8030
- Added DNS over TLS options for DNS Resolver Domain Override #8431
- Fixed editing DNS Resolver ACLs in non-English languages #8539
- Added a DNS Resolver status page #8430
- Clarified that "Register DHCP leases in the DNS Resolver" only works for IPv4 addresses #8592
- Added IPv6 representation of IPv4 addresses in DNS Resolver DNS Rebinding checks #8750
- Fixed disabling the DHCP Server on interfaces when the DNS Resolver **DHCP Registration** option is enabled (Only one enabled interface is required) #8120
- Added advanced option for `qname-minimization` to the DNS Resolver #8028
- Fixed an issue with IDs when editing or deleting DNS Forwarder host override entries #8767

3.29.6 Dynamic DNS

- Added Dynamic DNS client for DigitalOcean DNS #8478
- Fixed Dynamic DNS clients usage of custom check IP services #8664
- Added Dynamic DNS client for Azure #7769
- Updated DNSimple Dynamic DNS client to use DNSimple API v2 #8071
- Fixed handling of username and password fields for custom Dynamic DNS entries #8782

3.29.7 Routing/Gateways

- Added the ability to set a Gateway Group as the default gateway. #3781 #8187
- Extended the maximum Gateway monitoring **Probe Interval** #8593
- Fixed handling of Gateway Group **Trigger Level** #8586
- Fixed inconsistency in display and usage of units for Gateway latency #8477
- Upgraded FRR to 5.0.1 for compatibility with FreeBSD 11.2 #8449
- Fixed FRR BGP MD5 support #8407
- Fixed handling of Router Advertisement preferences #6237

3.29.8 IPsec

- Added routed IPsec using FreeBSD `if_ipsec(4)` VTI #8544
- Added a GUI option to the IPsec **Advanced Settings** tab for Asynchronous Cryptography which can dramatically improve IPsec crypto operation performance on multi-core hardware #8772
- Added IPsec identifiers to **Status > IPsec** #8598
- Fixed a JavaScript variable issue in IPsec IKE Phase 1 causing the Key Length field to be blank in some browsers such as IE #8543
- Added IPsec mobile client options to configure different (virtual) IP addresses per user #8292
- Added IPsec mobile client options to configure different DNS servers per user #8644
- Updated default cryptographic settings for IPsec #8594
- Changed the default behavior of an IPsec Phase 1 to rekey as needed #8540
- Fixed handling of per-user IPsec rules from an authentication server #8765
- Added warnings and hints to IPsec encryption and hash choices about potentially insecure selections #8766
- Fixed an issue with handling IP Alias VIPs with CARP parent after an interface up/down event #8768

3.29.9 OpenVPN

- Disabled compression by default for new OpenVPN client and server instances for security reasons #8788
- Changed OpenVPN Authentication to use an asynchronous authentication plugin which avoids stalling server traffic during the authentication process, especially noticeable on down/broken authentication servers #7905
- Fixed display of **Bridge Route Gateway** options on OpenVPN tap bridge servers #8658
- Fixed handling of LDAP fields in the OpenVPN wizard and brought the options in line with current LDAP server options #8605
- Updated default cryptographic settings for OpenVPN #8594
- Added missing OpenVPN compression options (`stub-v2` and plain `compress`) #8788

3.29.10 DHCP Server

- Fixed validation of custom DHCP options #8534
- Fixed a situation where DHCPv6 was configured for LAN when the LAN interface was not assigned #8048
- Fixed an issue with XMLRPC synchronization of DHCP static mappings #8721

3.29.11 Interfaces / VIPs

- Removed IPv4 and IPv6 settings from the Interface configuration for assigned OpenVPN/GIF/GRE/Routed IPsec instances, since the IP addresses are managed by the parent config not interfaces.php #8687
- Fixed an HTTP_REFERER issue when changing the LAN IP address in the Setup Wizard #8524
- Fixed an HTTP_REFERER issue when changing an interface IP address while accessing the GUI from the same interface #8822
- Fixed handling of the FreeBSD 11.2-BETA dhclient MTU value #8507
- Added PPPoE multi-link over single link to allow users with a supported provider to have a larger MTU #8737
- Fixed a PPPoE MTU issue with ORANGE FR #8595
- Fixed QinQ interface assignment #8446
- Fixed radvd/IPv6 when using a LAN bridge #8429
- Fixed deleting IP Alias VIPs outside an interface subnet where a gateway exists in the same subnet #4438
- Fixed handling of IP Alias and CARP VIP subnet mask/prefix autodetection #8741
- Fixed a panic in IPv6 fragment logging #8499
- Fixed handling of DHCP option 77 in the DHCP client #7425
- Fixed deleting Interface Group members which are disabled #8800
- Fixed MAC address spoofing for bridge interfaces #8138
- Fixed an issue with string termination when creating interfaces through the pfSense PHP module #8683
- Fixed an issue where changing a LAGG could cause a VLAN using that LAGG as a parent interface to lose its association with the LAGG #8527

3.29.12 Integrated Switches

- Added GUI controls to configure LAGG on integrated switch ports (Load Balance mode only)
- Added GUI controls to configure Speed/Duplex for switch ports on integrated switches
- Added the ability to tie the status of an assigned VLAN interface to a switch port for integrated switches
- Added Switch Status to status.php for platforms with a switch #8525
- Fixed an issue switching between Port VLAN and 802.1q VLAN mode on integrated switches #8422
- Fixed an SNMP error on hardware with integrated switches #8600
- Added **Preserve Switch Configuration** option when restoring `config.xml` to keep the current active switch settings instead of those from the imported configuration to help with hardware transitions

3.29.13 Hardware/Platform

- Added support for the new SG-5100
- Fixed an issue with ARM hardware not completely halting when shut down (SG-3100 and SG-1000)
- Fixed HDMI hotplug issues on Minnowboard Turbot hardware (MBT-2220 and MBT-4220)
- Fixed SG-1000 autonegotiation for 10baseT speed and duplex #7532

3.29.14 User Management / Authentication

- Added a visible warning to the user when default password has not been changed #8596
- Fixed configuration descriptions user management operations and added logging #8548
- Fixed escaping of LDAP search parameters #8626
- Fixed an OS issue with adding a group to a user when creating the user #8553
- Fixed handling of LDAP bind credentials #8583
- Removed some legacy code from `auth.inc` #8742
- Fixed Group selections after an input error in the User Manager #8622
- Fixed inconsistent usage of `sshdkeyonly` in `system_advanced_admin.php` #8403
- Added SSH configuration option to require **both** Key **and** Username+Password authentication at the same time #8402
- Replaced `radius.inc` by pear-Auth_RADIUS #7024
- Fixed synchronization of User Manager group scope and operating system groups #7013
- Fixed logging and display of GUI user authentication source IP address when the user logs in through a proxy #8813
- Fixed logging and display of GUI user authentication sources to show what source authorized the login (e.g. LDAP, RADIUS, Local, Fallback) #8816

3.29.15 Captive Portal

- Integrated Captive Portal authentication into the User Manager to enable support for LDAP #5112
- Updated Captive Portal HTML/CSS to a modern design and added controls to customize images and ToS without uploading custom HTML #8793
- Fixed deleting **Allowed Hostnames** and **Allowed IP Addresses** entries in Captive Portal when a zone is disabled #8530
- Added support for setting Captive Portal traffic quotas #8202
- Added display of a custom username when Captive Portal is set to *None* for the authentication type #8361
- Changed handling of Called-Station-Id/Calling-Station ID to send a MAC address instead of IP address when using RADIUS authentication #4294
- Changed to a standardized NAS-Identifier when using RADIUS authentication #3686
- Corrected accounting updates not being sent when expected #8655
- Fixed an issue with XMLRPC synchronization of Captive Portal settings #8806

3.29.16 WebGUI / Dashboard

- Enabled HTTP2 for the Web GUI server #8552
- Updated the text and links in the HTML footer #8733
- Fixed display of available swap with multiple swap disks in the **System Information** Dashboard widget #8587
- Updated text in the Setup Wizard #8753
- Moved the simplepie RSS reader code to a FreeBSD port for easier updates #6998
- Fixed handling of the **Inverse** option in the Traffic Graphs Dashboard Widget #8367
- Fixed issues with the GUI following upgrade progress #8519
- Added a line to display the current GUI user viewing the Dashboard in the System Information Widget #8817

3.29.17 Firewall Rules / NAT / Shaping

- Added CoDel, FQ-CoDel, PIE and FQ-PIE AQMs to limiters #6620
- Fixed firewall ruleset errors related to VIPs and outbound rules #8518 #8408
- Added validation for IPv6 NPt input #8575
- Fixed a race condition in NAT reflection filter rules that could lead to a ruleset load failure #8604
- Fixed viewing the list of Port Forwards when a user only has the "WebCfg - Firewall: NAT: Port Forward" privilege #8563
- Fixed an issue with default field selection when editing Firewall Rules #8597
- Added code to prevent nested alias loops #8101
- Added interface groups support for NAT rules #1933
- Fixed a case where invalid IPv6 NAT rules could be generated #8437
- Fixed a case where IPv6 Neighbor Discovery and other similar valid messages sent from the unspecified address (`::`) were not allowed by default #8791

- Added **Select All** functionality to firewall and NAT rules #8812
- Fixed IPv6 address form field format tooltip #8834

3.29.18 Packages

- Fixed situation where the firewall would get stuck attempting to reinstall packages after restoring a configuration when there is no Internet connection #7604
- Added a new tag for package services, `<starts_on_sync/>`, to allow packages to declare that they start themselves during the sync process, which lets packages opt out of a (second) forced start at boot and during interface events #8850

  See also: #8620

3.29.19 Miscellaneous

- Fixed display of stored Load Balancer custom settings #8704
- Fixed handling of `loader.conf` and `loader.conf.local` so it will not remove customized options that override defaults #8571
- Fixed the restoration process for a `config.xml` from USB during install to remove RRD data so that the data does not indefinitely stay in `config.xml` #7634
- Fixed handling of special characters in L2TP user passwords #7623
- Fixed handling of sample bounds with custom timer periods on **Status > Monitoring** #6477
- Changed the crash reporter so that users can download the reports locally rather than submitting to a server #8764
- Added more redacted XML tags to status.php #8819
- Changed status.php to use `ifconfig -va` to show more detail, including attached SFP devices with certain network interface drivers #8860

3.30 2.4.3-p1 New Features and Changes

New features and changes for this release of pfSense® software:

3.30.1 Security / Errata

- FreeBSD SA for CVE-2018-8897 FreeBSD-SA-18:06.debugreg
- FreeBSD EN for CVE-2018-6920 and CVE-2018-6921 FreeBSD-EN-18:05.mem
- Fixed a potential LFI in pkg_mgr_install.php pfSense-SA-18_04.webgui #8485
- Fixed a potential XSS in pkg_mgr_install.php pfSense-SA-18_05.webgui #8486

3.30.2 Misc

- Added a check to avoid creating route-to rules for proxy ARP addresses
- Corrected alias name input validation text referring to well-known and registered ports #8409
- Corrected the list of pf reserved keywords to prevent aliases from using invalid custom names #8445
- Fixed an issue with Captive Portal access rules being left behind on disconnect #8441
- Fixed an issue with pressing Enter in the filter field of diag_pftop.php #8494
- Fixed an issue with invalid rules generated due to the presence of IPv6 Alias VIPs #8408
- Fixed an issue with IPsec mobile Pre-Shared Keys and iOS devices #8426
- Fixed an issue with selecting a gateway when switching a firewall rule away from IPv4+IPv6 mode #8447
- Fixed firewall rules generated by the OpenVPN wizard #8391
- Fixed handling of OpenVPN RADIUS attribute firewall rules #8480
- Fixed handling of XMLRPC user/group synchronization when that section is disabled on the primary #8450
- Fixed input validation to allow named services to be used in firewall rules rather than numbers alone #8410
- Fixed issues with IP alias VIPs on Localhost at boot time #8393
- Increased the default Firewall Maximum Table Entries value to 400000 to cope with the increased size of the IPv6 bogon address lists #8417
- Updated SimplePie RSS to 1.5.1 #8423
- Added more fields to the list that status.php uses to redact private information #8394

3.31 2.3.5-p2 New Features and Changes

New features and changes for this release of pfSense® software:

3.31.1 Security / Errata

- FreeBSD SA for CVE-2018-8897 FreeBSD-SA-18:06.debugreg
- FreeBSD EN for CVE-2018-6920 and CVE-2018-6921 FreeBSD-EN-18:05.mem
- Fixed a potential XSS vector in RRD error output encoding #8269 pfSense-SA-18_01.packages
- Fixed a potential XSS vector in diag_system_activity.php output encoding #8300 pfSense-SA-18_02.webgui
- Fixed a potential LFI in pkg_mgr_install.php #8485 pfSense-SA-18_04.webgui
- Fixed a potential XSS in pkg_mgr_install.php #8486 pfSense-SA-18_05.webgui
- Changed sshd to use delayed compression #8245
- Added encoding for firewall schedule range descriptions #8259

3.31.2 Misc

- Added an option to disable HSTS for the GUI web server #6650
- Added filtering to pfTop page
- Added ospf6d to the routing log
- Change get_interface_subnet() to use configured value if available
- Corrected sethelp call on firewall_rules_edit.php #8242
- Fixed an issue with selecting a gateway when switching a firewall rule away from IPv4+IPv6 mode #8447
- Fixed an issue with the address family selection for remote syslog servers using IPv6 #8323
- Fixed a problem when IPsec bypasslan was enabled while the LAN interface is disabled or doesn't have an IP address #8239
- Fixed config.xml corruption handling
- Fixed input validation for Certificate SAN values to disallow IP addresses for FQDN/Hostname entries`#8275 <https://redmine.pfsense.org/issues/8275>`__
- Fixed issues with OpenVPN when using a /31 IPv4 Tunnel Network #8261
- Fixed NTP Status server time for zones with minute offsets (fractions of an hour) #8129
- Fixed selection of IPv6 gateways when creating a new firewall rule #8053
- Fixed various pf "busy" errors when the ruleset is reloaded
- Improved handling of aliases that mix IP addresses and FQDNs #8290
- Improved update repository controls
- Increased the default Firewall Maximum Table Entries value to 400000 to cope with the increased size of the IPv6 bogon address lists #8417

3.32 2.4.3 New Features and Changes

New features and changes for this release of pfSense® software:

3.32.1 Security / Errata

- FreeBSD-SA-18:01.ipsec
- Kernel PTI mitigations for Meltdown (optional tunable) FreeBSD-SA-18:03.speculative_execution.asc
- IBRS mitigation for Spectre V2 (requires updated CPU microcode) FreeBSD-SA-18:03.speculative_execution.asc
- Added a CPU Microcode update mechanism (cpuctl module, sysutils/devcpu-data port)
- Imported a FreeBSD patch to fix boot issues when running as a hypervisor guest on AMD Family 15h processors (FreeBSD PR #213155)
- Added validation for RRD parameters to ensure passed filenames are valid #8269
- Fixed a potential XSS vector in RRD error output encoding #8269 pfSense-SA-18_01.packages Fixed a potential XSS vector in diag_system_activity.php output encoding #8300 pfSense-SA-18_02.webgui
- Fixed a potential XSS vector in traffic_graphs.widget.php settings #8302 pfSense-SA-18_03.webgui

- Fixed a potential CSRF issue in service control request processing #8296
- Enabled CSRF protection for all dashboard widgets #8301
- Added encoding for firewall schedule range descriptions #8259
- Changed sshd to use delayed compression #8245
- Increased PHP-FPM resources on systems with over 1GB RAM to improve performance #8125
- Imported a netstat fix for ARM platforms to improve performance and reduce CPU usage, especially on the Dashboard #8237
- Fixed a memory leak in the pfSense_getall_interface_addresses() function in the pfSense PHP module #8249
- Hardware support for the XG-7100, including:
  - C3000 NIC support (factory installations only)
  - C3000 SoC support (factory installations only)
  - Marvell 88E6190 switch support (factory installations only)

3.32.2 Traffic Shaping / Limiters

- Fixed hangs due to Limiters and pfsync in HA #4310
- Added the Chelsio *cxl* driver to the list of ALTQ capable interfaces #7607
- Fixed an issue with limiters that had fractional bandwidth values #8091
- Changed status_queues.php to provide 'realtime' statistics #8185

3.32.3 IPsec

- Changed IPsec Phase 1 to allow selecting both IPv4 and IPv6 so the local side can allow inbound connections to either address family #6886
- Changed IPsec Phase 1 to allow configuration of multiple IKE encryption algorithms, key lengths, hashes, and DH groups #8186
- Fixed a problem when IPsec bypasslan was enabled while the LAN interface is disabled or doesn't have an IP address #8239
- Added IPv6 LAN Network to the IPsec LAN bypass list #8321

3.32.4 OpenVPN

- Fixed an error message encountered by a few users when manually killing OpenVPN connections #8266
- Added an OpenVPN tap bridge configuration option to push the bridged interface address to clients as a route-gateway for routes/redirects #8267
- Added an option to the DNS Resolver which allows registering the CN of OpenVPN clients as hostnames #6847
- Added an option to OpenVPN clients and servers to suppress creation of IPv4 or IPv6 gateway addresses for an interface #6848
- Fixed issues with OpenVPN when using a /31 IPv4 Tunnel Network #8261
- Updated the OpenVPN wizard with the current UDP and TCP protocol selections #8298 Added the interface for a VPN to the OpenVPN client and server list screens

3.32.5 Notifications

- Changed SMTP notifications handling so they are batched, to avoid sending multiple e-mail messages in a short amount of time #4031
- Added a notification when the firewall boot sequence is complete #7643

3.32.6 Dashboard

- Fixed issues with the IPsec dashboard widget causes GUI failure #6318
- Changed the Dynamic DNS Widget so it shows the description of custom entries to identify them #7843
- Fixed a reference to deprecated updateGatewayDisplays() function in the Gateways dashboard widget #8303
- Added a setting to the temperature widget to display readings in Fahrenheit 8205
- Changed the picture widget so the picture is stored on the firewall filesystem and not in config.xml to reduce the size of backup data #8371
  - On upgrade, pictures will be moved out of config.xml, so backup this file separately if it is important

3.32.7 DHCP

- Added an option to the DHCP Server Dynamic DNS configuration to set the server key algorithm #6621
- Added DDNS Client Updates option to DHCPv4 #7131
- Fixed handling of the DHCPv6 DDNS reverse zone key #6319
- Fixed DHCPv4 static mappings so that multiple MAC for same DHCP address or hostname are allowed`#8220 <https://redmine.pfsense.org/issues/8220>`__
- Fixed a potential issue in detecting primary/secondary node in a failover configuration
- Improved DHCP relay destination interface discovery
- Fixed DHCPv6 lease display for entries that were not parsed properly from the lease database #7413

3.32.8 Dynamic DNS

- Added an option for RFC 2136 Dynamic DNS server key algorithm #8244
- Added an option for RFC 2136 source address used to send updates #8278
- Fixed issues with Dynamic DNS updates using a gateway group when the primary route is down #8333
- Added GoDaddy Dynamic DNS provider

3.32.9 Interfaces / VIPs

- Fixed issues on assign_interfaces.php with large numbers of interfaces #6400
- Fixed handling of CARP VIPs on disabled interfaces at boot time #6677
- Fixed issues with radvd being enabled on a disconnected interface #6974
- Fixed issues with rtsold on VLAN interfaces #7412
- Fixed issues with dhcp6c lock files after unclean shutdown when using "Do not wait for an RA" on IPv6 WAN interface #8106
- Added a feature to allow pppoe on a CARP VIP so it will only be active on whichever node is master #8184
- Fixed an error when editing PPP interfaces on a system with no VIPs #8322
- Added VLAN priority tagging for DHCPv6 client requests #8200
- Added support for configuring the DUID type for an IPv6 interfaces #8191
- Allow custom INIT string for PPP modem SIM Pin and APN settings
- Added an indicator for disabled interfaces on status_interfaces.php
- Fixed an issue with the PPP linkup and linkdown scripts and cellular modems
- Fixed an issue where the combination of CARP with bridging could lead to a deadlock #8056

3.32.10 Packages

- Fixed reinstall process for missing packages #8183

3.32.11 Captive Portal

- Fixed Pass-through MAC automatic additions so it does not add duplicate entries #8226
- Fixed a missing global definition in Captive Portal pass-through MAC removal #8238
- Fixed Captive Portal voucher sync errors when vouchers are expired or disconnected while the secondary node is master #8317
- Fixed Captive Portal voucher synchronization between HA nodes #7972

3.32.12 Certificates

- Fixed automatic SAN handling when the CN of a certificate contains a space #8252
- Fixed input validation for Certificate SAN values to disallow IP addresses for FQDN/Hostname entries #8275

3.32.13 Gateways/Routing

- Fixed handling of the Router Lifetime value on services_router_advertisements.php so it allows a value of *0* #7502
- Added ospf6d to the routing log
- Allow recursive aliases to be used with static routes

3.32.14 Rules/NAT

- Fixed various pf "busy" errors when the ruleset is reloaded
- Fixed issues with editing firewall rules in non-English languages that contain single quotes in translated strings #8219
- Added an option to disable drag-and-drop of firewall and NAT rules
- Added a check to prevent 1:1 NAT rules with missing information from being added to the ruleset
- Added firewall rule tracking ID to rule list (in counter tooltip) and firewall rule edit page #8348
- Fixed cases where automatic or scripted rules were not getting tracking IDs #8353
- Added a check to prevent automatic outbound firewall rules with missing information from being added to the ruleset #8360

3.32.15 Users/Authentication

- Fixed issues with XMLRPC user account synchronization causing GUI inaccessibility on secondary HA nodes #7469
- Fixed an issue where a user with no privileges could not logout #8297
- Increased maximum username length from 16 to 32 characters to catch up to the current allowed length in FreeBSD
- Fixed required field markings on LDAP authentication server configuration fields #8337
- Fixed display of the LDAP host when testing the GUI authentication source #8338

3.32.16 Misc

- Fixed NTP Status server time for zones with minute offsets (fractions of an hour) #8129
- Added support for custom shutdown scripts in /usr/local/etc/rc.d #8182
- Fixed a references to an undefined function while restoring a config.xml file from an older version #8231
- Added support to diag_packet_capture.php to capture traffic on the loopback interface #8257
- Fixed an issue with the RAM disk warning pop-up appearing when no changes were made #8268
- Fixed an issue with the address family selection for remote syslog servers using IPv6 #8323
- Silenced warnings from sysctl that otherwise went to stderr
- Added a disk size check to ZFS to prevent it from being used on disk which are too small to contain the OS and swap space #7308
- Added a check to prevent pfSense-upgrade from running as a non-root user #7762

- Added an option to disable the IGMP Proxy service #8356
- Fixed an issue with package handling when restoring a configuration that contains a branch configuration that is not valid for the target system version #8208

3.33 2.4.2-p1 New Features and Changes

New features and changes for this release of pfSense® software:

3.33.1 Security / Errata

- Updated OpenSSL to address CVE-2017-3737 and CVE-2017-3738 FreeBSD-SA-17:12.openssl
- Fixed a potential authenticated command execution issue in certificate data handling #8153 pfSense-SA-17_10.packages.asc
- Fixed a potential XSS issue in status_filter_reload.php #8143 pfSense-SA-17_11.packages.asc

3.33.2 Misc

- Fixed an issue with the subnet mask not being preserved properly when editing existing 1:1 NAT entries #8112
- Fixed an indexing issue when deleting Host Override entries from the DNS Forwarder #8159
- Fixed logging for L2TP and PPPoE server login/logout events #8164
- Removed ix from the ALTQ interface list since ALTQ support for the ix driver is not currently viable #7378
- Fixed a premature session timeout issue on pages which update exclusively using AJAX, such as status_graph.php #8116
- Fixed ping_hosts.sh so it does not unnecessarily run a CARP check when there are no IPsec hosts to ping #8172
- Fixed a missing global variable declaration in interface IP address detection
- Fixed issues with local authentication when using translated languages

3.34 2.4.2 New Features and Changes

New features and changes for this release of pfSense® software:

3.34.1 Security / Errata

- Updated to OpenSSL 1.0.2m to address CVE-2017-3736 and CVE-2017-3735
- FreeBSD-SA-17:10.kldstat
- FreeBSD-SA-17:08.ptrace
- Fixed a potential XSS vector in status_monitoring.php #8037 pfSense-SA-17_07.packages.asc
- Fixed a potential XSS vector in diag_dns.php #7999 pfSense-SA-17_08.webgui.asc
- Fixed a potential XSS vector on index.php via widget sequence parameters #8000 pfSense-SA-17_09.webgui.asc

- Fixed a potential XSS in the widgetkey parameter of multi-instance dashboard widgets #7998 pfSense-SA-17_09.webgui.asc
- Fixed a potential clickjacking issue in the CSRF error page

3.34.2 Interfaces

- Fixed PPP interfaces with a VLAN parent when using the new VLAN names #7981
- Fixed issues with QinQ interfaces failing to show as active #7942
- Fixed a panic/crash when disabling a LAGG interface #7940
- Fixed issues with LAGG interfaces losing their MAC address #7928
- Fixed a crash in radvd on SG-3100 (ARM) #8022
- Fixed an issue with UDP packet drops on SG-1000 #7426
- Added an interface to manage the built-in switch on the SG-3100 Trimmed more characters off the interface description to avoid console menu output line wrapping on a VGA console
- Fixed handling of the VIP uniqueid parameter when changing VIP types
- Fixed PPP link parameter field display when a VLAN parent interface was selected #8098

3.34.3 Operating System

- Fixed issues resulting from having a manually configured filesystem layout with a separate /usr slice #8065
- Fixed issues updating ZFS systems created ZFS using an MBR partition scheme (empty /boot due to bootpool not being imported) #8063
- Fixed issues with BGP sessions utilizing MD5 TCP signatures in routing daemon packages #7969
- Updated dpinger to 3.0
- Enhanced the update repository selection choices and methods
- Updated the system tunables that tell the OS not harvest data from interrupts, point-to-point interfaces and Ethernet devices to reflect the new name/format for FreeBSD 11
- Changed ruleset processing so that it retries if another process is in the middle of an update, rather than presenting an error to the user
- Fixed some UEFI boot issues on various platforms

3.34.4 Certificates

- Fixed invalid entries in /etc/ssl/openssl.cnf (only affected non-standard usage of openssl in the cli/shell) #8059
- Fixed LDAP authentication when the server uses a globally trusted root CA (new CA selection for “Global Root CA List”) #8044
- Fixed issues creating a certificate with a wildcard CN/SAN #7994
- Added validation to the Certificate Manager to prevent importing a non-certificate authority certificate into the CA tab #7885

3.34.5 IPsec

- Fixed a problem using IPsec CA certificates when the subject contains multiple RDNs of the same type #7929
- Fixed an issue with enabling IPsec mobile client support in translated languages #8043
- Fixed issues with IPsec status display/output, including multiple entries (one disconnected, one connected) #8003
- Fixed display of multiple connected mobile IPsec clients #7856
- Fixed display of child SA entries #7856

3.34.6 OpenVPN

- Added an option for OpenVPN servers to utilize "redirect-gateway ipv6" to act as the default gateway for connecting VPN clients with IPv6, similar to "redirect-gateway def1" for IPv4. #8082
- Fixed the OpenVPN Client Certificate Revocation List option #8088

3.34.7 Traffic Shaping

- Fixed an error when configuring a limiter over 2Gb/s (new max is 4Gb/s) #7979
- Fixed issues with bridge network interfaces not supporting ALTQ #7936
- Fixed issues with vtnet network interfaces not supporting ALTQ #7594
- Fixed an issue with Status > Queues failing to display statistics for VLAN interfaces #8007
- Fixed an issue with traffic shaping queues not allowing the total of all child queues to be 100% #7786
- Fixed an issue with limiters given invalid fractional/non-integer values from limiter entries or passed to Captive Portal from RADIUS #8097

3.34.8 Rules/NAT

- Fixed selection of IPv6 gateways when creating a new firewall rule #8053
- Fixed errors on the Port Forward configuration page resulting from stale/non-pfSense cookie/query data #8039
- Fixed setting VLAN Priority via firewall rules #7973

3.34.9 XMLRPC

- Fixed a problem with XMLRPC synchronization when the synchronization user has a password containing spaces #8032
- Fixed XMLRPC Issues with Captive Portal vouchers #8079

3.34.10 WebGUI

- Added an option to disable HSTS for the GUI web server #6650
- Changed the GUI web service to block direct download of .inc files #8005
- Fixed sorting of Services on the dashboard widget and Services Status page #8069
- Fixed an input issue where static IPv6 entries allowed invalid input for address fields #8024
- Fixed a JavaScript syntax error in traffic graphs when invalid data is encountered (e.g. user was logged out or session cleared) #7990
- Fixed sampling errors in Traffic Graphs #7966
- Fixed a JavaScript error on Status > Monitoring #7961
- Fixed a display issue with empty tables on Internet Explorer 11 #7978
- Changed configuration processing to use an exception rather than die() when it detects a corrupted configuration
- Added filtering to the pfTop page
- Added a means for packages to display a modal to the user (e.g. reboot required before package can be used)

3.34.11 Dashboard

- Fixed display of available updates on the Installed Packages Dashboard widget #8035
- Fixed a font issue in the Support Dashboard widget #7980
- Fixed formatting of disk slices/partitions in the System Information Dashboard widget
- Fixed an issue with the Pictures widget when there is no valid picture saved #7896

3.34.12 Packages

- Fixed display of packages which have been removed from the repository in the Package Manager #7946
- Fixed an issue displaying locally installed packages when the remote package repository is unavailable #7917

3.34.13 Misc

- Fixed interface binding in ntpd so it does not erroneously listen on all interfaces #8046
- Fixed a problem where restarting the syslogd service would make sshlockout_pf process orphans #7984
- Added support for the ClouDNS dynamic DNS provider #7823
- Fixed an issue in the User and Group Manager pages when operating on entries immediately after deleting an entry #7733
- Changed the setup wizard so it skips interface configuration when run on an AWS EC2 Instance #6459
- Fixed an IGMP Proxy issue with All-multicast mode on SG-1000 #7710

3.35 2.4.1 New Features and Changes

3.35.1 Security / Errata

- Fixes for the set of WPA2 Key Reinstallation Attack issues commonly known as KRACK #7951
- Changed upgrade handling to use the pkg-static binary to prevent errors when moving to new major FreeBSD version Fixed a VT console race condition panic at boot on VMware platforms (especially ESXi 6.5.0U1) #7925
- Fixed a bsnmpd problem that causes it to use all available CPU and RAM with the hostres module in cases where disk drives are present without media inserted #6882
- Fixed an upgrade problem due to FreeBSD 11 removing legacy ada aliases, which caused some older installs to fail when mounting root post-upgrade #7937
- Changed the boot-time fsck process the ensure the disk is mounted read-only before running fsck in the preen mode.

3.35.2 Known Issues

- The VLAN changes mentioned in the **Interfaces** section may prevent PPP sessions from working on VLANs in some cases, see #7981

3.35.3 Interfaces

- Changed the VLAN interface names to use the 'dotted' format of FreeBSD, which is shorter and helps to keep the interface name smaller than the limit (16) This fixes the 4 digit VLAN issues when the NIC name is 6 bytes long.
- Improved the 'Assign Interfaces' console process to automatically stop when there are no more interfaces to assign
- Improved the 'Set interface IP address' console process to accept 'IP/mask' notation
- Fixed wireless client interfaces so they do not reconfigure wireless on a link up event, or else they can get stuck in a loop #7960
- Fixed setting VLAN Priority in VLAN interface configuration #7748

3.35.4 Dashboard

- Fixed a problem with the Picture Dashboard widget when it does not have a picture defined #7896
- Fixed time display for UTC in the NTP Dashboard Widget #7714
- Fixed an IPsec widget error when it would get back null data after a session ended #6318
- Improved error checking to prevent dashboard widget parsing errors

3.35.5 DNS

- Added an option for the DNS Resolver (Unbound) to serve expired records from the cache after their TTL expires which can improve speed in some cases #7814
- Fixed the DNS Resolver (Unbound) to allow snoop from localhost by default, otherwise "dig +trace" or "drill -T" queries from the firewall itself fail #7884

3.35.6 XMLRPC

- Fixed XMLRPC Sync to prevent a lock that would never be unlocked
- Fixed XMLRPC sync to ensure a proper empty array is returned instead of NULL, so that the last item of a section can be removed without error #7953

3.35.7 Misc

- Fixed Captive Portal voucher test and expire pages #7939
- Added UEFI 32 and UEFI 64 filenames defined inside a pool to dhcpd.conf #7949
- Fixed operation of the "Reset All States on WAN IP Change" GUI setting #7921
- Changed OpenVPN to retry client auth when it fails by default (auth-retry nointeract) #7506
- Changed the Cryptographic Accelerator module options to allow both the AES-NI and Crypto modules to be loaded at the same time #7810
- Added URL fingerprinting to the login page CSS
- Added the device serial/id to the console and SSH menu banner #7968
- Fixed "Unknown Step Values" in certain RRD graph cases #6860

3.36 2.4 New Features and Changes

3.36.1 Operating System / Architecture changes

Warning: **32-bit support has been deprecated and removed** – There are no images available for 32-bit (x86/i386) Intel architecture systems

Warning: **NanoBSD has been deprecated and removed** – There are no images available for NanoBSD, use a full install instead

- Upgrade of base OS to FreeBSD **11.1-RELEASE-p1**
- Added support for Netgate ARM-based systems such as the SG-1000
- Started using the FreeBSD installer instead of the old style installer (installation procedures have all changed)
  - The installer now supports UEFI #4044
  - If the new installer image will not boot on a specific piece of hardware, see *Troubleshooting Boot Issues*

  - The installer now supports ZFS
  - Added support to the new installer to copy an existing config.xml off an MS DOS formatted USB drive (formerly known as "PFI") #7689
  - Added support to the new installer to optionally recover config.xml off an existing installation drive (works with UFS and ZFS) #7708
- Fixed issues with major version base upgrades via pkg
- Changed cryptodev to load as a kernel module #5976

3.36.2 Security / Errata

- Converted various parts of the GUI to use POST instead of GET when performing actions that change the firewall state (e.g. delete or enable/disable an item) to avoid potential issues with cross-site request forgery and unintentional repeating of actions #4083
- FreeBSD 11.1 includes MAP_GUARD protection to protect against attacks such as Stack Clash
- pfSense-SA-17_07.packages
- A number of base system packages have been updated to address security issues, including dnsmasq, perl, cURL, and others.

3.36.3 Firmware Branch Behavior / Upgrading From Snapshots

- To control how a firewall obtains updates, visit **System > Update**, **Update Settings** tab

3.36.4 Known Issues

- Some systems may not be able to boot 2.4 installation images, for example, due to UEFI compatibility changes. These are primarily BIOS issues and not issues with the installer images. Upgrading from 2.3.x should still work on affected hardware.
- Users with ESXi or VMware Workstation may experience a boot-time crash during hardware detection, due to a race condition in the FreeBSD VT console code. This crash is infrequent and does not affect most users or most boot attempts, but since it is a race condition it can manifest randomly. To avoid the crash, configure the VM to use the `syscons` console rather than `vt` by editing `/boot/loader.conf.local` and adding this line:

```
kern.vty=sc
```

3.36.5 Cleanup

- Misc code cleanup, removal of patches that were no longer necessary or were inefficient
- Replaced multiple local copies of PHP PEAR libraries with updated copies using their official sources #3734
  - Notably, local static copies were replaced by their FreeBSD ports counterparts: pear, pear-XML_RPC2, pear-Net_IPv6, pear-Crypt_CHAP, pear-Mail, pear-Net_Growl
  - Code that relied on the old files was updated to use the current or replaced versions
- Removed all references to GLXSB (it was 32-bit only) #6755
- Removed all code in the builder and pfSense for handling the NanoBSD platform

- Removed all calls to conf_mount_rw / conf_mount_ro, since they were only required for NanoBSD
- Improved help text in various parts of the GUI

3.36.6 Wireless

- FreeBSD 11 contains an updated 802.11 stack with numerous improvements

> **Warning:** Wireless interfaces **must** be created on the **Wireless** tab under **Interfaces > Assignments** before they can be assigned! #6770

3.36.7 Firewall / Rules / NAT / Aliases

- Fixed issues with synproxy rules on a WAN/LAN style bridge #6769
- Fixed issues with limiters on rules that utilize NAT #4326
- Fixed issues with limiters used in conjunction with a transparent proxy or other local redirect rule #7050
- Fixed expansion of "Other" type VIP subnet entries in NAT destination drop-down selections #6094
- Fixed NAT rules so that when a port forward is disabled, its associated firewall rule is also disabled #6472
- Fixed handling of "URL Table (IPs)" and "URL (IPs)" when the file is hosted a server using HTTPS with a self-signed certificate #4766
- Show firewall rule descriptions in a column when viewing the log on new installs, upgrades retain their existing setting #7323
- Fixed firewall states showing a negative value for total bytes processed #7075
- Fixed handling of Port Forwards so they do not make up new destination information when a configured against a DHCP interface that does not currently have an address
- Fixed VLAN Priority pf syntax #7744
- Fixed a problem where pf scrub did not properly re-fragment unusual but valid IPv6 fragments, resulting in overlapping fragments #7485
- Fixed confirmation prompt handling when deleting a firewall state from diag_dump_states.php #7827
- Changed display of 1:1 NAT rules to match other firewall pages #7728

3.36.8 Traffic Shaping

- Added extra warnings to traffic shaping pages when the firewall has no interfaces capable of using ALTQ shaping #7032
- Fixed handling removal of shaping rules when deleting an interface #7231
- Added upgrade code to work around broken shaper rules from older wizard code #7434
- Fixed the Traffic Shaper so it shows interface names for disabled interfaces, rather than an 'empty' placeholder.
- Fixed handling of the priority field for different ALTQ shaper types

3.36.9 OpenVPN

- Upgraded OpenVPN to 2.4.x. #7054
  - This is a significant upgrade which includes support for a wide variety of new features, including AEAD ciphers such as AES-GCM.
  - AES-GCM can be accelerated by AES-NI, and is supported in SSL/TLS modes (not shared key) #7068
  - Added support for TLS Encryption as an optional TLS Key usage type. This encrypts the control channel, providing privacy and protocol obfuscation #7071
  - Added ECDH options to OpenVPN server and client options ("ECDH Only" choice for DH, ECDH Curve selection) #7063
  - Restructured the compression options to include LZ4 support and the new "compress" directive which replaces "comp-lzo" which has been deprecated. The old options remain for now, but are labeled "Legacy" #7064
  - Changed protocol selection for OpenVPN clients and servers because OpenVPN 2.4 treats "udp" and "tcp" as dual stack now #7062
    - Added "multihome" option in relevant protocol cases so OpenVPN will reply back using the address used to receive a packet #7062
  - Changed the DNS Server fields in the OpenVPN server options so they can define either IPv4 or IPv6 DNS servers to push to clients`#7061 <https://redmine.pfsense.org/issues/7061>`__
  - Added IPv6 support to status_openvpn.php and the OpenVPN widget #2766
  - Removed uses of the deprecated "tun-ipv6" OpenVPN directive, OpenVPN now always assumes IPv6 is enabled #7054
  - Replaced uses of the deprecated "client-cert-not-required" directive with its functional replacement "verify-client-cert none" #7073
  - Added support for Negotiable Crypto Parameters (NCP) to control automatic cipher selection between clients and servers #7072

Note: OpenVPN 2.4 handles CRL verification differently than previous versions, passing through validation to the library rather than handling it internally. This can cause some certificates to fail validation that may have passed previously. In particular, if a certificate is removed from a CRL, it may still fail validation until all copies of the CRL have been rewritten.

- Improved the help text on OpenVPN Client-Specific Overrides #7053
- Fixed issues with OpenVPN clients on dynamic or tunneled IPv6 interfaces (e.g. GIF) #6663
- Added locking to prevent issues with OpenVPN instance startup #6132
- Check OpenVPN server/client option visibility changes per mode #7331 #7451
- Added an OpenVPN GUI option for "fast-io" to clients and servers #7507
- Added an OpenVPN GUI Option for "sndbuf" and "rcvbuf", using the same value for both #7507
- Removed references to the defunct OpenVPN client manager port #7568
- Removed references to unused "Address Pool" setting in OpenVPN #7567
- Fixed OpenVPN server port validation to disallow "0", while still allowing it for a client port, which is the same meaning as blank/empty #7565

- Fixed OpenVPN help text for route_no_exec #7575
- Fixed description of the address assignment behavior for Tunnel Network fields in OpenVPN clients and servers #7573
- Remove the GUI option for "resolv-retry infinite" from OpenVPN, it is always enabled #7572
- Fixed the OpenVPN wizard so it better handles a user choosing a different type of authentication server than a previous run of the wizard #7569
- Fixed OpenVPN Auth Digest Algorithm selection so it does not use duplicate/alias names in the list, and added upgrade code to fix existing entries on upgrade so they use the actual digest name and not an alias #7685
- Fixed show/hide behavior of fields on vpn_openvpn_client.php in chrome #7451
- Changed OpenVPN wizard certificate input validation and encoding so it matches the standards of the current certificate manager #7854
- Fixed the OpenVPN wizard so it creates an OpenVPN server instance using current proper defaults #7864

3.36.10 IPsec

- Upgraded strongSwan to version 5.6.0
- Changed the default strongSwan logging levels such that IKE SA, IKE Child SA, and Configuration Backend all default to "Diag" #7007
- Added an option to set the Rekey Margin for IPsec tunnels in the Phase 1 settings
- Added RADIUS accounting support for mobile IPsec when accounting is enabled on the Authentication Server entry
- Added checks to prevent simultaneous/repeated calling of vpn_ipsec_configure() by /etc/rc.newipsecdns
- Added DH Groups 22, 23, 24 to IPsec Phase 2 selection for compatibility, but they should not normally be used for security reasons #6967

3.36.11 Certificate Management

- Added a check to ensure that the public key of the Certificate matches its private key when importing Certificate Authority and Certificate entries to prevent mismatching keys from being imported #6953
- Fixed error handling when creating a Certificate from the User Management section, failed actions will no longer fail silently #6953
- Fixed handling of Certificates generated from an imported CA when no starting serial number was set #6952
- Fixed handling of Certificate Authority deletion so that it does not remove associated certificates #6947
- Added "in-use" testing for Certificate Authority entries and disabled the delete action for CAs which are actively in use #6947
- Fixed choosing an existing user certificate when adding a certificate to an existing user #7297
- Added the ability for the certificate manager to sign a CSR using an internal CA #7383
- Added the ability to set the certificate type and SAN attributes in a Certificate Signing Request #7527
- Restructured how certificate types and SANs are handled in the cert manager when making a Cert/CSR/Signing, so each section can properly use the controls #7527 #7677

  It is now possible to add SANs and EKUs to certificates when signing using the certificate manager

Note: Attributes such as SANs and KU/EKU cannot be copied from a CSR when signing due to a deficiency in OpenSSL's x509 functions (they do not support "copy_extensions" at this time); These attributes must be specified manually when signing

- Fixed "server" certificate detection to key off of the EKU For "TLS Web Server Authentication" since nsCertType has been deprecated
- Added SAN, KU, and EKU information in an info block for each entry in the certificate list #7505
- Added the ability to use a wider range of characters in certificate fields as laid out by RFC 4514 #7540
- Added a useful error message when there is no private CA with which to create a new user certificate from within the user manager #7585
- Fixed the User Manager so it adds the username as the first SAN when making a user certificate at the same time a user is created #7666
- Added another possible Certificate Signing Request Armor string when validating on import #7383

3.36.12 Dynamic DNS

- Fixed response parsing for DNSimple Dynamic DNS #6874
- Fixed handling of password in Dynamic DNS entries to allow special characters #6688
- Changed CloudFlare and GratisDNS to use separate hostname and domain entry to handle TLDs with multiple components #6778
- Fixed the Save and Force Update button for RFC2136 Dynamic DNS #7291
- Fixed RFC2136 Dynamic DNS updates at boot time #7295
- Added the 'local' directive to RFC2136 Dynamic DNS so updates are sourced correctly #7446
- Fixed options text and display for IPv4 DNS and Verify SSL on Dynamic DNS clients #7588
- Fixed issues with Dynamic DNS entries utilizing gateway groups for their interface #7719
- Added DreamHost Dynamic DNS support #7321

3.36.13 DHCP Server / Relay

- Fixed handling of DHCPv6 lease status when there are no leases #6717
- Fixed issues with DHCP Relay not working #6658
- Added input validation to prevent the DHCP server from being configured on interfaces that do not have enough addresses for clients (/31, /32) #6930
- Fixed issues with the DHCP Relay options display getting out of sync with checkbox settings #7155
- Fixed static DHCP lease edits updating BIND zones #3710
- Fixed checks for DHCP Relay when editing additional DHCP pools
- Fixed handling of forced Dynamic DNS hostnames for DHCPv6 static mappings #7324

3.36.14 ARP / NDP

- Fixed static ARP handling when creating or editing DHCP static mappings #6821
- Added error checking for static ARP entries to ensure both an IP address and MAC address are entered, and to ensure that both exist before an entry is applied #6969
- Improved the detail displayed on the ARP table view #6822
- Added an expiration field to the NDP list

3.36.15 Captive Portal

- Adapted Captive Portal to work without multi-instance ipfw patches #6606
- Fixed Captive Portal instances to select "No Authentication" for a zone by default, since it is the default behavior #7591
- Fixed links to the Captive Portal MAC management page so they include the zone name #7591

3.36.16 XMLRPC

- Switched to pear-XML_RPC2 and removed the outdated static client files
- Fixed handling of XMLRPC sync using a username other than "admin" #809

3.36.17 Routing/Gateways

- Removed "route change" patches and updated code that relied on the deprecated behavior #6828
- Fixed handling of default routes when a default gateway is removed or disabled #6659
- Fixed discovery of IPv6 gateway for assigned OpenVPN interfaces #6016
- Fixed issues with a missing default gateway/route on certain PPPoE links after reconnect or IP address change #6495
- Fixed some 'route: writing to routing socket: Invalid argument' warnings during boot time
- Added a log message for gateway events that shows that an alarm was raised/cleared
- Added a check to not run dpinger when an IPv6 address has the tentative flag even after the timeout
- Added a delay to allow dpinger time to properly initialize before using results

3.36.18 Interfaces / Virtual IP Addresses

- Removed Device Polling as it was no longer useful #7021
- Improved stability of the igb(4) driver #7149 #7166
- Fixed handling of rc.newwanipv6 when run from dhcp6c so it only runs when required and not for any change #7145
- Fixed handling of SIGTERM and SIGKILL in dhcp6c #7185
- Fixed dhcp6c not starting until an RA is received #5993
- Fixed a PPP service name error with certain providers, such as T-Mobile #6890

- Fixed 3G service status so it does not report misleading information #4287
- Added support for the IPv6 AUTO_LINKLOCAL flag on bridge interfaces
- Disabled DAD on stf interfaces to fix problems with dpinger
- Added an option to use static IPv6 over an IPv4 PPP parent (e.g. PPPoE) #7598
- Removed unused WINS code for L2TP #7559
- Improved L2TP Server DNS input validation #7560
- Added a test to disable internal L2TP users when activating RADIUS, to follow the behavior stated in the GUI #7561
- Fixed L2TP section log shortcut #7564
- Fixed upgrade handling of wireless interfaces #7809

3.36.19 NTP

- Added support for the ntpd "pool" directive to make better use of servers in NTP pools #5985
- Fixed time display on the NTP widget to show server time #7245
- Added support for NTP to process PGRMF NMEA sentences (Garmin-specific) #7193
- Added an absolute offset statistic to NTP monitoring graph display #7548

3.36.20 User Management / Authentication

- Fixed delays during bootup when LDAP is enabled for user authentication #6367
- Added privileges to control which users can view and/or clear notices #7051
- Added an authentication cache mechanism for GUI authentication from a remote server (e.g. LDAP, RADIUS) so the authentication is checked periodically (default: 30s) instead of on each page load #7097
- Added protocol selection (PAP, MD5-CHAP, MS-CHAPv1 and MS-CHAPv2) to RADIUS authentication server options #7111
- Added the username to the page to display when adding user privileges #7586
- Standardized privilege page and sorting between users and groups #7587
- Added a log message if a user tries to save the configuration but has the 'deny config write' permission
- Added "auth_check" type of simple test that a page can use to verify a user is logged in and has access, using less cpu, which is better for AJAX data polling
- Fixed certificate chain verification issues with LDAP authentication using intermediate CAs #7830
- Fixed PHP errors when STARTTLS fails for LDAP authentication

3.36.21 Packages

- Fixed issues with snort, squid/clamav, and squidGuard when /var is in a RAM disk #6878
- Fixed handling of custom_php_deinstall_command during post-deinstall of a package #7401
- Changed package related calls to get_pkg_info() to use the new pkg metadata mechanism

3.36.22 Console / Menu

- Added options to the console reboot menu selection to reboot into single user mode or run filesystem checks #6639

3.36.23 OS Upgrade

- Fixed issues when upgrading to 2.4 with a stale package .inc that caused a PHP error #6920
- Changed the upgrade script to use reroot instead of reboot for updates that do not include a new Kernel #6045

3.36.24 SNMP

- Added a workaround to prevent the hostres module from being used with bsnmpd on VMware Virtual Machines that have a cd0 device, which caused 100% CPU usage #6882

3.36.25 Services

- Converted all mpd-based features (e.g. PPPoE and L2TP server) to MPD5 if they used an older version #4706
- Removed unused and non-functional SMART service handling and e-mail configuration #6393
- Fixed IGMP Proxy failing to recognize an upstream interface #6099

3.36.26 WebGUI

- Added support for multiple languages, currently that list includes:
  - US English (Default), Bosnian, Chinese (Simplified, China), Chinese (Taiwan), Dutch, German, Norwegian Bokmal, Polish, Portuguese (Brazil), Russian, Spanish, Spanish (Argentina)
- Changed the design of the login page for the WebGUI to a more modern style, with several color choices available
- Added URL fingerprinting to JavaScript and CSS file references to improve client-side behavior when files change between versions #7251
- Updated Logo to the new logo and made it a vectorized SVG image for better scaling
- Updated favicon to the new logo and added multiple sizes for different platforms
- Completed work to mark required fields on GUI pages #7160
- Fixed long hostnames overlapping the “time” title in the monitoring graphs #6138
- Fixed CIDR/Prefix selector handling for IPv4/IPv6 #7625
- Removed the Gold menu
- Fixed handling of info block content inside tables #7504

- Improved handling of PHP errors for user-entered PHP code on diag_command.php
- Fixed alignment of the IPv6 over IPv4 input fields #7128
- Optimized retrieval of Traffic Graph data to reduce spikes in the graphs and load on the firewall
- Fixed a problem with the traffic graphs not respecting the theme colors #6746
- Revised setup wizard wording and links

3.36.27 Dashboard

- Rewrote Dashboard AJAX updating in a centralized and optimized way to reduce load, improve accuracy, and increase speed
- Added a new Customer Support dashboard widget, enabled by default and on upgrade
- Changed the way AJAX updates are handled on the Dashboard widgets to improve efficiency and fix issues with some widgets refreshing in a timely manner
- Added filters to more dashboard widgets #7122
- Added customization for dashboard widget names
- Fixed Interface Statistics dashboard widget issues with interfaces in a "down" state
- Fixed formatting issues with the Interface Statistics dashboard widget #7501
- Added the ability to place multiple copies of widgets on the dashboard, optional for each widget
- Added a line to display detected CPU cryptographic hardware, such as AES-NI, in the System Information dashboard widget even if the module isn't loaded #7529
- Fixed CPU package/core count displayed on the System Information dashboard widget
- Changed how pkg metadata is handled to reduce the load on the Dashboard and reduce unnecessary calls to the pkg server for the System Information dashboard widget update check, and for the Installed Packages dashboard widget
- Changed CPU usage calculation in the System Information dashboard widget to avoid sleep() in an AJAX call
- Fixed the IPsec widget tunnel status to handle newer strongSwan childid format #7499
- Fixed error when saving Wake on LAN dashboard widget without any WoL entries
- Fixed a problem where traffic could be counted twice in traffic graphs #7751
- Fixed a problem with the Installed Packages dashboard widget when no packages are installed #7811
- Changed date formats of some fields on the Dashboard to be more consistent #7805
- Added an option to the Interface Statistics dashboard widget to rotate the table (put interfaces in rows instead of columns) to improve the display on firewalls with numerous interfaces #7501

3.36.28 pftop

- Removed the "size" option from pftop as it had no effect, use the "bytes" option instead #7579
- Removed the 'peak' and 'rate' views for pftop since they only work in interactive mode with cached data, not batch mode which is used by the WebGUI #7580
- Fixed path to an old copy of the pftop WebGUI page in obsolete list #7581

3.36.29 DNS

- Changed /etc/hosts such that the FQDN is listed first, except for localhost, so that dnsmasq will properly reverse resolve hostnames #7771
- Fixed a problem where the DNS Search Domain List was not being populated into radvd.conf #7081
- Enabled Python support for Unbound #7549
- Added a control to disable automatically added host entries in Unbound
- Changed the way unbound is started at boot time on systems with DHCP6 WANs

3.36.30 Misc

- Added hardware support and detection for new Netgate models
- Changed the User Agent passed to outbound requests from pfSense to include more accurate host information
- Added the User Agent to the request data when updating the Bogons list
- Fixed growl and SMTP notifications so performing a test saves first, so the new settings are used as expected #7577
- Fixed loading issues with PHP extensions #6628
- Removed symbolic links for configuration files that redirected items from /etc/ to /var/etc/ #5538
- Added the ability to filter Packet Captures by MAC address #6743
- Updated status.php with new info and changed its output organization #7047
- Fixed a problem where a proxy defined for use by the firewall could not use HTTPS when using proxy authentication #6949
- Improved RAM disk backups and file management #7098
- Changed the way RAM disk contents are handled when enabled #5897
- Changed various support functions to better facilitate translation to additional languages
- Fixed interface name display on the Router Advertisement configuration page #7133
- Fixed various issues with handling of unusually formatted, but valid, IPv6 addresses #7147
- Improved error handling when a client is logged when it attempts to poll data via rrd_fetch_json.php #6748
- Fixed various issues when the configuration backup count was set to 0 (disabled) #7273
- Fixed handling of "0" for the number of backups to retain in the configuration history #7273
- Fixed an issue with long configuration change descriptions leading to wrapping issues in certain cases such as AutoConfigBackup #6363

- Fixed an issue with installing packages from a backup when restoring using the External Configuration Locater on the first boot post-install #7914

3.37 2.3.5-p1 New Features and Changes

New features and changes for this release of pfSense® software:

3.37.1 Security / Errata

- Updated OpenSSL to address CVE-2017-3737 and CVE-2017-3738 FreeBSD-SA-17:12.openssl
- Fixed a potential authenticated command execution issue in certificate data handling #8153 pfSense-SA-17_10.packages.asc
- Fixed a potential clickjacking issue in the CSRF error page
- Fixed a potential XSS issue in status_filter_reload.php #8143 pfSense-SA-17_11.packages.asc

3.37.2 Misc

- Fixed an issue in the User and Group Manager pages when operating on entries immediately after deleting an entry #7733
- Fixed sorting of Services on the dashboard widget and Services Status page #8069
- Fixed display of available updates on the Installed Packages Dashboard widget #8035
- Fixed display of packages which have been removed from the repository in the Package Manager #7946
- Fixed the OpenVPN Client Certificate Revocation List option #8088
- Fixed an issue with the Pictures widget when there is no valid picture saved #7896
- Fixed an indexing issue when deleting Host Override entries from the DNS Forwarder #8159
- Fixed a premature session timeout issue on pages which update exclusively using AJAX, such as status_graph.php #8116
- Fixed ping_hosts.sh so it does not unnecessarily run a CARP check when there are no IPsec hosts to ping #8172
- Fixed a missing global variable declaration in interface IP address detection

3.38 2.3.5 New Features and Changes

The pfSense® software version 2.3.x release is a Security and Errata maintenance release. 2.4.x is the primary stable supported branch. If the firewall hardware is capable of running 2.4.x, consider upgrading to that release instead.

Updating to 2.3.5 from 2.3.4 on an amd64 installation that could otherwise use 2.4.x requires configuring the firewall to stay on 2.3.x as follows:

- Navigate to **System > Update**, **Update Settings** tab
- Set **Branch** to **Security / Errata Only**
- Navigate back to the **Update** tab to see the latest 2.3.x update

If the update system offers an upgrade to 2.3.5 but the upgrade will not proceed, ensure the firewall has correct versions of the repository configuration and upgrade script for 2.3.x by running the following commands from the console or shell:

```
pkg install -fy pfSense-repo pfSense-upgrade
```

Firewalls running 32-bit (i386) installations of pfSense software do not need to take any special actions to remain on 2.3.x as they are unable to run later versions.

3.38.1 Operating System / Architecture changes

- Upgrade of base OS to FreeBSD **10.3-RELEASE-p20**
- Fixed issues with major version base upgrades via pkg

3.38.2 Security / Errata

- pfSense-SA-17_07.packages
- Fixes for the set of WPA2 Key Reinstallation Attack issues commonly known as KRACK in wpa_supplicant and hostapd (FreeBSD-SA-17:07.wpa)
- A number of base system packages have been updated to address security issues, including dnsmasq, perl, cURL, and others.

3.38.3 Interfaces

- Added support for the IPv6 AUTO_LINKLOCAL flag on bridge interfaces
- Added an option to use static IPv6 over an IPv4 PPP parent (e.g. PPPoE) #7598
- Added IPv6 Prefix Delegation interface selection
- Improved input validation for GIF interfaces #7789

3.38.4 Dashboard

- Rewrote Dashboard AJAX updating in a centralized and optimized way to reduce load, improve accuracy, and increase speed
- Added a new Customer Support dashboard widget, enabled by default and on upgrade
- Changed the way AJAX updates are handled on the Dashboard widgets to improve efficiency and fix issues with some widgets refreshing in a timely manner
- Changed how pkg metadata is handled to reduce the load on the Dashboard and reduce unnecessary calls to the pkg server for the System Information dashboard widget update check, and for the Installed Packages dashboard widget
- Improved error checking to prevent dashboard widget parsing errors
- Fixed a variable conflict in the NTP Status Dashboard widget #7795
- Fixed a problem with the Picture Dashboard widget when it does not have a picture defined #7896
- Changed IPsec Dashboard Widget tunnel status to handle newer strongSwan childid format #7499
- Fixed time display for UTC in the NTP Dashboard Widget #7714

3.38.5 WebGUI

- Changed the design of the login page for the WebGUI to a more modern style, with several color choices available
- Added URL fingerprinting to JavaScript and CSS file references to improve client-side behavior when files change between versions #7251
- Updated Logo to the new logo and made it a vectorized SVG image for better scaling
- Updated favicon to the new logo and added multiple sizes for different platforms
- Added an option for sorting the Interfaces menu by description
- Added "auth_check" type of simple test that a page can use to verify a user is logged in and has access, using less cpu, which is better for AJAX data polling
- Improved handling of PHP errors for user-entered PHP code on diag_command.php
- Changed Interfaces menu "(Assign)" to "Assignments" and added support for menu divider bars
- Fixed automatic selection of '128' as prefix/mask for IPv6 address fields #7625
- Replaced Math.trunc with Math.floor to make IE properly handle traffic graphs #7804
- Changed nginx configuration so it does not allow direct download of .inc files #8005
- Fixed hostname input handling on diag_dns.php

3.38.6 Gateways

- Added a delay to allow dpinger time to properly initialize before using results
- Added a log message when gateway alarms are raised/cleared to show the parameters that triggered the alarm
- Reset All States on WAN IP Change option #1629

3.38.7 Rules/NAT/Shaper

- Fixed handling of Port Forwards so they do not make up new destination information when a configured against a DHCP interface that does not currently have an address
- Fixed ALTQ Traffic Shaper PRIQ priority number validation

3.38.8 IPsec

- Added an option to set the Rekey Margin for IPsec tunnels in the Phase 1 settings
- Added RADIUS accounting support for mobile IPsec when accounting is enabled on the Authentication Server entry
- Added checks to prevent simultaneous/repeated calling of vpn_ipsec_configure() by /etc/rc.newipsecdns

3.38.9 Misc

- Fixed an issue with installing packages from a backup when restoring using the External Configuration Locater on the first boot post-install #7914
- Fixed handling of forced Dynamic DNS hostnames for DHCPv6 static mappings #7324
- Fixed several issues with cron job updating and removal
- Added the device serial/id to the console and SSH menu banner #7968
- Changed /etc/hosts such that the FQDN is listed first, except for localhost, so that dnsmasq will properly reverse resolve hostnames #7771

3.39 2.3.4-p1 New Features and Changes

The pfSense® software version 2.3.4-p1 errata release is a minor release after *2.3.4* and contains beneficial security and bug fixes.

3.39.1 Security / Errata

- pfSense Security Advisories
  - pfSense-SA-17_05.webgui:
    - Fixed a potential XSS issue in the diag_edit.php file browser #7650
    - Fixed a potential XSS in handling of the 'type' parameter on diag_table.php #7652
    - Fixed validation and a potential XSS in interface names on firewall_nat_edit.php #7651
  - pfSense-SA-17_06.webgui:
    - Added a warning screen to the GUI and prevent access if the client IP address is currently in the lockout table, and also remove the client's connection states #7693

3.39.2 Bug Fixes

Captive Portal

- Fixed Captive Portal RADIUS Authentication to only cache credentials when required to perform reauthentication #7528
- Restored the captive portal feature to view the captive portal page directly from the portal web server as an additional button #7646

Dynamic DNS

- Fixed issues with wildcard CNAME records disappearing from Loopia when doing a DNS update
- Fixed issues with CloudFlare Dynamic DNS
- Fixed Hover Dynamic DNS updates so they Verify the SSL Peer

Logging

- Added syslogd service definition to enable status display and control #4382
- Fixed issues with syslogd stopping when installing or uninstalling some packages #7256

Virtual IP Addresses

- Fixed issues with CARP status display overmatching some VIP numbers #7638
- Fixed pid file handling for choparp (Proxy ARP Daemon)
- Added the ability to sort the Virtual IP address list

DNS

- Fixed diag_dns.php so it will not create an empty alias if name does not resolve
- Fixed diag_dns.php to not show Add Alias if the user does not have privileges to add an alais
- Fixed diag_dns.php to change the update alias button text after adding an alias
- Fixed diag_dns.php to disable the Add Alias button when the host field is changed
- Fixed calls to unbound-control to have the full configuration path specified so they do not fail #7667
- Fixed handling of "redirect" zone entries in the DNS Resolver so they do not produce invalid zones #7690
- Changed the way the DNS Resolver code writes out host entries, so the zones are more well-formed #7690
- Changed the way the DNS Resolver process (unbound) is stopped, to allow it to exit cleanly. #7326

Interfaces

- Fixed DHCPv6 to request a prefix delegation even if no interfaces are set to track6 #4544
- Updated handling of original MAC address retention for interfaces with spoofed MACs
- Fixed an array handling problem when working with gateway entries on the Interface configuration page #7659
- Fixed handling of MSS clamping values for PPPoE/L2TP/PPTP WANs #7675

DHCP

- Fixed an issue where some DHCP Lease information was encoded twice with htmlentities/htmlspecialchars
- Fixed an issue where in some edge cases, a variable was not properly set in a loop, leading to a previous value being reused

Misc

- Removed "/usr/local/share/examples" from obsolete files list, some packages rely on the files being there
- Added a few more items to status.php for support purposes, such as a download button, socket buffer info, and the netgate ID
- Fixed status.php to redact BGP MD5 password/key in output #7642
- Fixed OpenVPN to use is_numeric() to make sure $prefix is not 0
- Changed the "Rule Information" section so it is consistent between firewall and NAT rule pages
- Fixed APU2 detection for devices running coreboot v4.x
- Fixed the tunable description for net.inet.ip.random_id #6087
- Fixed some outdated links for help and support
- Fixed some issues with empty config tags in packages #7624
- Fixed issues with entry IDs after deleting Authentication Server instances #7682

3.40 2.3.4 New Features and Changes

3.40.1 Security / Errata

- Updated base OS to FreeBSD 10.3-RELEASE-p19
- FreeBSD/ports Security Advisories
  - Updated ntpd to 4.2.8p10_2 FreeBSD-SA-17:03.ntp.asc
  - Updated cURL to 7.54.0 (CVE-2017-7407, CVE-2017-7468)
  - Updated libevent to 2.1.8 (CVE-2016-10197, CVE-2016-10196, CVE-2016-10195)
- pfSense® Software Advisories
  - Fixed encoding of displayed values from DHCP leases to prevent a badly formatted DHCP lease hostname from causing a potential XSS #7497 (pfSense-SA-17_04.webgui)
- See the Certificates section below for an important note about GUI certificate errors on Chrome 58 and later

3.40.2 Certificates

- Improved certificate generation to always include the CN as the first Subject Alternative Name (SAN), which fixes issues with Chrome 58+ #7496

  To work around an error with the firewall GUI certificate on Chrome 58+, take one of the following actions:

  - Generate and activate a new GUI certificate automatically, from the console/shell: pfSsh.php playback generateguicert
  - Utilize the *ACME package* to generate a trusted certificate for the GUI via Let's Encrypt
  - Create a own new CA/Server certificate and use that for the GUI
- Fixed linking of a certificates to its CA after submitting the signed version of a CSR #7512

3.40.3 Firewall Rules/NAT/Shaper

- Fixed restarting the Load Balancer (relayd) clearing system tables/aliases #7396
- Fixed ruleset generation to notify when an unresolvable alias is encountered by the parser #7421
- Fixed handling of a rule using an empty port alias #7428
- Fixed the traffic shaping wizard handling of SMB rules in Raise/Lower Other Protocols, it was producing an invalid rule #7434
- Fixed handling of alias renaming after input validation #7473
- Fixed handling of long rule descriptions #7294

3.40.4 Dashboard

- Improved formatting in the gateways widget by reducing the numeric precision of displayed values #6841
- Fixed the NTP widget to show the server time instead of client time #7245
- Added a "None" option to Widgets with filtering options #7318
- Added PPPoE uptime display on the Interfaces dashboard widget #6032
- Added filters to more dashboard widgets #7122
- Added BIOS information to the System Information widget
- Added Netgate Unique ID to the System Information widget

Note: This identifier for support services is only displayed on the Dashboard for information purposes and is **not** transmitted anywhere automatically by default. In the future, customers can use this identifier when requesting support information from Netgate staff or systems.

3.40.5 Configuration

- Fixed issues restoring a configuration containing packages when the firewall does not have Internet connectivity #6594
- Fixed factory reset when Captive Portal has Vouchers enabled #7508
- Cleaned up unused code in diag_backup.php

3.40.6 Interfaces

- Changed interface handling so it retains the original vendor MAC address at power up when spoofing, so it can be restored without a reboot #7011
- Fixed interface assignment of QinQ interfaces #4669
- Fixed errors in PPP service provider selection when a country without providers is selected #7399
- Fixed input handling when editing static IP address fields on interfaces #7493
- Added the ability for DHCP Client WANs to specify a list of IP addresses from which to reject leases #7510

3.40.7 User Manager / Authentication

- Added a warning to system_authservers.php to warn that RADIUS does not work with IPv6 #4154
- Added a status icon to the User Manager to show if a user is enabled or disabled #7517

3.40.8 General GUI

- Added navigation links to breadcrumbs #7099
- Improved service name support and error handling in pfSenseHelpers.js #7445

3.40.9 DHCP

- Changed dhcpleases so it does not start when DHCP Relay is enabled #6750
- Fixed checks for DHCP Relay being enabled/disabled so they are skipped when editing an additional pool

3.40.10 ARP / NDP

- Added the ability to delete NDP entries #7513
- Added expiration field to NDP listing #7514

3.40.11 Misc

- Fixed DNS issues when upgrading NanoBSD #7345
- Fixed the Reset Demotion Status for CARP to function when the demotion value is negative #7424
- Fixed editing of Host Overrides in the DNS Resolver/Forwarder pages #7435
- Fixed service handling (start/stop/restart) for Captive Portal #7444
- Fixed display of the ALTQ "queue" view in pfTop due to recent changes in the pfTop port #7461
- Added support for the Dynamic DNS Client Hover #7511
- Fixed UTF-8 handling in Base64 decoding on diag_edit.php
- Fixed handling of traffic graph data irregularities #7515
- Added visual separation to the legend on the installed packages list #7203
- Changed SMTP and Growl notification test to use the new, unsaved settings #7516

3.41 2.3.3-p1 New Features and Changes

The pfSense® software version 2.3.3-p1 errata release is a minor release after *2.3.3* and contains beneficial security and bug fixes.

3.41.1 Security / Errata

- Updated to FreeBSD 10.3-RELEASE-p17
  - FreeBSD-SA-17:02.openssl (CVE-2016-7055, CVE-2017-3731, CVE-2017-3732)
- Upgraded cURL to 7.53.0 (CVE-2017-2629)

3.41.2 Bug Fixes

- Fixed issues with the upgrade check seeing the version of pfSense-upgrade instead of pfSense in some circumstances. #7343
- Fixed handling of domain-only (@ record) updates for CloudFlare Dynamic DNS #7357
- Fixed a problem with the Dynamic DNS Widget where RFC2136 entries showed an incorrect status #7290
- Fixed Dynamic DNS status widget formatting for medium with browser window #7301
- Fixed a problem with HTML tags showing in certificate description drop-down lists in the Certificate Manager #7296
- Fixed an error loading some older rules with ICMP types #7299
- Fixed display of selected ICMP types for old rules without an ipprotocol option set #7300
- Fixed Log widget filter interface selection with custom interface descriptions #7306
- Fixed the widget Filter All button so it does not affect all widgets #7317
- Fixed the password reset script so it resets the expiration date for the admin account when run, to avoid the user still being locked out #7354

- Fixed the password reset script so it properly handles the case when the admin account has been removed from config.xml #7354
- Fixed input validation of TCP State Timeout on firewall rules so it is not arbitrarily limited to a maximum of 3600 seconds #7356
- Fixed console settings for XG-1540/XG-1541 to use the correct default console #7358
- Fixed initial setup handling of VLAN interfaces when they were assigned at the console before running the Setup Wizard #7364
- Fixed display of OpenSSL and input errors when working in the Certificate Manager #7370
- Fixed Captive Portal "disconnect all" button
- Fixed pkg handling timeouts #6594
- Updated blog URL in the RSS widget
- Removed whirlpool from the list of CA/certificate digest algorithms since it does not work #7370

3.42 2.3.3 New Features and Changes

3.42.1 Security / Errata

- Updated to FreeBSD 10.3-RELEASE-p16
  - FreeBSD Security Advisories
    - FreeBSD-SA-16:29.bspatch
    - FreeBSD-SA-16:31.libarchive
    - FreeBSD-SA-16:33.openssh
    - FreeBSD-SA-16:35.openssl
    - FreeBSD-SA-16:37.libc
    - FreeBSD-SA-16:38.bhyve
    - FreeBSD-SA-16:39.ntp
    - FreeBSD-SA-17:01.openssh
  - FreeBSD Errata Notices
    - FreeBSD-EN-16:17.vm
    - FreeBSD-EN-16:18.loader
- pfSense® Software Advisories
  - pfSense-SA-17_01.webgui
    - Fixed validation and encoding on Captive Portal status pages #7019
  - pfSense-SA-17_02.webgui
    - Fixed update_config_field() in wizard.php so it does not pass user input through eval() #7230
  - pfSense-SA-17_03.webgui
    - Added encoding for 'from' and 'to' before output on pkg_mgr_install.php #7225
    - Added encoding for the contents of pkg_filter before output #7227

    * Converted easyrule.php to use a confirmation landing page so that the parameters can be submitted via POST #7228

- Updated numerous third-party libraries and supporting programs
- Changed behavior of fsck during bootup to improve filesystem stability #6340
- Added protection to /etc/ttys to prevent corruption or missing lines

3.42.2 Known Issues

- The Captive Portal **Disconnect All Users** button does not fully disconnect all users PR#3565
- RFC 2136 Dynamic DNS Entries will show red on the Dashboard widget even when correctly updated #7290
- Firewall rules without an IP protocol set in the configuration which also have an ICMP type set may not load or display correctly. #7299 #7300

3.42.3 General Info

- Added Packages: tinc, cellular, LCDproc, TFTP Server
- Fixed numerous typos and wording issues
- Added marking for required fields on various pages #7083
- Input validation fixes on various pages
- Cleaned up some unneeded files/pages/functions
- Fixed broken/outdated links

3.42.4 OpenVPN

- Changed OpenVPN RADIUS authentication to send proper NAS-Port-Type, NAS-Port, and NAS-Identifier values #6609
- Added compression option to handle connecting to OpenVPN peers which do not have LZO compiled into their OpenVPN executable #6739
- Added a workaround to block outside DNS on Windows 10 OpenVPN clients to prevent DNS leaks #6719
- Improved OpenVPN server handling when using CARP VIPs in Gateway Groups
- Improved handling of chained/intermediate CAs in OpenVPN #2800
- Changed OpenVPN widget so it updates dynamically #6723
- Adapted the encryption cipher list to the new output format in OpenVPN 2.3.12, also now displays key and block lengths #6849
- Changed OpenVPN server list to display more information
- Improved error message to explicitly state allowable characters for certificate fields in the OpenVPN wizard #6432
- Fixed handling of OpenVPN authentication when the backend server name contains special characters (e.g. '&') #7002
- Fixed saving an OpenVPN instance on a DHCP interface that does not currently have an IP address #7031
- Added an IPv6 Tunnel Network field to OpenVPN Client-Specific Overrides #7053

- Fixed changing between tun and tap mode for OpenVPN Clients
- Changed OpenVPN startup to avoid overwriting its configuration, and to wait for its PID file to be written
- Fixed OpenVPN binding to an IP Alias VIP #7136
- Fixed display of disabled OpenVPN clients #7180
- Fixed handling of "redirect-gateway" in Client-Specific Overrides #6633

3.42.5 IPsec

- Clarified IPsec Key Exchange Version drop-down to specify IKEv1/IKEv2 #6898
- Fixed handling of static routes for IPsec peers on tunnels bound to IP Aliases VIPs with CARP parents
- Fixed MSS clamping for mobile IPsec clients #7005
- Added IPsec to the State Table interface list

3.42.6 Interfaces

- Fixed handling of LAGG MTU when child QinQ interfaces are present #6227
- Improved behavior when using DHCP before RA #5993
- Added the ability to send a DHCP Release from Status > Interfaces, rather than only stopping dhclient
- Fixed issues adding/editing QinQ entries
- Fixed input validation of QinQ entries
- Fixed validation to prevent an interface, interface group, and alias from using the same name #6976
- Updated interface group name validation rules to match limits of the operating system
- Prevented interface group names, interface names, and aliases from starting with `pkg_` to reserve it for packages use (e.g. tinc) #7173
- Added validation to prevent Interface Group Names from containing a dash #7173
- Added validation to prevent Interface Groups from being renamed to an existing name #7183
- Fixed issues with Interface Statistics widget display #7134
- Fixes for interfaces_ppps_edit.php to fix MTU validation, interface friendly names, advanced options expansion
- Changed linkup event handling to ignore events for interfaces that are member of bridges which have no IP address configured
- Fixed input validation for L2TP and PPTP WAN type interfaces #6732
- Added validation to prevent adding duplicate gateways from the Interface configuration page
- Fixed handling of IPv6 checksum options for "Disable hardware checksum offload" #5321
- Fixed handling of the confirmation dialog when deleting a VLAN #6916
- Fixed handling of wireless MAC address spoofing
- Fixed wireless channel changing #6833
- Improved labels and help text for IPv6 tunneling options
- Added the ability for an L2TP or PPTP WAN to use a hostname for the remote gateway #6899

3.42.7 Certificate Management

- Added missing recommended key lengths and digests to certificate manager
- Fixed CRL editing so that certificates already contained the CRL are not displayed

3.42.8 Users / Authentication / Privileges

- Fixed SSH Keyboard-Interactive authentication #6963
- Added STARTTLS to LDAP Authentication Server Configuration
- Improved WebGUI usability when a remote LDAP server is not available
- Fixed issues with local_sync_accounts failing during boot when using an LDAP server on a non-local network or hostname #6857
- Fixed port build options for scponly #7012
- Fixed notifications so that the Mark All as Read button is not shown to users who do not have sufficient privileges to use it #3454
- Added privileges to control display of notices #7051
- Standardized privilege name capitalization
- Fixed issues with low-privilege users accessing Help pages #7139 #7140
- Added a privilege for UPnP & NAT-PMP configuration #7141
- Simplified tcsh prompt and changed the prompt so it respects default terminal colors

3.42.9 Firewall / Rules / NAT / Aliases / States

- Fixed restoring rule type selection after input errors while saving firewall rules
- Fixed a copy/paste error in variable test when validating firewall rule ports.
- Corrected the descriptions and behavior of the Adaptive Start and Adaptive End settings for firewall state handling
- Fixed display of the number of states in the Firewall Rules page
- Moved “Any” to top of protocol list in firewall rules
- Fixed issues with hidden fields on firewall_rules_edit.php #7057
- Fixed issues with moving rules that required scrolling while dragging #6895
- Enhanced ICMP type handling in rules
- Fixed issues when hovering the mouse pointer over aliases on disabled rules making the hint difficult to read #6448
- Fixed handling of firewall rule separators when a NAT associated rule is deleted #6676
- Added field to specify source-hash key for outbound NAT rules
- Fixed issues with **Firewall > NAT > Edit** forgetting destination type selection when input errors occur #6224
- Removed “self” as a destination from NAT 1:1 rules
- Fixed NAT rules so that when a port forward is disabled, its associated firewall rule is also disabled #6472
- Fixed 1:1 NAT address family validation #6927

- Fixed problems with nested aliases containing FQDNs #6982
- Changed the Status > Filter Reload page so it shows the entire filter reload progress, rather than only the last state #6931
- Fixed labels on diag_states_summary.php #6711
- Fixed initial state of confirmation checkboxes on diag_resetstate.php
- Changed Diag > States so it can optionally require a filter before displaying states, to improve handling with large state tables #7069

3.42.10 Traffic Shaping

- Added Chelsio network cards (cxl) to the list of drivers that are capable of using ALTQ #6830
- Fixed the traffic shaper wizard so it uses whole numbers instead of decimals #6779

3.42.11 HA / CARP

- Fixed issues when XMLRPC synchronizes IP Alias type Virtual IP addresses bound to Localhost #7010
- Fixed a bug where the CARP VIP status was incorrect when the interface has more than one CARP VIP

3.42.12 DHCP/DHCPv6 Server / Router Advertisements

- Updated the ISC DHCP Daemon to fix issues with missing hostnames in leases, and removed workarounds that are no longer needed #6840
- Fixed reversed behavior of "Change DHCPv6 display lease time from UTC to local time" #6640
- Fixed incorrect index for edit action on DHCP Leases #7233
- Added an option to force a Dynamic DNS hostname in DHCP/DHCP6 Server settings
- Changed DHCP lease times to always display in 24-hour clock format
- Added an option to allow BOOTP to be specifically disabled in the DHCP Server settings #4351
- Fixed validation to allow URLs for TFTP Server in DHCP Server settings #6634
- Improve dhcpd and dhcpleases reload handling
- Fixed DHCP NTP Server form validation to allow hyphens #6806
- Fixed restore of DHCP6 leases on full install when using MFS /var
- Fixed a problem with the DHCP range being reset if the Setup Wizard was re-run when a custom DHCP range already exists #4820
- Fixed issues with DHCP traffic being blocked with DHCP Relay enabled #6996
- Changed the DHCP/DHCPv6 server GUI so it can be configured (but not run) while DHCP Relay is enabled #6997
- Added Client ID to DHCP Leases display, if present
- Added Client ID to DHCP Mapping list, if present
- Disabled DHCP server on interfaces with subnet >= 31 #6930
- Changed DHCP6 client to allow a prefix size of /59

- Changed DHCP6 server to allow a prefix size of /59 and /61
- Added new “Ignore client identifiers” option to DHCP Server
- Fixed handling of DNS entries for IPv6 static mappings when using delegated prefixes #6768
- Improved the help text for Router Advertisement configuration #6889

3.42.13 DNS / Resolver / Forwarder

- Allow a variable number of DNS servers #5549
- Changed interface boxes in the DNS Resolver so they can be resized
- Fixed sorting of DNS Forwarder hosts and domains in config.xml #6903
- Fixed DNS Resolver (unbound) logging after clearing logs #6915
- Added support for “deny_non_local” and “refuse_non_local” ACLs in the DNS Resolver #6914
- Fixed DNS Server Gateway validation
- Changed behavior of DNS Resolver overrides to only add FQDN entries, not short hostnames #6064
- Fixed issues with DNS Resolver Host Overrides not being updated properly #6712

3.42.14 NTP / GPS

- Fixed display of Prefer/No Select checkboxes invisible when adding entries in NTP Server settings #6788
- Fixed handling of NTP IPv6 restrict clauses
- Fixed setting default NTP access restrictions when there are no custom restrictions #6454
- Fixed NTP status widget IPv6 address handling so addresses are not truncated #4815
- Fixed the NTP Orphan Mode stratum field #7034
- Fixed issues with NTP GPS status
- Fixed a case that could result in an empty ‘restrict’ line in the NTP configuration #7110
- Added a limit for NTP time source fields so they cannot exceed the maximum number saved to configuration #7164
- Fixed display and behavior issues with NTP ACLs #6984
- Improved parsing of GPS initialization and output, and add support for more GPS output formats and extended status
- Added an autocorrect tool for checksums on GPS initialization commands #7159

3.42.15 Captive Portal

- Changed Captive Portal MACs page to be sortable #6786
- Fixed handling of Captive Portal user bandwidth set to 0 #6872
- Changed Captive portal to send "Admin Reset" as termination cause when disconnecting a user from the WebGUI
- Added option to Captive Portal to include idle time in total session time
- Fix bandwidth limitation settings in Captive Portal MAC passthrough
- Fixed links to view current Captive Portal page for all interfaces #6391
- Converted Captive Portal active sessions to a sortable table
- Added code to hide the client MAC address column in Captive Portal status when MAC filtering is disabled, rather than displaying an empty column
- Added popup with session details to the Captive Portal active sessions list on the status page
- Added button to disconnect all Captive Portal users
- Worked around race condition between captiveportal_disconnect_all() and captiveportal_prune_old()
- Added locking to avoid race conditions between rc.prunecaptiveportal and captiveportal_disconnect_all()
- Reworked logging and RADIUS accounting when disabling a Captive Portal zone or rebooting
- Increased speed of captiveportal_disconnect_all()

3.42.16 Dynamic DNS

- Added the ability to change the URL queried by Dynamic DNS entries to check the external IP address (Services > Dynamic DNS, Check IP Services tab) #6591
- Added support for All-Inkl Dynamic DNS provider
- Added support for duiadns.net Dynamic DNS provider
- Added support for CloudFlare Proxy to Dynamic DNS
- Added Cloudflare Dynamic DNS IPv6 support #6623
- Fixed status checking on Dynamic DNS (RFC2136), updates were always considered successful even on failure #6357
- Fixed handling of multiple RFC2136 entries #6153
- Fixed links in RFC2136 entries in the Dynamic DNS widget #7126
- Fixed HTTP header processing for Dynamic DNS updates
- Fixed handling of custom IPv6 Dynamic DNS in the widget #6922
- Changed Cloudflare and Gratis plus Dynamic DNS to store passwords in base64
- Updated Route 53 Dynamic DNS to fix several reported issues #3973 #6751 #5054
- Fixed handling of ZoneEdit Dynamic DNS when used with a CARP VIP #6992
- Removed excess loops from the Dynamic DNS Widget

3.42.17 Gateways / Routing

- Added the ability to disable gateway monitoring actions without disabling gateway monitoring #3151
- Changed gateway notifications to notify by email and syslog when a gateway goes up or down
- Improved gateway notification mechanisms
- Fixed handling of deleting or disabling static default gateways so they are properly removed from the routing table #6659
- Fixed L2TP WAN dynamic gateway naming #6980
- Fixed status display for unmonitored gateways
- Fixed static blackhole route handling
- Fixed handling of long hostnames on Diagnostics > Routes #6869
- Corrected behavior of disabled static routes #3560
- Created a PHP Shell playback script to view the gateway status from the shell and status output #7046

3.42.18 Notifications

- Fixed SMTP settings test so it properly displays results
- Fixed validation of secure SMTP Connection Modes (SSL/TLS and STARTTLS are mutually exclusive)
- Removed validation of password mismatches when SMTP or Growl notifications are disabled #7129
- Changed format of file_notice() alerts in webgui for easier reading

3.42.19 Graphs / Monitoring

- Changed traffic graphs to use d3.js (Dashboard and Status > Traffic Graphs)
- Moved export button to heading for Status > Monitoring page
- Moved graph labels so long hostnames do not overlap as easily #6138
- Improved error checking in case JSON isn't returned when building graphs #6748
- Added a missing RRD step value to lookup table #6860
- Added support for multiple views in Status > Monitoring graphs (Adds tab shortcuts to different graph views)
- Added a per-view "Refresh Interval" option to Status > Monitoring graphs
- Fixed fix null acronyms and axis label for queues/queuedrops graph in Status > Monitoring
- Enabled Area and Bar graph types for Status > Monitoring graphs

3.42.20 WebGUI

- Added an option to allow display of the firewall hostname on the login page
- Added filtering to widgets where appropriate
- Standardized PHP memory limit configuration
- Fixed formatting issues with the Installed Packages widget #6601
- Improved Compact-RED theme
- Changed service running/stopped icons
- Fixed issues with JavaScript confirmation prompts missing words (e.g. “Are you sure you wish to?”) #6972
- Fixed issues with packages that toggle visibility of advanced options areas #7100
- Removed the crash reporter link from the dashboard when a user does not have crash_reporter page access #7043
- Fixed display of Package installation message #7226
- Fixed “” tag processing in package XML handling
- Fixed inconsistent handling of empty/null configuration settings in config.xml #6893

3.42.21 Logging

- Increased filtering tail limit for logging to ensure enough entries will be displayed #6652
- Added a means for packages to request a syslogd socket inside a chroot environment #4898
- Added BIND logging to proper facility #5524
- Improved handling of the TFTP Proxy/xinetd process when it is disabled, to reduce log messages #6308

3.42.22 Misc

- Updated simplepie (RSS Parsing library) to 1.4.3
- Fixed storing of IPv6 addresses so they are always saved in lower case #6864
- Fixed bsnmpd “printcap” log errors #6838
- Fixed a foreach error when restoring a configuration without packages
- Fixed handling of signal traps in the console menu #6741
- Fixed “Goto line #” action on diag_edit.php so pressing the enter key also activates the function
- Changed the PHP Execute feature of Diagnostics > Command so that it does not generate a crash report from a syntax error #6702
- Added enable link to Status > UPnP & NAT-PMP error message if disabled #6689
- Changed the time zone help text to clarify and warn against the use of the Etc time zones that use POSIX style signs, which are the opposite of what most users expect #7089
- Added validation to prevent duplicate Wake on LAN entries
- Fixed permissions on /var/tmp when /var is a RAM disk #7120
- Added a fallback for get_pkg_info() to use pkg info if there is no local copy of the repository catalog

- Removed spurious output from the PHP Shell executable when running a playback script from a command prompt #7045
- Updated status.php with new info and changed its output organization #7246

3.43 2.3.2-p1 New Features and Changes

3.43.1 2.3.2 Update 1

- FreeBSD-SA-16:26.openssl - Multiple vulnerabilities in OpenSSL. The only significant impact on pfSense® software is OCSP for HAproxy and FreeRADIUS.
- Several HyperV-related Errata in FreeBSD 10.3, FreeBSD-EN-16:10 through 16:16. See https://www.freebsd.org/relnotes/10-STABLE/errata/errata.html for details.
- Several built-in packages and libraries have been updated, including:
  - PHP to 5.6.26
  - libidn to 1.33
  - curl to 7.50.3
  - libxml2 to 2.9.4
- The hardware serial number is now displayed in the system information widget, or a host UUID if a serial number is not found. This is for display purposes and facilitates users seeking support in identifying their hardware.
- Added encoding to the 'zone' parameter on Captive Portal pages.
- Added output encoding to diag_dns.php for results returned from DNS. #6737
- Worked around a Chrome bug with regular expression parsing of escaped characters within character sets. Fixes "Please match the requested format" on recent Chrome versions. #6762
- Fixed DHCPv6 server time format option #6640
- Fixed /usr/bin/install missing from new installations. #6643
- Increased filtering tail limit for logging so searching will locate sufficient entries. #6652
- Cleaned up Installed Packages widget and HTML. #6601
- Fixed widget settings corruption when creating new settings. #6669
- Fixed various typos and wording errors.
- Removed defunct links to the devwiki site. Everything is on https://www.netgate.com/docs/pfsense/ now.
- Added a field to CA/Cert pages for OU, which is required by some external CAs and users. #6672
- Fixed a redundant HTTP "User-Agent" string in DynDNS updates.
- Fixed the font for sortable tables.
- Added a check to verify if an interface is active in a gateway group before updating dynamic DNS.
- Fixed wording of the "Reject leases from" option for a DHCP interface (it can only take addresses, not subnets.) #6646
- Fixed error reporting for SMTP settings test.
- Fixed saving of country, provider, and plan values for PPP interfaces
- Fixed checking of invalid "Go To Line" numbers on diag_edit.php. #6704

- Fixed off-by-one error with “Rows to Display” on diag_routes.php. #6705
- Fixed description of the filter box on diag_routes.php to reflect that all fields are searchable. #6706
- Fixed description of the box for the file to edit on diag_edit.php. #6703
- Fixed description of the main panel on diag_resetstate.php. #6709
- Fixed warning dialog when a box is unchecked on diag_resetstate.php. #6710
- Fixed log shortcut for DHCP6 areas. #6700
- Fixed the network delete button showing when only one row was present on services_unbound_acls.php #6716
- Fixed disappearing help text on repeatable rows when the last row is deleted. #6716
- Fixed dynamic DNS domain for static map DHCP entries
- Added control to set dashboard widget refresh period
- Added “-C /dev/null” to the dnsmasq command line parameters to avoid it picking up an incorrect default configuration which would override our options. #6730
- Added “-l” to traceroute6 to show both IP Addresses and Hostnames when resolving hops on diag_traceroute.php. #6715
- Added note about max ttl/hop limit in source comment on diag_traceroute.php.
- Clarified language on diag_tables.php. #6713
- Cleaned up the text on diag_sockets.php. #6708
- Fixed display of VLAN interface names during console assignment. #6724
- Fixed domain-name-servers option showing twice in pools when set manually.
- Fixed handling of DHCP options in pools other than the main range. #6720
- Fixed missing hostnames in some cases with dhcpdv6. #6589
- Improved pidfile handling for dhcpleases.
- Added checks to prevent accessing an undefined offset in IPv6.inc.
- Fixed the display of the alias popup and edit options on source and destination for both the address and port on outbound NAT.
- Fixed handling of backup config count. #6771
- Removed some dangling PPTP references that are no longer relevant.
- Fixed up/caught up remote syslog areas. Added “routing”, “ntpd”, “ppp”, “resolver”, fixed “vpn” to include all VPN areas (IPsec, OpenVPN, L2TP, PPPoE Server). #6780
- Fixed missing checkboxes in some cases when adding rows on services_ntpd.php. #6788
- Revised service running/stopped icons.
- Added a check to CRL management to remove certificates from the drop-down list that are already contained in the CRL being edited.
- Fixed rule separators moving when multiple firewall rules are deleted at the same time. #6801

3.44 2.3.2 New Features and Changes

3.44.1 SSH Daemon

Note: The ssh host keys were made more secure, and if a client remembers an older, weaker key, the ssh client may refuse to connect. Remove the older key and then make the ssh client learn the new key.

- Changed sshd to use stronger Key Exchange algorithms and disabled some older, weaker algorithms. Clients may need to be updated to handle the new Key Exchange methods.

  Currently allowed Key Exchange Algorithms: `curve25519-sha256@libssh.org`, `diffie-hellman-group-exchange-sha256`
- Removed the ECDSA host key from the sshd configuration
- Added ED25519 host key to the sshd configuration
- Changed the list of available ciphers.

  Current allowed ciphers: `chacha20-poly1305@openssh.com,aes256-gcm@openssh.com, aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr`
- Changed the list of available Message Authentication Code methods,

  Current MAC list: `hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com, hmac-ripemd160-etm@openssh.com,umac-128-etm@openssh.com,hmac-sha2-512,hmac-sha2-256, hmac-ripemd160,umac-128@openssh.com`

3.44.2 Backup/Restore

- Don't allow applying changes on interface mismatch post-config restore until the reassignment is saved. #6613

3.44.3 Dashboard

- Dashboard now has per-user configuration options, documented in *User Manager*. #6388

3.44.4 DHCP Server

- Disabled dhcp-cache-threshold to avoid bug in ISC dhcpd 4.3.x omitting client-hostname from leases file, which makes dynamic hostname registration fail in some edge cases. #6589
- Note that DDNS key must be HMAC-MD5. #6622

3.44.5 DHCP Relay

- Imported fix for dhcrelay relaying requests on the interface where the target DHCP server resides. #6355

3.44.6 Dynamic DNS

- Allow * for hostname with Namecheap. #6260

3.44.7 Interfaces

- Fix "can't assign requested address" during boot with track6 interfaces. #6317
- Remove deprecated link options from GRE and gif. #6586, #6587
- Obey "Reject leases from" when DHCP "Advanced options" is checked. #6595
- Protect enclosed delimiters in DHCP client advanced configuration, so commas can be used there. #6548
- Fix default route on PPPoE interfaces missing in some edge cases. #6495

3.44.8 IPsec

- strongSwan upgraded to 5.5.0.
- Include aggressive in ipsec.conf where IKE mode auto is selected. #6513

3.44.9 Gateway Monitoring

- Fixed "socket name too large" making gateway monitoring fail on long interface names and IPv6 addresses. #6505

3.44.10 Limiters

- Set pipe_slot_limit automatically to maximum configured qlimit value. #6553

3.44.11 Monitoring

- Fixed no data periods being reported as 0, skewing averages. #6334
- Fix tooltip showing as "none" for some values. #6044
- Fix saving of some default configuration options. #6402
- Fix X axis ticks not responding to resolution for custom time periods. #6464

3.44.12 OpenVPN

- Re-sync client specific configurations after save of OpenVPN server instances to ensure their settings reflect the current server configuration. #6139

3.44.13 Operating System

- Fixed pf fragment states not being purged, triggering "PF frag entries limit reached". #6499
- Set core file location so they can't end up in /var/run and exhaust its available space. #6510
- Fixed "runtime went backwards" log spam in Hyper-V. #6446
- Fixed traceroute6 hang with non-responding hop in path. #3069
- Added symlink /var/run/dmesg.boot for vm-bhyve. #6573
- Set net.isr.dispatch=direct on 32 bit systems with IPsec enabled to prevent crash when accessing services on the host itself via VPN. #4754

3.44.14 Router Advertisements

- Added configuration fields for minimum and maximum router advertisement intervals and router lifetime. #6533

3.44.15 Routing

- Fixed static routes with IPv6 link local target router to include interface scope. #6506

3.44.16 Rules / NAT

- Fixed "PPPoE Clients" placeholder in rules and NAT, and ruleset error when using floating rules specifying PPPoE server. #6597
- Fixed failure to load ruleset with URL Table aliases where empty file specified. #6181
- Fixed TFTP proxy with xinetd. #6315

3.44.17 Upgrade

- Fixed nanobsd upgrade failures where DNS Forwarder/Resolver not bound to localhost. #6557

3.44.18 Virtual IPs

- Fixed performance problems with large numbers of virtual IPs. #6515
- Fixed PHP memory exhaustion on CARP status page with large state tables. #6364

3.44.19 Web Interface

- Added sorting to DHCP static mappings table. #6504
- Fixed file upload of NTP leap seconds. #6590
- Added IPv6 support to diag_dns.php. #6561
- Added IPv6 support to filter logs reverse lookup. #6585
- Package system - retain field data on input error. #6577
- Fixed multiple IPv6 input validation issues allowing invalid IPv6 IPs. #6551, #6552
- Fixed some DHCPv6 leases missing from GUI leases display. #6543
- Fixed state killing for 'in' direction and states with translated destination. #6530, #6531
- Restore input validation of captive portal zone names to prevent invalid XML. #6514
- Replaced calendar date picker in the user manager with one that works in browsers other than Chrome and Opera. #6516
- Restored proxy port field to OpenVPN client. #6372
- Clarify description of ports aliases. #6523
- Fixed translation output where gettext passed an empty string. #6394
- Fixed speed selection for 9600 in NTP GPS configuration. #6416
- Only allow IPv6 IPs on NPT screen. #6498
- Add alias import support for networks and ports. #6582
- Fixed sortable table header wrap oddities. #6074
- Clean up Network Booting section of DHCP Server screen. #6050
- Fix "UNKNOWN" links in package manager. #6617
- Fix missing bandwidth field for traffic shaper CBQ queues. #6437

3.44.20 UPnP

- UPnP presentation URL and model number now configurable. #6002

3.44.21 User Manager

- Prohibit admins from deleting their own accounts in the user manager. #6450

3.44.22 Other

- Added PHP shell sessions to enable and disable persistent CARP maintenance mode. "playback enablecarpmaint" and "playback disablecarpmaint". #6560
- Exposed serial console configuration for nanobsd VGA. #6291

3.45 2.3.1 New Features and Changes

3.45.1 Security/Errata

- FreeBSD Security Advisories
  - FreeBSD-SA-16:17 OpenSSL
  - FreeBSD-SA-16:18 atkbd
  - FreeBSD-SA-16:19 sendmsg
- **OpenVPN upgraded from 2.3.10 to 2.3.11. Fixes two potential security** issues.
  - OpenVPN 2.3.11 Change Log
- pfSense® Software Advisories
  - pfSense-SA-16_03.webgui
  - pfSense-SA-16_04.filterlog
  - 2.3.1 update 1 patches pfSense-SA-16_05.webgui.

3.45.2 Config Upgrade

- Fixed config upgrade for CARP VIPs on gateway groups, GRE and gif for uniqid format. #6222
- Fixed config upgrade for IP aliases with CARP IP parent. #6164
- Correct OpenVPN topology config upgrade to retain 2.2.x and prior net30 topology. #6140
- Correct and adjust apinger parameters to dpinger parameters automatically on upgrade. #6142

3.45.3 Gateways

- Fix static route for IPv6 monitor IP with link-local gateway. #6353
- Fix default gateway switching with IPv6 and link-local gateways. #6258

3.45.4 OS / Backend

- NanoBSD is now permanent read-write, to avoid issues with slow rw->ro mount times and systems getting stuck read-only mounted. #6184
- Systems using a RAM disk for /var/ have their alias tables backed up and restored during bootup. #6189
- Set console settings (serial configuration, password protection, etc.) post-upgrade. #6120
- Ensure package repo is updated with latest metadata when checking for latest version. #6115
- Display consistent firmware version on dashboard and in update checker. #6320
- Correct description of update branch options. #6136
- Prevent update checking failures from killing webGUI. #6177
- Make pkg use configured proxy server settings where they exist. #6149

3.45.5 Web GUI

- Fix row delete button on unsaved aliases, NTP, UPnP and other screens. #6101
- Captive portal MAC passthrough credits waiting period box restored. #6290
- Outbound NAT edit screen destination field alias auto-completion restored. #6287
- Captive portal allowed IPs direction selection on edit fixed. #6267
- Restored input validation on port forwards to prohibit IPv6. #6265
- Restored input validation on firewall rules to prohibit IPv6 IPs in IPv4 rules and vice versa. #6211
- Fixed PHP error on edit of PPP interfaces. #6264
- Fixed radio button placement on gateways dashboard widget settings. #6259
- Fixed display post-refresh of system information dashboard widget. #6251
- Restored in/out bytes counters on Status>Interfaces. #6244
- Correctly show and hide OpenVPN topology field as applicable. #6236 #6214
- Correct voucher character set input validation. #6231
- Disable background update checking on dashboard update check is disabled. #6212
- Restore input validation of IP address family and rule type, verifying IPv6 IPs with IPv6 rules, and IPv4 for IPv4 rules. #6218
- Add validation of address family and protocol combinations on packet capture page. #6219
- Add validation of IP aliases with CARP parent interfaces to ensure matching address family. #6218
- Restore GET parameters on status_graph.php. #6192
- Fixed PHP error on input validation failure with floating rules in some cases. #6175
- Use CDATA for firewall rule separator descriptions so non-English characters work. #6174
- Fix port forward edit destination field filling when virtual IPs configured. #6173
- Fix load balancer monitor edit. #6171
- Restore “none” in load balancer fall-back pool. #6170
- Restore use of aliases in load balancer. #6169
- Fix duplicate for load balancer pools and virtual servers. #6168
- Restore description field on lagg edit page. #6163
- Fix saving of bogons update frequency. #6162
- Restore description field on captive portal IP passthrough. #6161
- Fix saving of sticky connections timeout field. #6146
- Show all restore areas in backup/restore screen. #6144
- Fix moving of rule separator before saving. #6128
- Use consistent up and down arrow formats on dashboard widgets. #6123
- Fix typo on OpenVPN server description. #6102
- Fix missing string on notification “mark as read” button. #6104
- Fix firewall rule separator positioning with easy rule addition. #6105

- Prevent closing of info box on monitoring page. #6106
- Add custom date range option to monitoring page. Use infoblock on IPsec PSK screen. #6107
- Fixed loss of "Do not NAT" enable on edit on outbound NAT. #6112
- Correct label of 1:1 NAT edit screen. #6114
- Add AJAX updates to NTP status page. #6117
- Fix button spacing on Edit File and Command pages. #5995
- Fix specification of port in DNS Resolver domain overrides. #6091
- Fix moving of multiple items to bottom of list on firewall, NAT and IPsec screens. #6092
- Fix setup wizard with only WAN assigned and using static IP. #6093
- Remove logo from wizard since it's now redundant. #6095
- Fix gateway widget cut-off with 3 column dashboard. #6096
- Fixed force update on RFC 2136 DDNS. https://redmine.pfsense.org/issues/6359
- Fix reboot prompt when changing RAM disk setting and encountering an input error. #6349
- Fix highlighted tab when editing IPsec mobile P1. #6341
- Fix selection of configured speed and duplex on interface page. #6331
- Fix division by zero in status_queues.php. #6329
- Fix alignment issues in forms. #6327
- Fix entry of CIDR range in host aliases for conversion to IPs. #6322
- Allow use of # and ! again in DNS Forwarder domain overrides. #6310
- Restored hostname infobox in menu bar. #6306
- Fixed editing and deleting of additional DHCP pools. #6303
- Fixed requests to diag_system_activity.php piling up on slow systems. #6166

3.45.6 Interfaces

- Unset LAN DHCPv6/RA configuration if LAN interface is removed. #6152

3.45.7 IPsec

- Fix starting of strongswan twice. #6160

3.45.8 DNS Resolver

- Switched domain overrides from stub-zone to forward-zone so domain overrides don't require the target server provide recursion. #6065
- Allow adding 0.0.0.0/0 to access lists. #6073
- Added 100,000 and 200,000 options for Unbound cache limit. #6230
- Fix Unbound startup where both DNS Forwarder and Resolver are enabled. #6354

3.45.9 DHCP Server

- Hostnames now allowed for NTP servers. #6239

3.45.10 IPsec

- Fixed LAN interfaces stopping functioning when IPsec is in use. #6296
- Mobile PSK matching issue with multiple PSKs fixed. #6286
- leftsendcert=always specified for all RSA types. #6082
- rc.newipsecdns fixed to check correct enabled status. #6351

3.45.11 Notifications

- Fixed growl notifications to unresolvable hostname generating crash report. #6187
- Fixed growl notification test with no password. #6221

3.45.12 Captive Portal

- Fixed error handling captive portal username with single quote. #6203
- Fixed issues with mixed-case zone names. #6278

3.45.13 OpenVPN

- Prevent leading space in tunnel network configuration causing invalid configuration. #6198

3.45.14 User Manager

- Fix RADIUS login with attribute class (25) when the server returns multiple attribute entries with different data. #6086
- Honor deny config write for RADIUS users. #6088

3.45.15 Package System

- Uninstall all packages pre-upgrade from <= 2.2.x to 2.3 to avoid problems from old packages. Reinstall them post-upgrade. #6137
- Fix reinstall of renamed packages post-upgrade to 2.3. #6118
- Fix package reinstallation getting stuck in loop when there is no Internet connectivity post-upgrade. #6180

3.45.16 Other

- Removed lua support from nginx to not deprecate old CPUs lacking CMOV support. #6185
- Added validation to console menu interface assignment to prevent creating duplicate VLANs. #6183
- Blacklisted S.M.A.R.T. options with Hyper-V to prevent crash. #6147
- Silence SSH host key log spam. #6143
- Fix order of gateway and gateway group name in gateway down log message. #6134
- Allow use of @ in hostname field for Namecheap DDNS. #6122
- Fix console error where $nat_if_list isn't an array. #6307
- Include patch number in version display. #6309
- Fix pw groupdel error in log during boot. #6352
- Fixed stale xmlrpc.lock preventing config sync from functioning. #6328
- Fixed failed chown on startup with /var as a RAM disk. #6131
- Crash reporter now ignores warnings in release versions. #6178
- Fixed crash reporter to show full PHP warnings in development versions. #6097

3.45.17 Update 1

2.3.1 update 1 (2.3.1_1) was released on May 25, 2016 with the following fixes/changes since 2.3.1-RELEASE.

- Security issue pfSense-SA-16_05.webgui patched.
- Lowered default LDAP timeout from 25 seconds to 5 seconds. #6367
- Fixed handling of IPsec negotiation mode with IKE version set to auto. #6360
- Increase PHP's memory limit to 512 MB on 64 bit versions to better accommodate systems with a large number of active states. #6364
- Set request_terminate_timeout the same as max_execution_time to prevent many possible circumstances of "504 gateway error" from occurring. #6396
- Fix use of URL IP type aliases in firewall rules. #6403
- Fix show/hide fields Javascript in Chrome on macOS. #6401
- Fixed save of "IPv6 over IPv4 Tunneling" address on System>Advanced, Networking. #6381

3.45.18 Update 2 through 4

These were internal-only versions that weren't publicly-released.

3.45.19 Update 5

2.3.1 update 5 (2.3.1_5) was released on June 16, 2016 with the following fixes/changes since 2.3.1_1.

- Fixed command injection vulnerability in auth.inc via User Manager. #6475
- Fixed command injection vulnerability in pkg_mgr_install.php id parameter. #6474
- Upgraded PHP to 5.6.22
- Fixed Captive Portal redirect hangs caused by longer keepalive_timeout in nginx. #6421
- Fixed DDNS PTR zone in dhcpd.conf with third octet of 0. #6413
- Fixed save and reset buttons on load balancer status page. #6254
- Fixed schedule editing on firewall rules page. #6428
- Allow "-" character in TFTP server field on DHCP Server page. #6433
- Allow "-" and "_" characters in system tunables. #6438
- Fixed changing of link type on PPPs edit screen. #6439
- Fixed setting of "RADIUS issued IPs" on L2TP page. #6440
- Restored apply changes button for interface mismatch post-config restore. #6460
- Fixed display of Outbound NAT port aliases. #6463
- Fixed schedule edit allowing invalid time range. #6468

3.46 2.3 New Features and Changes

3.46.1 Security/Errata

- FreeBSD Security Advisories:
  - FreeBSD-SA-16:01.sctp
  - FreeBSD-SA-16:02.ntp
  - FreeBSD-SA-16:05.tcp
  - FreeBSD-SA-16:07.openssh
  - FreeBSD-SA-16:09.ntp
  - FreeBSD-SA-16:11.openssl
  - FreeBSD-SA-16:12.openssl
  - FreeBSD-SA-16:15.sysarch
- pfSense® Security Advisories:
  - pfSense-SA-16_01.webgui
  - pfSense-SA-16_02.webgui

Several obsolete items were removed from this release. The items are noted again in the sections below, but worth emphasizing:

- The PPTP **VPN Server** has been completely removed. The protocol has been broken for over three years.

  The PPTP WAN client remains for use with ISPs still using PPTP.

- Layer 7 classification support has been removed from the traffic shaper.

  It was rarely used, had been broken for all of 2.2.x, had absurdly high CPU usage, and snort filters better/faster

- WEP support has been removed from Wireless interfaces. #5123

  No reason to still be using this in this day and age. If it is still needed, use external AP.

- Single DES support has been removed from IPsec (3DES remains).

  It should not be used, it is not secure.

- 1GB NanoBSD images have been removed, as they were not large enough to proper accommodate the system and upgrade data. The supported sizes for NanoBSD images are now 2GB and 4GB.
- The default system password hash has been changed to bcrypt. Current passwords will continue to work. Existing users need to reset their password to convert to the new hash. More info below under "Authentication". #4120
- The LiveCD platform has been removed. The ISO is a bootable installer, as always, but it cannot run a live system.
  - The installer ISO image is now named "pfSense–RELEASE-.iso", with the .iso extension signifying the type of image it is (optical media installer).
  - For the very few people who were still using LiveCD, if the hardware can boot from USB, install to a USB thumb drive and run from it instead. If the options to keep /var and /tmp in RAM are active, and no packages are installed, the net result should be similar but ultimately more functional.

3.46.2 Dashboard/Widgets/GUI

- Converted GUI to the Bootstrap framework, completely new look
- Changed the GUI and Captive Portal web server to nginx; removed lighttpd. #5719
- Cleaned up a lot of GUI code, option text, etc
- TLS v1.0 disabled for the GUI. #5984
- Removed old style themes, introduced new CSS-based themes
- Refactored JavaScript and CSS, moved included items to more convenient locations
- Added more AJAX updating in widgets and other places
- Changed to more intuitive and modern icons and action buttons rather than the old confusing icon set (now using font-awesome icons)
- Changed log display to be more consistent (single page for most logs, common filtering options)
- Removed obsolete fifolog support. It was never used or fully implemented, and had no GUI option.
- Improved notices in the GUI
- Made breadcrumbs and page title handling more consistent
- Added an option to have the top menu follows the user when scrolling
- Renamed several GUI file names to match menu structure. #5628
- Fixed AES-NI hardware display in the system information widget. #4911
- Added widescreen support to the Dashboard. #5195
- Improved password field handling security. Stored passwords are not presented back to the user in HTML. A masked value is returned instead. All password fields have also been changed to require confirmation.
- Many pages have been reworked for improved internationalization

- Changed info box functions, removed print_info_box_np, now print_info_box and print_apply_box are used to print appropriate boxes without problematic automatic detection
- Moved RRD graphs to **Status > Monitoring** #5498
- Changed RRD GUI interface to D3 rather than using the RRD graph command, so that a newer rrdtool base could be used with minimal added dependencies. #5498
- Monitor IP added to gateways widget. #4782
- Increased max_input_vars from 1000 to 5000 to accommodate larger aliases. #4780
- Fixed NTP RRD graphs to accommodate negative values. #4423

3.46.3 OS/Backend

- Moved to a FreeBSD 10.3-RELEASE base
- Added tryforward() support to get (nearly all of) the performance of fastforward with IPsec enabled
- Overhauled the build system
  - Eliminated the -tools repository
  - Removed Patches, changes are now applied a vendor branch of FreeBSD
  - Rewrote/changed the build scripts significantly
  - Moved the new build scripts to the main pfSense repository
- PHP Upgraded to 5.6
- Replaced pecl-APC with opcache. #4744
- Added support for -c parameters to /etc/rc.initial. #4422
- Added optional package for kernel debug symbols. #5330
- Rewrote system_set_harddisk_standby() for the current CAM-based ATA stack. #4569
- Fixed a Panic/Crash with “sbflush_internal: cc 4294967166 || mb 0 || mbcnt 0”. #4689
- Fixed a kernel panic with AES-NI. #4702
- Updated AES-GCM/AES-NI bits from FreeBSD -HEAD. #4841
- Removed zoneinfo.tgz file for Time Zones, move to the same format as FreeBSD. #4726
- Fixed tcpdump with zerocopy enabled (net.bpf.zerocopy_enable=1). #5257
- Added ability to disable PV disks and NICs on Xen. #5452
- Removed the built-in but unused MySQL PHP modules and added them to the pkg server instead. They may be added as package dependencies or manually installed as needed.
- Followed FreeBSD (r294560) in ceasing generation of rsa1 and dsa ssh server host keys by default
- Removed support for nanobsd images < 2GB #5836
- Overhauled IP address handling code in various parts of the system
- scponly package is included by default. #5190
- Shortened F1 boot prompt delay on nanobsd. #3426

3.46.4 Packages

Note: The list of available packages in pfSense 2.3 has been significantly trimmed. Netgate has removed packages that have been deprecated upstream, no longer have an active maintainer, or were never stable.

- Removed use of PBI-based packages, moved to pkg(ng)
- Fixed installation and handling of packages to use pkg, now works identically in the GUI and shell/console
- Changed packages to use the FreeBSD ports format/layout to work with pkg
- XMLRPC calls for package information and installation have been removed, replaced with native pkg functions. #4575
- Added support for packages to be (re)built automatically by Poudriere
- Added search capability to Available Packages list to filter packages by keywords. #5324
- Fixed the version comparison code in the Package manager. #4924
- Added support for tags in listtopic fields for use by packages
- Factory reset now completely uninstalls packages. #5829
- Improved handling of package install post-upgrade. #3597

3.46.5 System Updates

- Major changes to update management
- Removed "full update" or "full slice" upgrade for systems on 2.3 to later versions

  These files will remain available for use by older versions updating to 2.3.
- The "Full Backup" feature has been deprecated.
- Changed system updates to be handled via pkg
- Changed Base, kernel, and standard pre-installed binares to packages
- Removed "Firmware" nomenclature, now only referred to as "Update"
- Fixed updating of base to work the same from the console or the GUI
- Added preliminary support for restarting system services without rebooting in cases when the base is updated but the kernel is the same.

3.46.6 Gateways/Routing

- Replaced apinger with dpinger(!). #5624
  - This fixes many gateway monitoring related issues, including incorrect latency and loss in various edge cases.
  - Eliminates status file race conditions that caused update failures on services bound to gateway groups in some edge cases. #5180 and #3818 among others.
  - Fixed gateway monitoring startup at boot time with assigned OpenVPN interfaces. #4587
  - Check gateway monitor settings after upgrade, dpinger has different options than apinger.
- Added code to allow gateways outside of an interface subnet. #972

- Corrected "State Killing on Gateway Failure" description. #4709
- Fixed disabling of a static route set to use a disabled gateway. #4813
- Added standard deviation to gateway status and widget
- Fixed dynamic gateway logic to prevent GIF/GRE from making dummy/unusable gateways that show up for monitoring/routing/etc #5766
- Changed static routes handling for DNS servers so they are removed when a gateway is disabled #4921
- Increased gateway weight limit from 5 to 30. #5843
- Fixed issues with PPP type WANs and the Default Gateway Switching option. #1837
- Fixed dynamic gateway handling for OpenVPN tap clients. #5981
- Fixed display of full interface name in Diagnostics>Routes. #5484

3.46.7 Rules/NAT/pf

- Added drag-and-drop rule reordering for firewall and NAT rules.
- Fixed a situation where pf drops IPv6 packets with fragment header followed by a last fragment only. #2762
- Fixed "LAN network" in v6 rules not working when a link-local address is assigned to LAN. #3656
- Added reordering for 1:1 NAT rules. #3888
- Improved handling of firewall rule tracker IDs for port forward associated rules
- Added support for a separator bar in firewall and NAT rules for use as a visual reference. #5373
- Standardized the NPt options in the GUI so their options and appearance are more similar to 1:1 NAT
- Added a "no binat" checkbox to 1:1 NAT screen for exclusions. #3887
- Limited pfsync syncpeer to IPv4 since it does not support IPv6 #4648
- Changed the default CARP pass rules to use "no state" to avoid issues with broken L2 gear that duplicates packets #5800
- Added sorting to Alias lists #4195
- Added a hit counter to the firewall rule display with states and bandwidth consumed by packets matching rules.
- Fixed issues with the DNS Forwarder and DNS Resolver being enabled concurrently (on different ports) in an HA environment #5882
- Added a visual indication in the rule list for floating rules with the "quick" property set #5860
- Improved state display on **Diagnostics > States**, now shows packets and bytes for each state
- Fixed aliases containing both FQDNs and IPv6 subnets. #5872
- Fixed removal of downloaded URL table alias contents when alias is deleted. #5856
- Significantly improved validation of downloaded data for URL Table aliases. #5848
- Fixed possibilities for creating an invalid ruleset with missing URL Table Ports aliases. #5845
- Fixed filterdns issues with significant system clock time jumps. #4166
- Added firewall rules hit counter. #3504

3.46.8 Interfaces/VIPs

- Fixed pfSense_getall_interface_addresses truncating IPv6 link local IP addresses. #4062
- Add GUI setting for VLANs PCP. #4133
- Fixed GRE interfaces failing to have a RUNNING state after reboot. #4191
- Fixed setting non-default MTUs in some edge cases. #4397
- Added input validation on bridges to prevent adding the same interface to multiple bridges. #4595
- Fixed CARP not working under bhyve. #4623
- Improved input validation for 6RD, GRE and gif interfaces, helping prevent invalid configurations.
- Changed input validation to allow /31 to be used for CARP VIPs since that is now supported and works in FreeBSD. #5533
- Added debug logging option for DHCP6 client. #4534
- Fixed cases where DHCP6 client (dhcp6c) was being launched multiple times in some circumstances. #5621
- Upgraded dhcp6c. #5734
- Upgraded DHCP client to ISC dhcpd 4.3.3P1.
- Fixed applying of non-default MTU on gif interfaces post-boot with dynamic IP WANs. #5842
- Added support for PPPoE with MTU/MRU > 1492, RFC 4638. #4542
- Fixed issues with link cycling on some Intel 10G ix NICs #5913
- Corrected ALTQ test to show that ix/ixgbe NICs are capable of traffic shaping. #5923
- Improved handling of default interface assignment for some hardware. #4535
- Corrected input validation for invalid IPv6 IPs with leading or trailing colon. #6024
- Fixed orphaning of VLANs on lagg interfaces after editing the lagg. #6014
- Fixed loss of some dhcpleases and dhcpleases6 logs. #5968
- Fixed adding of routes immediately post-reboot for delegated IPv6 prefixes to sub-routers. #5957
- Fixes to DHCPv6 leases status page and prefixes.php. #5944 #4206
- Fixed loss of IPv6 IP on track6 interfaces when saving and applying changes on that interface. #5945
- Fixed incorrect interface mismatch prompt post-config restore when using VLANs on lagg. #5892
- Added support for multiple span interfaces on bridges. #5871
- Prevent naming conflicts between interfaces and interface groups. #5795
- Prevent naming conflicts between interfaces and aliases. #5778
- Fixed use of IP aliases with GRE tunnels. #4450
- Fixed application of bridge advanced options after interface added to bridge. #4312
- Set MTU back to default after clearing the field. #3926
- Fixed IPv6 IP aliases on CARP IPs. #3716
- Fixed IP alias on CARP IPs where IP alias above CARP parent in list. #3257
- Fixed modifying unassigned VLAN interfaces changing assigned VLAN. #3209

3.46.9 Authentication

- Fixed the WebGUI becoming slow or unusable when an LDAP server used for GUI auth is unreachable. #3383
- Fixed a problem with using 'local' as the name of an authentication server 'Descriptive Name'. #4469
- Fixed default Auth Server selection on system_usermanager_settings.php. #5440
- Added support for bcrypt as a passwd hash and enabled it as the system default #4120
- Replaced the default passwd hash for root/admin using bcrypt (blowfish).
  - Existing user passwords will continue to work in their existing format until the user's password is changed.
  - User passwords cannot be automatically converted as they are not stored plain text. To convert the password hash of an existing user to bcrypt, edit the user and change their password.
- Added the ability to filter privileges when adding them to a user or group, to make finding them easier.
- Fixed updating of group file for renamed groups. #6013
- Fixed handling of groups with spaces in their names. Local group names can no longer contain spaces. New group scope option "Remote" added for LDAP and RADIUS use where spaces in group names are valid. #6012
- Added support for RFC2307 style LDAP groups. #4923

3.46.10 Services

- Fixed handling of the SNMP Bind Interface. #3883
- Fixed ntpd crashes on 32 bit with dynamic WAN reconnections and OpenVPN client configured. #4155
- Fixed a kernel panic with APU and SNMP with mibII. #4403
- Updated igmpproxy to the latest version. #4672

  The old version had some custom patches, so be wary of behavior changes
- Added encoding for DHCP/DHCPv6 server additional BOOTP text options to preserve data when stored in XML #5623
- Fixed duplication action for Load Balancer Monitor entries #4441
- Upgraded DHCP Server and Relay to ISC dhcpd 4.3.3P1
- Added statistics gathering for DHCP Server leases. #5387
- Fixed DDNS key issues with DHCP and DHCPv6 Server enabled on multiple interfaces. #5603
- Added custom ACLs for NTP (restrictions by network) #4463
- Prevent starting of radvd in circumstances where it shouldn't. #5812
- Added description column to DHCP leases status screen. #5729
- inetd replaced with xinetd (used for proxy mode NAT reflection and TFTP proxy). #5707
- DHCP lease counters added to Status>DHCP Leases. #5186
- Allow configuration of RAs when DHCPv6 Relay is enabled. #6063
- Fixed DHCPv6 Server's DDNS. #4675
- DHCP Server menu item now defaults to the first interface with an enabled DHCP Server instance. #4647
- Allow configuring DHCPv6 and RAs on track6 interfaces. #3029

- Fixed RADIUS NAS IP in PPPoE server. #185
- Deprecated ntpdate_sync_once.sh, replacing with ntpd -g. #6053

3.46.11 DNS

- Fixed Unbound IPv6 link local handling. #4021
- Added validation for advanced configuration directives in Unbound. #4411
- Upgraded dnsmasq to 2.76.0test8 to fix crashes in 2.75. #5341
- Fixed Unbound binding to IP alias virtual IPs. #5464
- Changed Namecheap dynamic DNS to use separate hostname and domain name fields #4366
- Added Multi-WAN support to RFC 2136 Dynamic DNS.
- Added RFC 2136 support to the Dynamic DNS widget
- Added input validation to prevent the same DNS server from being added multiple times on **System > General** #5915
- Fixed CloudFlare dynamic DNS to not configure 'proxiable' and 'proxied' parameters. #6005
- Fixed dnsmasq host overrides when both DNS Forwarder and Resolver are enabled. #5883
- Added RFC 2136 dynamic DNS to dashboard widget. #5862
- Added multi-WAN support to RFC 2136 dynamic DNS client. #5862
- Don't specify 127.0.0.0/8 IPs as forward-addr in Unbound configuration. #5750
- Added input validation to require configured DNS servers before enabling Resolver's forwarding mode. #4747
- Added Google Domains DDNS support. #4322
- Added DNS Made Easy DDNS support. #1258
- Allow @ in Dynamic DNS hostnames. #3900
- Improve IPv6 link local handling in DNS Resolver and Forwarder so it works across configuration restores and with HA config sync. #3802

3.46.12 IPsec

- Upgraded to strongSwan 5.4.0.
- Fixed multiple possibilities for IPsec status hangs. #5520
- Revised handling of IPsec reloading when strongswan.conf is changed. #4353
- Fixed problems with the search domain in IPsec mobile clients. #4418
- Added support for elliptic curve for IPsec on webconfigurator. #4683
- Added input validation for authentication backend when using EAP-RADIUS with IKEv2 Mobile IPsec. #5219
- Fixed unit display on IPsec status pages for time and data to be more human-friendly. #5364
- Removed support for single DES from IPsec #5543 (3DES Remains)
- Removed global IPsec disable flag as it is no longer necessary. On upgrade, if the IPsec enable box was unchecked, all Phase 1 entries are disabled individually instead.
- Changed IPsec 'up' commands to start in the backgound so they are non-blocking #5882

- Disabled the strongSwan unity plugin by default, and improved the method used to disable the plugin #4178
- Removed unnecessary and troublesome 'pass out' rules for mobile IPsec #5819
- Fixed "no valid leases object found" log spam with IPsec dashboard widget. #5855
- Fixed automatically added WAN rules (UDP 500, 4500, ESP) when using IPsec with IP aliases. #5500
- Fixed IKEv2 to Cisco ASA resulting in traffic selector mismatch when initiated by traffic. #4719
- Added "split connections" option to phase 1 for IKEv2 for interoperability with third party devices that do not support multiple traffic selectors on one child SA (Cisco ASA, others). #4704
- Added dynamic AJAX update to status_ipsec.php. #6049

3.46.13 OpenVPN

- Changed the default behavior of the OpenVPN server to use topology subnet, not net30. #5526
- Changed Client-Specific Overrides so they can be set to apply to specific servers rather than being globally set. #5526
- Fixed OpenVPN Server validation of self-signed certificates with a depth of 2. #4329
- Fixed overwriting of custom `/etc/dh-parameters.*` on upgrade. #4816
- Fixed invalid rules generated with some AVPair-defined ACLs. #5451
- Improved display of server certificates on OpenVPN servers to help avoid users incorrectly picking user certificates for servers. #5602
- Fixed OpenVPN client specification of auth-user-pass in shared key modes where it's not valid. #5941
- Fixed problems with OpenVPN and some use of special characters in the username or password. #4605

3.46.14 MPD/PPP VPN/Services

- Removed PPTP Server. #4226
- Add MS-CHAPv2 option to L2TP Configuration. #4732
- Fixed editing of multiple PPPoE connections with dial on demand enabled changing the port assignment. #4378
- Added a user login count option to the PPPoE server

3.46.15 UPnP/NAT-PMP

- Enabled port-in-use checking in miniupnpd. #4320
- Enabled IPv6 for miniupnpd. #4321
- Set secure_mode=yes in miniupnpd configuration #5627

3.46.16 Wireless

- Removed WEP. #5123
- Improved default settings for Wireless interfaces

3.46.17 Captive Portal

- Fixed Captive Portal to support more than 120 VLAN interfaces. #4150
- Added an option in Captive Portal for FreeRADIUS-friendly stop/start RADIUS accounting updates that solves problems with user session time limits. #2164
- Fixed selection of RADIUS NAS IP with VIPs when editing Captive Portal zone. #5656

3.46.18 Traffic Shaping

- Fixed CODELQ scheduler defaults. #4692
- Removed Layer 7 classification support from the traffic shaper #5508
- Relaxed the shaper wizard interface validation when there are no interfaces with gateways selected #4524
- Fixed traffic shaper failure with "bandwidth for q… higher than interface" in some edge cases. #5721

3.46.19 Misc

- Allow wildcards in Certificate Subject Alternative Names. #3733
- Removed the "Certificate Authority" option on the **Certificates** tab of the Cert Manager when creating a **Certificate**. To make a Certificate Authority, use the **CAs** tab instead. #5924
- Adapted gitsync to new repo structure. #4999
- Changed the packet capture output in the GUI so that when the protocol is set for CARP, tcpdump interprets it as CARP for more accurate output
- Added pfsync protocol option to packet capture page. #5866
- Added "GoTo line #" control to Diagnostics > Edit File
- Corrected help in pfSsh.php to properly reflect how recording works
- Fixed validation of playback file passed to pfSsh.php #5657
- Fixed disabling of filter.log logging where local logging is disabled. #6018
- Updated included software on licenses.php page. #5903
- Internationalization improvements. #5777
- Fixed use of IP aliases on Test Port page. #5185
- Fixed key map, screen map and font selection in installer. #4387
- Prevent deletion of certificates in use by packages. #4142

3.46.20 Update Patches

This section lists the changes contained in patch updates post-release.

2.3_1

The 2.3_1 update upgrades NTP to fix FreeBSD security advisory SA-16:16.ntp. The only change is upgrading ntpd from 4.2.8p6 to 4.2.8p7.

3.47 2.2.6 New Features and Changes

3.47.1 Security/Errata Notices

- Updated to FreeBSD 10.1-RELEASE-p25
  - FreeBSD-SA-15:26.openssl Multiple vulnerabilities in OpenSSL
- Updated to strongSwan 5.3.5
  - Includes fix for CVE-2015-8023 authentication bypass vulnerability in the eap-mschapv2 plugin.
- pfSense-SA-15_09.webgui: Local File Inclusion Vulnerability in the pfSense® WebGUI
- pfSense-SA-15_10.captiveportal: SQL Injection Vulnerability in the pfSense captive portal logout
- pfSense-SA-15_11.webgui: Multiple XSS and CSRF Vulnerabilities in the pfSense WebGUI

3.47.2 Logging

- Fixed log duplication for some log entries. #5606

3.47.3 IPsec

- strongSwan 5.3.5 update fixes several bugs.

3.47.4 Config sync

- Fixed config synchronization failure in some circumstances. #5509

3.47.5 Captive Portal

- Fixed captive portal database handling issue that could reset database instead of waiting for lock to clear. #5622
- Fixed problem with 0 byte files in captive portal file manager. #5642

3.48 2.2.5 New Features and Changes

3.48.1 Security/Errata Notices

- Updated to FreeBSD 10.1-RELEASE-p24
  - FreeBSD-SA-15:25.ntp
  - FreeBSD-SA-15:14.bsdpatch:
  - FreeBSD-SA-15:16.openssh:
  - FreeBSD-SA-15:18.bsdpatch:
  - FreeBSD-SA-15:20.expat:
  - FreeBSD-SA-15:21.amd64:
  - FreeBSD-SA-15:22.openssh:
- **pfSense-SA-15_08.webgui:**
  Multiple Stored XSS Vulnerabilities in the pfSense® WebGUI

  The complete list of affected pages and fields is listed in the linked SA.
- Updated strongSwan to 5.3.3
- Updated PHP to 5.5.30
- Updated miniupnpd to 1.9.20150721 to address a potential vulnerability in miniupnpd.

3.48.2 User Management/Authentication

- Added support for GUI auth from RADIUS to obtain group names from the RADIUS reply attribute "Class" as a string (local groups must exist, similar to LDAP). #935
- Added an LDAP server timeout field to address GUI access issues when the LDAP server is down/unreachable. #3383
- Added support for LDAP RFC 2307 style group membership. #4923
- Worked around a chicken-and-egg problem in user syncing which was preventing users from using ssh the first time the account was saved. #5152
- Prevent deletion of system users and groups by authenticated, authorized users using manually crafted POSTs. #5294

3.48.3 OpenVPN

- Fixed an incorrect netmask being sent to OpenVPN clients with static IP addresses set in RADIUS. #5129
- Changed the calculation of the OpenVPN point-to-point server IP address obtained from RADIUS to be consistent with CSC/Overrides (Server should be one IP address below the Client)

3.48.4 IPsec

- strongSwan upgraded to 5.3.3. strongSwan's change log
- Fixed missing DH group 22-24. #4918
- Fixed handling of IPv4 IPsec Phase 1 endpoints that resolve to an IPv6 address. #4147 (Fixed by strongSwan update to 5.3.3)
- Brought back "auto" IKE version and fixed problems with its previous implementation.
- Pre-shared keys configured as "any" under VPN>IPsec, Pre-Shared Keys tab are added as %any to ipsec.secrets now, as described in the note on the page. #5246
- Resolved memory leak by switching printf hooks to vstr. #5149
- Change to vstr to fix memory leak broke SMP status plugin. Switched to vici for status output.
- ID selectors omitted from ipsec.secrets for mobile PSK+XAuth configurations. Fixes pre-shared key mismatches with Apple iOS Cisco IPsec and other mobile clients. #5245
- Fixed logging default settings and ability to set logging to silent. #5340
- Logging settings applied correctly on clean start and stop/start of service. #5242
- Remove deleted CAs, certificates and CRLs from strongswan configuration. #5238
- Prevent over-matching of auto-added firewall rules for mobile IPsec configurations. #5211
- Added IPv6 virtual address pool support for mobile. #5284
- Allow both IPv4 and IPv6 in phase 2 entries on a single phase 1 when using IKEv2. #5305
- Omit NAT rules for disabled phase 1 and 2 configurations. #5320
- Only display certificate authority field for methods where it's relevant. #5323
- Only write out CA certificates for those specified in a Phase 1 configuration. #5243
- Fixed Hybrid RSA + xauth. #5207
- Fixed configuration of split tunnel attribute. #5327
- Specify rightca in ipsec.conf where relevant. #5241
- Specify leftsendcert=always in ipsec.conf for mobile profiles using IKEv2 to better accommodate iOS and macOS manual configurations. #5353
- Fix IKEv2 mobile client pool status display with small number of active leases

3.48.5 Rules/NAT

- Fixed handling of url_port alias types when processing items that should be handled by filterdns. #4888
- Fixed handling of line endings when parsing a URL table ports file.
- Fixed handling of empty bogon lists on NanoBSD.
- Fixed handling of 6rd rules so they are only added when there is an IPv4 IP defined for the gateway, otherwise the ruleset ends up invalid. #4935
- Added support for port ranges on Outbound NAT. #5156
- Added a check to prevent renaming an alias to an existing name. #5162
- Improved the fix for increasing the "self" table size in pf.

- Imported fixes from FreeBSD for a situation that could result in a panic/crash due to source address limits in pf rules ("pf_hashsrc: unknown address family 0"). #4874

3.48.6 Captive Portal

- Implemented an alternate method to find VIP targets that should be allowed for Captive Portal. #4903
- Improved handling of the captive portal database files for zones in cases when the database files may be corrupt or unreadable. #4904
- Improved handling of vouchers that are too short. In certain cases they were not being properly rejected. #4985
- Fixed handling of voucher database files, initializing the database properly when necessary. #5113
- Fixed loading of allowed hostnames at boot time. #4746, #5345

3.48.7 Packages

- Fixed handling of package install errors and connect timeouts during the install process. #4884
- Improved package version comparison. #4924
- Fixed an issue with package editing where the default value was not being populated for new fields.
- Fixed removal of syslog.conf entries during package uninstall #5210

3.48.8 DHCP

- Fixed handling of DHCP pools that are out of range, preventing them from creating an invalid dhcpd configuration. #4878
- Added support for UEFI network booting with arch 00:09. #5046
- Fixed a situation where dhcpleases could miss updates for hostnames in the leases file, delaying functional hostname resolution of new and updated DHCP leases. #4931
- Automatically add firewall rules to permit DHCP traffic when DHCP Relay is enabled, matching the behavior for DHCP Server. #4558

3.48.9 Interfaces

- Fixed identification of IPv6 interfaces with PPP-type interfaces and DHCP6 #3670
- Removed "Could not find gateway for interface…" log messages as they were largely useless. #4102
- Added OpenVPN interfaces to the list of available interfaces when reassignment is necessary during config.xml restoration.
- Fixed interface assignment menus running off VGA screen.
- Fixed preservation of MLPPP settings when saving interface settings. #4568
- Correct handling of SLAAC, DHCP6 and DHCP-PD with PPP interfaces. #5297

3.48.10 Dynamic DNS

- Fixed Cloudflare support for Dynamic DNS updates.
- Fixed GratisDNS support for hosts without subdomains.
- Disabled DHS provider. It had never worked.
- Fixed IPv4 dynamic DNS registrations on dual stack hosts to providers with AAAA records. #3858
- Update Dynamic DNS using gateway groups upon enable and disable of gateways. #5214
- Fixed Dynamic DNS using gateway groups specifying a CARP IP. #4990

3.48.11 Misc

- Fixed the configuration version comparison in XMLRPC sync to prevent more invalid synchronization cases. #4902
- Cleaned up old unused platforms referenced in a few areas of the code that were no longer relevant.
- Fixed killing of individual states in cases when the source and destination were reversed. #4907
- Fixed killing of individual states for IPv6. #4906
- Changed the "enableallowallwan" script to also allow bogons, which makes the use of RFC 5735 / RFC 6890 test networks easier in lab environments.
- Fixed handling of VIPs in source address selection for Diagnostics > Test Port. #4986
- Updated status.php to include more information. #5304
- Fixed handling of the description in Traffic Shaping.
- Fixed pfSense base version comparison. #4925
- Fixed handling of multiple notices in the same second. #4879
- Removed the routed service as it is being handled by the package.
- Set MIME type for SVG in lighttpd configuration.
- Improved handling of the cron service reconfiguration process.
- Added option to display monitor IP on Gateways widget #4782
- Added "Description" as a display option on Traffic Graphs. #4783
- Fixed handling of L2TP server interface selection. #4830
- Added /usr/bin/dc back into the build. #5111
- Fixed a crash/panic "Sleeping thread owns a non-sleepable lock" in ARP code when using Proxy ARP type VIPs. #4685
- Added support for Sierra Wireless 7355. #4863
- Updated time zones. #5254
- Added fsync of Unbound's root.key to ensure the file isn't corrupted if power is lost shortly after writing of the file. Code added to detect corrupt root.key and delete and recreate it. #5334
- Fix changing outbound NAT modes and uploading/downloading files on exec.php with non-English languages. #5342, #5343
- Associate intermediate internal CA certificates with the signing CA. #5313

3.49 2.2.4 New Features and Changes

3.49.1 Security/Errata Notices

- pfSense-SA-15_07.webgui: Multiple Stored XSS Vulnerabilities in the pfSense® WebGUI

  The complete list of affected pages and fields is listed in the linked SA.

- FreeBSD-SA-15:13.tcp:
- Further fixes for file corruption in various cases during an unclean shut down (crash, power loss, etc.). #4523
  - Fixed pw in FreeBSD to address passwd/group corruption
  - Fixed config.xml writing to use fsync properly to avoid cases when it could end up empty. #4803
  - Removed the 'sync' option from filesystems for new full installs and full upgrades now that the real fix is in place.
  - Removed softupdates and journaling (AKA SU+J) from NanoBSD, they remain on full installs. #4822

Note: The forcesync patch for #2401 is still considered harmful to the filesystem and has been kept out. As such, there may be some noticeable slowness with NanoBSD on certain slower disks, especially CF cards and to a lesser extent, SD cards. If this is a problem, the filesystem may be kept read-write on a permanent basis using the option on **Diagnostics > NanoBSD**. With the other above changes, risk is minimal. The best practice is to replace the affected CF/SD media by a new, faster card as soon as possible. #4814

- Upgraded PHP to 5.5.27 to address CVE-2015-3152 #4832
- Lowered SSH LoginGraceTime from 2 minutes to 30 seconds to mitigate the impact of MaxAuthTries bypass bug.

Note: *Sshlockout* will lock out offending IP addresses in all past, current and future versions. #4875

3.49.2 Certificates

- Changed the built-in certificate manager to specify keyUsage and extendedKeyUsage in certificates. Windows will now correctly function with IKEv2 using certificates from the built-in certificate manager without disabling EKU. #4580

Note: This change applies only to new certificates, created on 2.2.4 or newer, and the CN of the certificate must match the hostname or IP address to which clients connect.

- Added authorityKeyIdentifier to CRLs generated by the built-in certificate manager. (strongSwan requires it to match.) #4860

3.49.3 IPsec

- Fixed non-GCM AES modes with AES-NI enabled. #4791
- Fixed issues with keyid and some mobile IPsec identifiers. #4811 #4806
- Fixed includes so PHP shell session restartipsec script works.
- Fix dashboard hardware crypto display where AES-NI is enabled. #4809
- Fixed issues with IPsec with certificates/ASN1.DN. #4792 #4794
- Added code to write out CRLs from the built-in certificate manager for use by strongSwan.
- Added option for enabling Strict CRL Checking (strictcrlpolicy in strongSwan config).
- Fixed saving Advanced IPsec options before IPsec is enabled.
- Changed LAN bypass to be from "LAN subnet" to "LAN subnet" rather than from "LAN subnet" to "LAN address" to allow it to work for VIPs on the interface.
- Remove "Auto" key exchange option, and change upgraded configurations to IKEv2. #4873
- Specify rightid for mobile IPsec non-PSK configurations. Add peer ID option "any" for excluding peer identifier checks for mobile IPsec circumstances where peer ID matching is impossible or undesirable.

3.49.4 OpenVPN

- Fixed handling of OpenVPN automatic stop/start when bound to gateway groups using CARP VIPs. #4854

3.49.5 DHCP

- Fixed issues with IPv6 Prefix Delegation caused by an invalid prefix/subnet check added to the ISC DHCP daemon. Reported upstream and patched the checks out in FreeBSD ports. #4829

3.49.6 DNS Resolver

- Changed Unbound to use interface-automatic where interface list is empty so it replies correctly in a default HA configuration. #4807
- Fixed selection of a CARP VIP for outgoing interface. #4852
- Fixed some inconsistencies in text across the GUI in places that specified DNS Forwarder vs. Resolver. #4551

3.49.7 Load Balancer

- Improved handling of port ranges in relayd. #4810
- Fixed references to Load Balancer Virtual Server *redirect_mode*.

3.49.8 Traffic Shaping

- Fixed adding of VoIP rules from traffic shaper wizard where IP/alias was not specified. #4838
- Fixed default CoDel values.
- Corrected inverted target/interval values for CoDel.

3.49.9 Rules/NAT/Aliases

- Fixed a foreach() error when saving an empty alias.
- Fixed input validation on Alias import page.
- Fixed inconsistencies in descriptions in Alias editing for URL Table aliases.
- Added labels to more default firewall rules.
- Avoid an error loading the rules with a numeric hostname in an alias.

3.49.10 Misc

- Removed unnecessary deletion of rc.conf; Added an empty rc.conf with a note.
- Removed a check for a QinQ interface existing when deleting. The check unnecessarily made QinQ un-deletable where the parent interface no longer existed.
- Fixed GratisDNS support.
- Fixed glob for serial devices to match more accurately.
- Fixed a foreach() warning when editing PPP entries.
- Fixed GRE and GIF interface input validation so required fields and descriptions match.
- Changed the behavior of Cancel buttons to be consistent (return to referring page).
- Fixed display of advanced DHCP settings when present.
- Removed old, unused *NetUtils.js*.
- Retain */usr/bin/fsync* from FreeBSD in images.
- Added “netstat -ni” to /status.php output.
- Fixed a typo in upgrade code for Captive Portal.
- Fixed limiter upgrade code to allocate pipe numbers even if no rules are present.
- **Fixed upgrade code to remove old CA/Cert config entries that were** moved/relocated.

3.50 2.2.3 New Features and Changes

3.50.1 Security/Errata Notices

- pfSense-SA-15_06.webgui: Multiple XSS Vulnerabilities in the pfSense® WebGUI

  The complete list of affected pages and fields is very large and all are listed in the linked SA.

- FreeBSD-SA-15:10.openssl:
- Fixes for filesystem corruption in various cases during an unclean shut down (crash, power loss, etc.). #4523
  - Changed new filesystems to use the 'sync' option to avoid loss of data.
  - Added upgrade code to activate the 'sync' option on the root slice for existing installations.
  - Changed new filesystems to use softupdates and journaling (AKA SU+J).
  - Changed the way fsck is handled at boot time:
    * Followed best practice of using fsck from FreeBSD rc.d/fsck script. (Run preen mode first and later try forcefully fixing issues.)
    * Added as much information during boot on the status of the filesystem as possible.
    * Changed fsck to run with -C flag and always in foreground during boot to prevent issues that might schedule background mode.

Note: The forcesync patch for #2401 was considered harmful to the filesystem and removed. As such, there may be some noticeable slowness with NanoBSD on certain slower disks, especially CF cards and to a lesser extent, SD cards. If this is a problem, the filesystem may be kept read-write on a permanent basis using the option on **Diagnostics > NanoBSD**.

3.50.2 Rules/Aliases/NAT

- Fixed a problem with more than 64 IP addresses in the "self" table in pf.
- Fixed issues with FQDNs in aliases causing static entries to be lost. #4296
- Added the tracker ID rule number lookup to dynamic firewall log. #4730
- Fixed alias rename and delete not being propagated to outbound NAT. #4701
- Fixed tracker IDs of policy route negation rules which had been duplicating the tracker ID of the rule they were based upon. This confused the log parser and displayed the negation rule rather than the actual rule. #4651
- Fixed logging of passed IGMP traffic when the rule is not set to log. #4383
- Fixed a situation where a combination of L2TP, overlapping subnets, port forwards and NAT reflection could cause an invalid ruleset. #4772
- Added a GUI field to control the size of the pf fragment limit #4775

3.50.3 IPsec

- Updated strongSwan to 5.3.2. #4750
- Integrated a patch from https://wiki.strongSwan.org/issues/951 to solve IPsec SA rekey issues on strongSwan+FreeBSD. #4686
- Added patches from FreeBSD PR 200282 to help address IPsec rekey issues.
- Backported FreeBSD r283146 and patch from FreeBSD PR 192774 to address PF_KEY ACQUIRE missing port and protocol information.
- Added reply-to/route-to rules for mobile-ipsec. #4235
- Removed the manual specification of reqid in the IPsec configuration because strongSwan 5.3.0 has fixed issues with its handling, which caused the existing code to misbehave. #4665
- Fixed the display and behavior of the LAN bypass option for IPsec. #4655
- Fixed IPsec LAN bypass toggling every time save is pressed. #4640
- Changed how charon is started and restarted to fix a various issues with IPsec configuration reloading. #4268
- Added new modes for IPsec Phase 1 according to RFC 5903 (Ecliptic Curve groups). #4260
- Implemented the "make before break" feature available in strongSwan 5.3.0, which is useful for IKEv2. #4626
- Fixed vpn_ipsec_configure so it always performs a filter reload to ensure the ruleset is updated where necessary in every IPsec change scenario. #4631
- Added support for EAP-RADIUS to IKEv2 Mobile Clients. #4614
- Fixed a panic/crash when accessing services on the firewall over mobile IPsec on 32-bit installations (set net.inet.ipsec.directdispatch=0 on i386). #4537
- Fixed an issue with FQDN hosts and PSKs. #4785

3.50.4 OpenVPN

- Added a space to the OpenVPN TLS Verify script to avoid appended parameters appearing the same as existing parameters.
- Fixed get_interface_ip() to return the IP address correctly for gateway groups specifying a VIP, which fixed OpenVPN clients not working with gateway groups specifying VIPs. #4661
- Changed the OpenVPN client settings to allow just one of either the username or password to be specified. #3633
- Fixed OpenVPN servers listening on an associated IPv6 addresses.

3.50.5 Captive Portal

- Fixed filterdns to use the proper API for ipfw changes on FreeBSD 10.1+ to correct captive portal allowed hostnames not being loaded into tables at boot time. #4746
- Fixed Captive Portal RADIUS accounting. #4131
- Fixed Captive Portal Idle-Timeout causing a value of 2147483647 for acctsessiontime. #4652
- Fixed disconnection of active voucher users, and corrected disconnection of users especially when triggered via XMLRPC. #4625

3.50.6 Operating System

- Fixed both the kernel and choparp to better handle I/O and prevent issues in the way it handles BPF, which can contribute to a panic when using Proxy ARP VIPs. #4685
- Merged a patch that avoids a panic on sockbuf module. #4689
- Fixed AESNI to be SMP friendly to avoid various decryption errors and possible encryption mistakes. Also present critical_enter/critical_exit to avoid preemption of the currentrunning thread which should fix panics. #4702
- Updated time zone data from FreeBSD 10.1-RELEASE. #4459
- Fixed creation of */var/spool/lock* on NanoBSD at boot time. #4532
- Removed boot_serial='yes' from loader.conf when serial is disabled. #4617
- Fixed an issue where mtree would fail during an upgrade from a previous version of FreeBSD when moving to 2.2.x. #4653

3.50.7 Interfaces/NIC Drivers

- Added support for Sierra Wireless MC7354.
- Added support for Intel X552, ixgbe changes from stable/10, and moved altq changes for ixgbe to the large ixgbe patch.
- Enabled ix/ixv/ixl modules in the kernel
- Fixed duplication of statistics on vlan(4) interfaces for outgoing bytes #3314
- Fixed updating wireless statistics so that the output bytes are not always zero. #4028
- Added a patch from FreeBSD PR 200722 for mpd5 to preventing it from printing a warning when renaming an interface to an existing name.
- Fixed SLAAC/DHCPv6 handling for cases where the global SLAAC IPv6 address might be present when using DHCPv6. #4483
- Corrected descriptions on Key Rotation and Master Key Regeneration for wireless interfaces.
- Removed the "insert my MAC" feature from interfaces.php.
- Defined $var_path as a global key since it is being used in interfaces.inc, but it was not declared.
- Fixed issues setting the MTU on certain interfaces. #4397

3.50.8 Packages

- Fixed various issues with PBI generation.
- Synchronized and cleaned up various pfPorts, eliminated several that had changes pushed back into FreeBSD ports.
- Fixed an issue where rebuild_package_binaries_pbi.php could fail due to missing build files. #4600
- Backported patches from FreeBSD stable/10 to fix a crash when stopping squid. #4592
- Fixed pfflowd to use the correct version for parsing the new pfsync header and corrected the pfsync version check. #4304
- Updated pkg_edit.php with fixes for usecolspan2 and combinedfields.

- Fixed pagination on pkg.php.
- Fixed boot-time log file initialization for package logs. #4603

3.50.9 DHCP/RA

- Clarified that DNS Forwarder and Resolver both apply in DHCP/DHCPv6 and router advertisements. #3730
- Removed unnecessary filtering on the DHCP static mappings table.
- Added appropriate RA Flags for "Stateless DHCP".
- Added error checking to avoid warnings about DHCP relay during boot.
- Fixed hostname validation for static DHCP leases such that only fully qualified hostnames must be unique, not only short names.
- Fixed adding DHCP static mappings from the DHCP leases view to non-default pools. #4649
- Stopped invalid DHCP settings from being applied when input errors exist.
- Removed DHCP static lease overlap cleanup and its associated function and killing of the DHCP daemon. This behavior could cause problems with failover scenarios, especially when adding/editing/removing static mappings.

3.50.10 Web GUI

- Fixed language selection. #4705
- Changes to status.php to make it easier to gather and submit support information:
  - Added sanitization of OpenVPN static/tls keys to status.php.
  - Cleaned up, organized, and expanded the info presented by status.php.
  - Changed status.php to additionally save the output to individual
  - text files and compress them into a .tgz for later download.
- Fixed setup wizard LAN DHCP pool calculation to avoid an invalid pool.
- Improved the setup wizard hostname check. #4712
- Fixed some minor text issues in wizards.
- Changed the wizard to use the current WAN gateway name rather than assuming the name. #4713
- Updated and corrected the wireless status flags and capabilities list. There are many more possible flags, now documented at *Wireless Status*.
- Added a fall back to look up local user privileges and groups if the groups could not be found from LDAP and there is a local user.
- Fixed Crash Reporter submissions when symlinks were present as part of crash report, which would fail to save the report on the server. #4650
- Set a user agent for the Crash Reporter.
- Cleaned up code logic in status_upnp.php.

3.50.11 CARP

- Changed CARP so that it does not trigger a carp demotion taskqueue if the value is 0, which can cause the cluster to misbehave.
- Fixed issues for CARP+Bridges where pfSense would crash or freeze. #4607
- Fixed the CARP plugin call for packages. The "interface" parameter was coming through as NULL during CARP events.
- Added INIT event for CARP in devd.conf as an alternate for 'backup', otherwise scripts would not take down services during a MASTER->INIT transition. (e.g. interface unplug, link loss)
- Fixed NTP so that it properly uses selected CARP IP addresses. #4370
- Fixed CARP packet flow after initial interface creation. #4633

3.50.12 Traffic Shaper/Limiters

- Fixed limiters when used with IPv6. #2526
- Corrected handling of NAT when RDR/BINAT is applied on packet and it is being sent to limiters. #4596

3.50.13 DNS

- Consistently handle clear_subsystem_dirty after an Unbound restart.
- Added a call to clear_subsystem_dirty('staticmaps') when using Unbound, otherwise DHCP static mappings would not fully apply when Unbound was in use. #4678
- Fixed an Unbound warning when "dnsallowoverride" was off and port forwarding was on. #4682
- Re-enabled verification for selfhost DynDNS since their chain issue has been resolved. #4545

3.50.14 Misc

- Updated PHP to 5.5.26
- Fixed various issues in the installer for GEOM mirrors (mirror slice detection, gmirror cleanup on non-clean disks.) #4658
- Fixed new user creation to use skel as the source of new user files rather than copying from the home directory of *root*.
- Changed growl so it will not be called if the configured address isn't an IP address or resolvable hostname. This avoids 1 minute timeout delay in fsockopen in growl.class. This change cuts that down to about a 20 second timeout. #4739
- Added a reboot after restoring a full backup in the GUI. #4107
- Deprecated */usr/local/bin/3gstat* as it was no longer used. It was replaced by *3gstats.php* long ago.
- Started using the "host!" flag when setting CURLOPT_INTERFACE, as recommended by the CURL documentation.
- Started passing the interface to CURLOPT_INTERFACE instead of the IP address, also started using the "if!" flag to avoid CURL trying to resolve the interface name.
- Fixed NTP serial configuration to setup the serial port before attempting to configure a GPS unit.

- Cleaned up various HTML/XHTML issues.
- Fixed a check for deleting a VIP when in use by OpenVPN.
- Fixed issues with backup/restore of a config.xml breaking the serial console on ADI installs. #4720
- Fixed several issues with boot speed when WAN was disconnected. #4442
- Removed some unused/obsolete files.

3.51 2.2.2 New Features and Changes

3.51.1 Security/Errata Notices

- pfSense-SA-15_05.webgui: Multiple XSS Vulnerabilities in the pfSense® WebGUI
- FreeBSD-SA-15:09.ipv6:
- FreeBSD-SA-15:06.openssl:

3.51.2 Rules / NAT

- Added hidden config option to disable blocking of link-local IPv4 (169.254.0.0/16) for the rare instances where it's required. Not recommended, violates RFC 3927.
- Fixed invalid ruleset generation when using port forwards with destination "any" on a DHCP client WAN-type interface, have pure NAT mode reflection enabled, and have the interface with link up but unable to reach a DHCP server for an extended period. #4564
- Allow the use of version IPv4+IPv6 on firewall rules without restrictions on protocol. The former restrictions date back to earlier base software versions, and are no longer applicable.
- Omit route-to from rules specifying a specific gateway when that gateway is forced down. #4566
- Use the subnet address when forming rules for networks, rather than the interface IP address
- Added SCTP to the protocol drop-down for firewall rules

3.51.3 IPsec

- Enforce disabling of "prefer old SAs" option.Having this option enabled will cause connectivity problems after rekeying in many circumstances. Upgrading to 2.2.2 will fix this.
- strongSwan upgraded to 5.3.0
- Don't apply mobile IPsec phase 2 PFS configuration to non-mobile IPsec. #4538
- Correct applying of uniqueid configuration. #4359
- Bring back automatic exclusion of LAN subnet to LAN IP for scenarios where remote IPsec overlaps with local LAN subnet. #4504
- Enable ike_name for daemon logging, adding connection identifiers to IPsec logs that can be correlated to output of 'ipsec statusall' (GUI log viewer integration to come).

3.51.4 DNS Forwarder/Resolver

- Fix DNS registration of hostname "0" #4573
- Domain overrides to multiple server IPs are possible in DNS Resolver. Add message noting this, and how to achieve it. #4350
- Always configure user-specified DNS servers in the Unbound configuration, to make its behavior consistent with dnsmasq
- Only list nameservers once in resolv.conf

3.51.5 Wireless

- Atheros wireless driver updated to latest from FreeBSD 11-CURRENT. Not many changes since 2.2.1-RELEASE. #4582
- Wireless cards removed from ALTQ-capable interfaces (traffic shaper capability) since that isn't supported at the moment. #4406
- New option "auto" added for Standard. This omits configuring mode with ifconfig, which currently can trigger driver problems that don't exist when not specified. Standard "auto" is preferred, and possibly required, for BSS and IBSS wireless modes with Atheros cards (at a minimum, potentially others).

3.51.6 IPv6

- Make sure 'DHCPv6 Prefix Delegation size' is provided if 'Send IPv6 prefix hint' flag is checked to avoid generating invalid dhcp6c configuration file.
- DHCPv6 Relay fixed. #4572
- Allow "0" for id-assoc na ID, id-assoc pd ID, sla-id and sla-len DHCP6 configuration options. #4547
- Fix the use of multiple prefixes in IPv6 router advertisements. #4468

3.51.7 Other

- Clean up logic in OpenVPN resync code. Discussion here and additional change here.
- SSL certificate validation disabled for selfhost - their certificate chain had a problem that made OpenSSL fail verification, making the service non-functional. #4545 The provider fixed the issue after 2.2.2-RELEASE, so verification has been re-enabled for 2.2.3 and newer.
- Fix error in traffic shaping wizard. #4529
- Fix broken image path. #4530
- A variety of minor text clean up in web interface.
- Remove some code no longer used in a few places.
- Clean up of code path when adding a new user. #4620
- Make sure RRD backup is not restored when /var memory disk is not in use. #4531
- Show friendly name of the interface on custom RRD graph drop-down selection
- PHP upgraded to 5.5.23
- Prevent a user from adding a VLAN using the invalid ID "0"

- Cleanup display of times in DHCP leases
- Use the correct field for voucher “expired” and “no access” messages
- Fix traffic shaper wizard bandwidth input validation calculations https://redmine.pfsense.org/issues/4259
- Changed Diagnostics > Sockets to display sockets bound to localhost
- Allow single interface bridges, useful for span ports and when migrating interfaces to a bridge

3.52 2.2.1 New Features and Changes

3.52.1 Security/Errata Notices

- pfSense-SA-15_02.igmp: Integer overflow in IGMP protocol (FreeBSD-SA-15:04.igmp)
- pfSense-SA-15_03.webgui: Multiple XSS Vulnerabilities in the pfSense® WebGUI
- pfSense-SA-15_04.webgui: Arbitrary file deletion vulnerability in the pfSense WebGUI
- FreeBSD-EN-15:01.vt: vt(4) crash with improper ioctl parameters
- FreeBSD-EN-15:02.openssl.asc: Update to include reliability fixes from OpenSSL

Potentially Relevant

The following updates are included from upstream in FreeBSD, but are not directly relevant. Neither pfSense software nor its packages include SCTP services, but such services may have been manually added by the user.

- FreeBSD-SA-15:02.kmem: SCTP SCTP_SS_VALUE kernel memory corruption and disclosure
- FreeBSD-SA-15:03.sctp: SCTP stream reset vulnerability

Not Relevant

- OpenSSL “FREAK” vulnerability:
  - Does not affect the web server configuration on the firewall as it does not have export ciphers enabled.
  - pfSense 2.2 already included OpenSSL 1.0.1k which addressed the client-side vulnerability.
  - If packages include a web server or similar component, such as a proxy, an improper user configuration may be affected. Consult the package documentation or forum for details.

3.52.2 Known Issues

- Some cases remain where filterdns does not properly handle hostnames in multiple aliases properly. Most of the cases have been fixed, so the situation is better than 2.2-RELEASE, but it is not 100% resolved. See issue #4296 for details. Placing hostname aliases into a separate alias so they are not mixed with static entries effectively works around the issue.

3.52.3 General

- Updated the default SSL cipher list to be stronger, obsoletes the need for a "BEAST protection" option #4230
- Fixed gen_subnet_max returning an incorrect result on 32 bit (i386) versions, which in turn fixed Wake on LAN and other areas on 32 bit (i386) versions. #4318
- Fixed crash on boot with some hardware, caused by gpioapu on systems where smbios.system.product is null. Mostly seemed to be the recycled Watchguard users affected by this issue. #4363
- Updated ufslabels.sh to handle a wider variety of disk layouts.
- Added a choice of SMTP authentication protocols for notifications, Office365 mail support. #4176
- Removed latin-1 encoding of RSS feed to fix display issues of RSS items.
- Fixed an issue where the GUI setting for PAP or CHAP in L2TP Server was not being respected.
- Fixed changing source tracking value separate from changing the Sticky option.
- Added input validation to force a minimum *100000* byte log file size to prevent undersizing the logs.
- Added more cleanup to the **Restart PHP-FPM** console menu action.
- Removed PTR records for aliases in host overrides.
- Fixed *diag_arp.php* to allow underscore in resolved host names.
- Fixed an issue in DHCP settings where the "add routers" value was not being preserved across a loop for each interface.
- Added capability to handle reverse lookup domain overrides.
- Fixed issues with NTP RRD graph state changes.
- Added input validation to require RADIUS protocol and server IP address/host in Captive Portal when RADIUS authentication is selected. #4384
- Fixed swap size calculation in the installer to avoid creating improperly sized partitions in systems with lots of RAM but not much disk space.
- Fixed test for comconsole when matching for enabling serial console. #4464
- Updated pfSense PHP shell help to current configuration structure. #4492
- Fixed switching from a PPP type WAN to "None" or "DHCP".
- Disables SNMP hostres module on APU boards due to crashes. #4403
- Removed -U from mtree call used to restore files permissions as it was breaking symlinks on upgrade. #4328
- Added input validation for Wireless configurations to prevent problematic combinations of settings. #4178
- Improved handling of FQDN entries in aliases with filterdns, but not 100% resolved. #4296
- Fixed various typo, style, and formatting issues.

3.52.4 Rules / NAT

- Fixed ordering of DHCPv6 client and bogon rules so the bogon rules can't block DHCPv6 requests. #3395
- Fixed a bug where applying NAT changes in Hyper-V could break the running NAT configuration. #4445
- Fixed a bug where marking a packet with only a number resulted in a broken rule. #4274
- Fixed DSCP choices that were non-functional and resulted in a broken ruleset. #4302
- Fixed PHP memory exhaustion on NAT pages with VIP ranges on a 32 bit (i386) versions. #4317 (Related to #4318)
- Fixed input validation on Outbound NAT to accept a port range. #4300
- Removed Carrier-Grade NAT subnet from "Block private networks" as it was in 2.0.x and earlier releases since it specifically notes RFC 1918 and CGN is more closely related to bogon networks. #4379
- Removed code that set adaptive.start and end to 0, now they are left at their defaults (60% and 120% of the state limit, respectively) if not user-overridden.
- Added configuration options for state timeout values under System>Advanced, Firewall/NAT. #4509

3.52.5 IPsec

- Added MOBIKE control, now disabled by default. #3979
- Fixed page rendering so MOBIKE is only shown with IKEv2 selected, NAT-T only shown with IKEv1 selected.
- Removed *Prefer older IPsec SAs* option from the GUI, and existing configurations with it enabled will not have that setting applied. #4349
- Added input validation to prevent use of AES key lengths larger than 128-bit when the *glxsb* cryptographic accelerator is enabled. #4361
- Added an option for an IPsec tunnel to act as a responder only. #4360
- Added a filter reload when IPsec is disabled. #4245
- Fixed RSA cert handling in IPsec to use double quotes on asn1dn specification so it is properly interpreted by strongSwan. #4275
- Added an option to allow controlling unique ID handling in IPsec advanced settings. #4359
- Fixed *restartipsec* command line script.
- Fixed handling of IPsec with Gateway Groups #4482
- Added a workaround to disable the strongSwan Unity plugin. #4178
- Added error logging when an IPsec Phase 1 cannot be located.

3.52.6 OpenVPN

- Added encoding for username and password to avoid issues with special characters. #4340
- Fixed issues with OpenVPN TLS and authentication scripts. #4329
- Fixed issues with handling of the Authentication Mode if the user changes the value after changing other incompatible settings.

3.52.7 DNS Resolver

- Upgraded to Unbound 1.5.3.
- Added correct scaling of *rrset-cache-size* in *unbound.conf*. #4367
- Added support for 0x20 DNS random bit. #4205
- Changed DNS Resolver default values to be a bit more strict: Enable Hide Identity, Hide Version, Harden DNSSEC data.
- Force harden glue configuration option, and remove GUI control of that option. Problem with Unbound pre-1.5.2 means in 2.2-RELEASE, having this option enabled, and DNSSEC disabled, could lead to DNS cache poisoning. #4402
- Added a check to test if Unbound is enabled and using the same port before allowing dnsmasq to be enabled. #4332
- Removed hard-coded value for *harden-referral-path*. #4399

3.52.8 Logging

- Fixed GUI log parser handling for IGMP log entries. #4343
- Fixed syslogd issues where the daemon stopped and failed to restart during boot in some cases. #4393

3.52.9 Traffic Shaping

- Fixed input validation errors in the Traffic Shaper wizard due to old data not being cleared. #4333
- Fixed handling of Upstream SIP Server in the Traffic Shaper wizard. #4314, #4427
- Fixed crash when using limiters and pfsync. #4310
- Fixed limiters used with IPv6. #2526

3.52.10 IPv6

- Fixed calculation of the 6rd default gateway honoring netmasks other than /32.
- Fixed recording of the IPv6 interface's new IP address and do not issue commands that cannot succeed. #3669
- Fixed not being able to save custom and custom-v6 DynDNS entries.
- Added IPv6 IP addresses to */etc/hosts* in the same manner IPv4 IP addresses are added. #4395
- Fix computation of the displayed DHCPv6 range start to be consistent with the actual check.
- Added *dhcp6.name-servers* option with DHCPD-PD regardless of PD length.

- Fixed Net_IPv6::compress() to properly handle all-zeros address.
- Enabled `UnicastOnly` in radvd for `ovpnX` interfaces. #4455
- Removed requesting a prefix delegation when there are no tracking interfaces setup to use it. #4436
- Added code to destroy stf interface when a 6rd or 6to4 tunnel is disabled. #4471

3.52.11 VIP/CARP

- Added input validation to prevent the VIP "interfaces" from being assigned since they are just an identification of the VIP for tracking and not actual interfaces. #4389
- Fixed functions to properly return the VIP subnet now that the CARP might not match its parent interface subnet. #4390
- Fixed a bug that caused the status icon from previous CARP VIP to be shown in cases where the IP address was not present on an interface.
- Changed the carp demotion factors slightly to avoid CARP transitions that are most likely unnecessary. (Do not demote on NIC send errors or pfsync errors)
- Expanded the CARP demotion error
- Added button to reset demotion status
- Fixed handling of IP Alias deletion from a secondary node using XMLRPC configuration sync #4446

3.52.12 Misc Binary/OS Changes

- Upgraded PHP to 5.5.22.
- Re-enabled Suhosin in PHP.
- Updated 802.11 code and Atheros wireless driver from FreeBSD 11-CURRENT
- Added patch to fix crash with Ralink wireless cards in access point mode. #4117
- Added athstats, cryptostats and cryptodev back. #4239
- Fixed AESNI module checks when used inside a virtual machine.

3.53 2.2 New Features and Changes

3.53.1 Special Notes

Due to CSS and JavaScript changes, forcing the browser to clear its cache or reload the pages after an update is advised. This is especially true if any cosmetic anomalies are observed, such as alignment problems or spurious bits of text in widgets.

3.53.2 Security Fixes

- Update to openssl 1.0.1k to address FreeBSD SA-15:01
- Multiple XSS vulnerabilities in web interface pfSense-SA-15_01
- OpenVPN update for CVE-2014-8104
- NTP update FreeBSD-SA-14:31.ntp though these circumstances don't seem to impact pfSense® software.

3.53.3 Default Configuration Changes

- **DNS Resolver (unbound) enabled for new installs**. #3396
- **DNS Forwarder (dnsmasq) disabled for new installs**. #3396
- **Change default NICs from vr to em** – vr is on the way out and em is the most common NIC in use today.
- Default config.xml has been cleaned up. Outdated comments have been removed that used to loosely document the config file, but had been neglected for quite some time and aren't all that useful anyway.
- Default sysctls have moved out of config.xml and now reside in globals.inc to reduce the size of config.xml
- Default sysctl values do not need to be set in config.xml. The default values are obtained from sysctl now. Also to reduce config.xml size.
- Tracking IDs added to default rules

3.53.4 Security Enhancements

- Verify SSL certificates for HTTPS URLs
- Detect if an unofficial package repository is in use and warn the user. Warning is displayed on the dashboard and package management pages. #484
- Check and verify the package server's SSL certificate if using HTTPS. #484
- For dyndns providers that support HTTPS, use it when performing updates.
- Replaced lots of GET actions with POST actions in various places in the GUI as they were touched.
- Update jquery to 1.11.1
- Remove almost all calls to history.back() and make Cancel button back to HTTP_REFERER
- Hide FreeBSD version from sshd banner. #3840
- Disable SSLv3 in lighttpd
- Disable RC4 ciphers in lighttpd
- Teach the certificate generation code how to make a self-signed certificate, and change the GUI cert generation code to use it.

  Also, move the GUI cert generation code to its own function so we can add a GUI option to regenerate it later. Also use some more sane defaults for the contents of the default self-signed certificate's fields so it will be more unique and less likely to trigger problems in browser certificate storage handling.
- Add command line script to generate and activate a new GUI certificate (generateguicert)
- Catch some more sensitive information when sanitizing the contents of config.xml output on /status.php.

3.53.5 OS Changes

- Updated base OS to FreeBSD 10.1-RELEASE
- PHP backend switched from FastCGI to PHP-FPM
- PHP Moved to 5.5
- Migrate captive portal code to SQLite3 PHP module
- Fix some lingering call-time pass-by-reference instances that fail on PHP 5.5
- Default serial speed is now 115200 #3715
- Sync gettytab and etc/ttys with FreeBSD 10-STABLE and reduce customizations
- Log pfSense version to syslog after bootup
- Set the sysctl net.inet.icmp.reply_from_interface to 1 to use the incoming interface to send ICMP replies. #3666
- Switched the hash method in pf to XXHASH for speed improvements

3.53.6 DNS

- Imported Unbound for use as the default DNS Resolver. The old dnsmasq DNS Forwarder is available as a non-default option. Upgraded systems will retain existing settings.
- Various changes to Unbound and supporting programs to complete its integration.
- Removal of bind from FreeBSD base necessitated the switch to alternate programs for DNS utilities (e.g. drill for dig, different nsupdate)
- AJAX DNS updates for firewall logs (when clicked)
- Make sure that the DNS Forwarder/Resolver is always capable of accepting queries on localhost before using it as a DNS server.
- If localhost is configured to be included in resolv.conf, force its selection in Unbound. The resolv.conf logic prevents that from being a problem, but users don't seem to realize they have to pick that to use Unbound for the host itself.
- IPv6 support in Unbound
- Check port of dnsmasq/unbound and skip 127.0.0.1 in resolv.conf if not port 53. #4022
- Add a note to the wizard about the DNS Resolver ignoring manual name servers by default. (They are still used as secondary/tertiary servers for the firewall itself, however)
- Domain and search should not both be defined in resolv.conf per FreeBSD man page and handbook (only the latter is actually used). Only search is set now.

3.53.7 CARP

- Changes to CARP for new FreeBSD 10 CARP system
- Provide a way to 'permanently' set CARP to 'maintenance mode' (advskew 254) persisting through a reboot.
- Key off net.inet.carp.demotion and display a warning to the user if the system has self-demoted its CARP status.
- Allow CARP IP address to be outside interface and alias subnets

3.53.8 Interfaces

- Implement an option to allow using the IPv4 connectivity interface for sending the dhcpv6 information. Usually useful for PPP[oE] type links and some ISPs
- Add gre and gif checks for IPv4 function interface_has_gateway($friendly), like they are already for IPv6
- Do not allow the user to set IPs for GRE interfaces on interface edit page. #3575
- On interfaces_assign.php, let user select network port to add instead of picking the first available #3846
- When changing an existing VIP, use previous configured interface for checking, this fixes the issue that happens when trying to change a VIP to a new interface. #3807
- Validate the GIF interface MTU (must be something between 1280 and 8192) #3927
- Properly set MTU for lagg(4) interface #3922
- Fix formatting of the Interfaces Widget on the Dashboard. #3937
- Don't allow interface descriptions that are strictly numbers as that generates an invalid ruleset. #4005
- Disable delete_old_states in dhclient-script. rc.newwanip handles this correctly in 2.2, and this killed states in multiple circumstances where that isn't necessary nor desirable.
- Do not unset configuration values from PPP config if not needed. #3727
- Overhaul handling of flags for hardware offloading and make it work correctly for system_advanced page settings. Lagg is still a special case that may require a reboot initially to apply. #1047
- Don't try to launch 3gstats unless it's on a valid device.
- Updated list of mobile service providers

3.53.9 Gateways/Routing

- Add an option to force a gateway to be down. #2847
- List GWGs in Interface to send DynDNS update from
- Allow reordering, batch delete, and disable of static routes
- Option to disable a gateway added
- Check gateway for IPv6 also for reply-to rules.
- Fix issue where ICMP6 messages sometimes have the wrong source IP address when a monitor IP address has been set #3607
- Improve look of gateways widget
- Provide a toggle for apinger debug messages to be logged to syslog
- Setting an interface IP to 0.0.0.0 with mask 0.0.0.0 overwrites the default route with that interface's link route. Follow FreeBSD 10.1 and use a /8 mask instead. #3941
- Use static route with -iface option for PPPoE to help when more than one PPPoE connection has the same gateway. #4040
- net.inet6.ip6.rfc6204w3 needs to be 1 for dhcpv6 to work correctly. #3361
- Add a route debug option to log info about route commands executed (where those aren't already logged) to help with troubleshooting various routing scenarios.
- Make sure srcip and target have scope when link-local addresses are used in apinger. #3969

- Properly generate and use the default gw for 6rd.

3.53.10 Firewall Rules

- Custom logging daemon that provides easy-to-parse output on a single line
- Persistent tracking ID for firewall rules so that logs may always be traced back to their corresponding rules
- Removed settings for maximum tables and maximum table entries since pf on FreeBSD 10 does not have any limits for these.
- Expose all p0f OS types that it supports so that subtypes of various Operating Systems can be detected (e.g. blocking Windows XP)
- The "(self)" concept of "Any IP address on this firewall" is now a choice for firewall rule destination (and floating rule source for out direction rules), port forward destination, and outbound NAT source.
- Can now optionally log default pass rules as well as default block rules
- Add IP alias subnets to interface subnet macro on GUI. #983
- Adjust states summary for new pfctl -ss output. #2121
- Add a more obvious note on group rules about how they do not work as expected for WANs
- Block IPv4 link-local/APIPA 169.254.0.0/16. #2073

Note: Per RFC 3927, hosts "MUST NOT send the packet to any router for forwarding", and "any network device receiving such a packet MUST NOT forward it". FreeBSD won't route it (route-to can override in some circumstances), so it can't be in use as a real network anywhere with the possible exception of local-only networks. Unlikely any such situation exists anywhere

---

- Fix JavaScript confirmation dialog for EasyRule.
- Use 'clog -f /var/log/filter.log' to view firewall log entries from the console so they are displayed in the new format.
- Set MSS clamping on VPNs in both directions rather than requiring it be set on both ends.
- Add option to kill all states on IP change, currently a hidden option for more testing. #1629
- Kill states associated with the old WAN IP when WAN IP has changed. #1629

3.53.11 NAT

- Hybrid outbound NAT style that allows the user to keep the existing automatic behavior but layer manual rules on top of it.
- Option to disable outbound NAT without disabling pf
- Display networks used in automatic outbound NAT when using that mode
- Allow reordering, batch delete, and disable of 1:1 NAT rules
- Take virtual IPs into consideration for automatic outbound NAT rules #983
- Outbound NAT can apply to any type of interface, make WAN-type specific reference generic

3.53.12 Aliases

- Allow individual line descriptions on alias bulk import
- Implement URL Table aliases for ports
- Optimizations for URL table aliases to use less memory and be more robust in general
- Alias name cannot have more than 31 chars, add maxlength to the field as an extra check. #3827
- Prevent Internal Server Error if an IP range is entered backwards.
- Expand range or subnet entered into a host type alias.
- Warn that IPv6 address ranges are not supported in aliases.
- When an alias contain hosts, add IPs and networks to filterdns too, otherwise the ruleset ends up with a predefined and non-persistent table. #3939

3.53.13 Dashboard & General GUI

- Various fixes for XHTML compliance
- Various fixes for typos
- Add a setting to allow the user to specify the clog file size so more (or less) entries may be kept in the raw logs.
- Add an option for users to be able to adjust how many configuration revisions are kept in the local backup cache.
- Show backup file size in config history.
- Display pfSense interface name on status interfaces
- Dashboard cleanups/fixes for jQuery
- Add “pfsense_ng_fs” full screen/widescreen theme
- GUI redirect works on both IPv4 and IPv6 #3437
- Disk usage section of the System Information widget now shows all UFS, ZFS, and cd9660 filesystems, not just the root (/) slice, and also indicates if they are a RAM disk.
- Add a message about premium content to the setup wizard and add a link in the menu to the signup page.
- Add pages missing from the Status > Traffic Graph privilege that are required for the full page to load
- Fix traffic graph widget default autoscale
- Be more strict on user and group removal to avoid removing accidentally removing additional users #3856
- Add an option to restart php-fpm from console
- Add .inc file for gmirror status widget to give it a better title and link to the management page.
- Allow the Virtual IP list table to be sorted (cosmetic only)

3.53.14 Translations

- Change default charset on pages to utf-8
- Updates to pt_BR translation
- Added Japanese translation
- Added Turkish translation
- Fixes for gettext

3.53.15 Captive Portal

- Add a way to download CP portal, error and logout html pages. #3339
- Add an option to restore default logout/error/portal custom pages on Captive Portal. #3362
- For more than 100 MAC pass-through entries create pipes in line with the rules file to speedup the process. #3932
- Zone backend changed from text-based (e.g. "cpzone") to using the zone id (e.g. "2") for specifying the context.
- ipfw_context has been removed. To list zones, use "ipfw zone list"
- Default lighttpd daemon port for a Captive Portal zone is based on the zone ID. For example, zone ID 2 uses port 8002. There may not be a daemon on port 8000.

3.53.16 IPsec

- IPsec backend changed from racoon to strongSwan
- IKEv2 settings have been enabled in the GUI
- Default IPsec configuration settings for newly created site to site configurations updated to use main mode and AES 256 on both phase 1 and 2.
- IPsec status page and dashboard widget changes to accommodate different output from strongSwan
- Move the IPsec settings from System > Advanced, Misc tab to "Advanced Settings" tab under VPN > IPsec.
- It is now possible to configure *L2TP/IPsec*
- Add AES-GCM and AES-XCBC to the list of available IPsec algorithms and hashes, respectively. Expand P1 DH groups up to 24.
- Allow hash algorithms to be empty for phase 2 where the encryption is AES-GCM
- Allow to reorder IPsec Phase 1 and Phase 2 items, remove multiple P1/P2 items, toggle enable/disable status of P1/P2 items #3328
- Provide a first implementation of EAP-TLS authentication with IKEv2. It is a start and might not work on all cases
- Do not accept non-ASCII characters on IPsec PSK #3931
- Fix ping_hosts.sh to not ping IPsec if CARP is in backup.
- Allow accept_unencrypted_mainmode_messages to be enabled for IPsec if needed.
- Check that subnet masks are equal when choosing binat type for IPsec to avoid errors on ruleset. #3198
- Change NAT Traversal options as strongSwan only has two options: force or auto.
- Don't allow P2 local+remote network combinations that overlap with interface+remote-gateway of the P1. #3812

3.53.17 OpenVPN

- Allow entering OpenVPN client credentials in the GUI
- Add fields for local (push route) and remote (iroute) network definitions in an OpenVPN client-specific override entry.
- Change OpenVPN compression settings to cover the full range of allowed settings on OpenVPN (unset, off, on, adaptive) rather than a simple off/on switch that either doesn't set the value or enables it with adaptive (OpenVPN's default).
- Add an Authentication Digest Algorithm drop-down to OpenVPN server/client and to the wizard (SHA1 is the default since that is OpenVPN's default)
- Add option to specify client management port for OpenVPN client export use
- Ensure e-mail address carries over from the CA screen to the Cert screen in the OpenVPN wizard.
- Allow the user to select "None" for OpenVPN client certificate, so long as they supply an auth user/pass. #3633
- Byte counts on OpenVPN status are now human readable rather than huge unformatted numbers.
- OpenVPN instances have new options: "Disable IPv6", route-nopull, route-noexec, verb selector
- Use stronger defaults in the OpenVPN wizard.
- Fix ovpn-linkup for tun + topology subnet case setting router as ifconfig_local envvar when route_vpn_gateway and ifconfig_remote are both not defined. #3968

3.53.18 DHCP

- Add code for UEFI booting and DHCP
- Advanced RFC 2136 configuration for DHCPd service
- Add ability to not supply a DHCP gateway to clients
- Allow defining DHCP static mappings using dhcp-client-identifier
- Do not call write_config() when Applying Changes on DHCP settings #3797

3.53.19 Packages

- Package signing to ensure validity/authenticity
- Single package manifest (XML) file rather than one per architecture
- Various improvements to PBI setup/structure from upstream (PC BSD)
- Added the capability for package hooks in /etc/rc.carpmaster and /etc/rc.carpbackup
- Split package category display into separate tabs for categories, and provide an "All" tab
- Move the fetching of a package's config file and additional files to separate functions.
- Clarify logs generated by newwanip(v6) when restarting packages, it's not only IP changes that end up here (by design).
- When reinstalling a package, try to start it after the install completes.

3.53.20 Dynamic DNS

- Added support for DynDNS Provider "City Network"
- Added support for DynDNS Provider "OVH DynHOST"
- Added support for DynDNS Provider "GratisDNS"
- Added support for DynDNS Provider "Euro DNS"
- Added support for DynDNS Provider "CloudFlare"
- Add support for custom IPv6 DDNS.
- Add backend support for HE.net AAAA record updates.
- Add additional options to Custom DynDNS
- Allow hostname to start with '@.' for namecheap #3568
- Do not disable certificate verification in DynDNS. Proper CA certificates are now in place to validate SSL in these cases.
- "+" is a valid character in some dynamic DNS providers' usernames. #3912

3.53.21 GEOM Mirrors (gmirror)

- New gmirror library to perform various gmirror tasks and get information, using some of the former widget logic to start.
- Added a Diag > GEOM Mirrors page that displays information about existing mirrors and performs various management tasks.
- Also included is a notification setup. Mirror status is polled every 60 seconds, and if any aspect of the mirror changes, notifications are issued that alert in the GUI and by SMTP, etc.

Warning: If a manual gmirror configuration was performed post-install and not using the pfSense software installer gmirror option before install, there is a chance that the mirror will not function on pfSense software version 2.2 because the manual post-install method did not create a completely proper mirror setup. If the upgraded mirror does not function on 2.2, the following `/boot/loader.conf.local` entry may be used to work around the integrity check that would otherwise fail:

```
kern.geom.part.check_integrity=0
```

If one of these configurations is present, the best practice is backing up the configuration and reinstalling using the built-in gmirror option in the pfSense software installer.

3.53.22 Traffic Shaping

- Fix DSCP values and provide a config upgrade to fix values stored in config.xml. #3688
- Remove 'multi lan/single wan' and 'multi wan/single lan' traffic shaper wizards, multi lan/wan can be used to replace any of them.
- Only show the correct type of interfaces (LAN/WAN) on traffic shaper wizards #3535
- Shaper wizard will automatically attempt to guess the correct number of WANs and LANs.
- Updated and expanded traffic shaping for games, game consoles, and other applications.

- Allow up to 2900 limiters. This was set to 30. #3213
- Fix logic to find available next number for limiters and queues. #3998
- Add vmx and hn to list of ALTQ capable interfaces.
- Remove the "Limiter burst" parameter as it currently doesn't work with dummynet in pf.

3.53.23 Misc

- Cleaned up various older files/scripts that were no longer being used
- Dropped all support for cvsup. cvs is dead, long live svn and git.
- Optimizations/changes to the XML Parsing code
- NTP updates to handle a wider ranges of GPS devices and more NTP options
- Move to zerocopy_enable for bpf to optimize logging which uses bpf interface. This should increase the general performance since pflog is always enabled.
- Add sshd service to list (if enabled)
- Add a "status" subcommand to the svc php shell script.
- When using the reset webConfigurator password option on the console, if authentication server is not Local Database, ask user if they want to revert back to it. #3341
- Fix interface selections on UPnP to show the customized descriptions entered by the user. While here, add an external interface selection knob. Fixes #3141
- Layer 7 Pattern: EAOrigin.pat
- Layer 7 Pattern: SWF (Flash)
- Remove some old obsolete code that referred to the now-defunct "embedded" platform that was replaced with NanoBSD back in 1.2.x.
- Sometimes fsck requires a second run, teach rc script to call it more than once when it's necessary
- Add column for internal port on UPnP status page
- Make listening on interface rather than IP optional for UPnP
- Use interface name for miniupnpd rather than IPv4 address #3874
- Packet Capture: Host field supports rudimentary boolean logic.
- Packet Capture: Protocol, host, and port now support negation.
- Added interface column to Diagnostics > States
- Change is_port() to only validate a single port, is_portrange() for specific cases. #3857
- Fix guess_interface_from_ip() to account for differences in netstat output. #3853
- Fix Certificate Authority SAN name handling #3347
- Add a basic command line password reset script.
- Use configured proxy URL/port for downloading bogon list. Does not use credentials. #3789
- Underscores are valid characters in domains. #3219
- Let user decide to proceed with upgrade when sha256 fails to download. #3576
- Remove the command number shown in the shell prompt.

- Use a better method of finding disks for SMART.
- Process obsolete files in shell script instead of PHP.
- Do not allow FQDN in fields that should only accept a hostname.
- Set proxy environment variables on interactive shell and also on crontab so that they may be used by all scripts. #3789
- Add input checkboxes to remove multiple users and groups
- Make sure an empty group or user is not created when editing
- Update URLs in help.php.
- Change wording at the end of the wizard to remove "donate" since that is no longer an option.
- Put the booting signal in globals.inc since it makes all the other scripts detect we are booting. Otherwise separate PHP instances will not detect that. rc.bootup clears this flag so all should work correctly
- Force serial console when it was selected by the installer. #4009
- Wait 10 minutes before retrying bogon fetch on soft failures to avoid us getting DoSed if something is wrong there (like someone's system can't validate the cert)
- Use IPv4 for ntpq if IPv6 is not allowed

3.53.24 HEADS UP for Xen Users

The FreeBSD 10.1 base used by pfSense 2.2 includes PVHVM drivers for Xen in the kernel. This can cause Xen to automatically change the disk and network device names during an upgrade to pfSense 2.2, which the hypervisor should not do but does anyway.

The disk change can be worked around by running /usr/local/sbin/ufslabels.sh *before* the upgrade to convert the fstab to UFS labels rather than disk device names.

The NIC device change issue has no workaround. Manual reassignment is required at this time. Note there have been performance issues reported in Xen with this NIC device change.

3.54 2.1.5 New Features and Changes

The pfSense® software version 2.1.5 release follows shortly after *2.1.4* and is primarily a security release.

3.54.1 Security Fixes

- pfSense-SA-14_14.openssl
  - See http://www.openssl.org/news/secadv_20140806.txt
  - Updated to OpenSSL 0.9.8zb and 1.0.1i
- pfSense-SA-14_15.webgui
- pfSense-SA-14_16.webgui
- pfSense-SA-14_17.webgui

3.54.2 Other Fixes

- Handle a missing DHCPD config section properly during a configuration upgrade
- Fix a regression that broke CARP+IP alias VIP functionality
- Fix the Pass, Block, Reject and Interface filters in the Firewall Logs Widget #3725
- Use HTTPS for dyndns providers that support it
- Avoid resetting the firewall hostname from a WAN DHCP server #3746
- Add missing qlimit keyword in some shaper rules
- Change Cancel button to call history.back() when editing firewall aliases to fix issues with IE 11 #3728
- Allow hostnames in bulk import since they are valid entries in a network type alias
- Fix input validation logic on diag_testport.php, escape more shell arguments for good measure
- Escape the individual dnsmasq advanced/custom options
- Encode the detail field of an alias entry before displaying its contents back to the user
- Encode interface/VIP descriptions before displaying them on the NTP daemon settings, and GIF/GRE interfaces
- Per the dhcpd.conf man page and other documentation from ISC, mclt must not be defined on the secondary
- Shorten the wait at "reload" in startup wizard to 5 seconds from 60
- Do not execute DNS lookups on GET, only pre-fill Host box so the user can press the button to execute
- Turn alias creation links from DNS lookups into submit buttons for POST
- Remove javascript alert DNS resolution action from the firewall log view. It was already removed from 2.2, and it's better not to allow a GET action to perform that action
- Require click-through POST confirmation when restoring or deleting a configuration from the backup history page
- Avoid a "Cannot use string offset as an array" error if the packages section of the config is missing
- Avoid generating an invalid IPsec (racoon) config if the user specified a mobile pool that is too small
- IPsec phase 2 pinghost was not used if the source IP was a virtual IP address #3798
- Move dhcp6c log to dhcpd.log #3799
- Do not reset source and destination port range values when it's an associated rule created by NAT port forward. #3778
- Added filter.so to list of extensions loaded for filter_var() support.
- The pfSense PHP module was setting the subnet mask of lo0 to /0, which could break some routes and cause other unintended routing side effects.

3.55 2.1.4 New Features and Changes

pfSense® software version 2.1.4 follows very shortly after *2.1.3* and is primarily a security release. Refer to the *2.1.1 release notes*, *2.1.2 release notes*, and *2.1.3 release notes* for other recent changes.

3.55.1 Security Fixes

- pfSense-SA-14_07.openssl
  - FreeBSD-SA-14:14.openssl
- pfSense-SA-14_08.webgui
- pfSense-SA-14_09.webgui
- pfSense-SA-14_10.webgui
- pfSense-SA-14_11.webgui
- pfSense-SA-14_12.webgui
- pfSense-SA-14_13.packages

Packages also had their own independent fixes and need updating. During the firmware update process the packages will be reinstalled properly. Otherwise, uninstall and then reinstall packages to ensure that the latest version of the binaries is in use.

3.55.2 Other Fixes

- Patch for Captive Portal pipeno leaking issue which leads to the 'Maximum login reached' on Captive Portal. #3062
- Remove text not relevant to Allowed IPs on the Captive Portal. #3594
- Remove units from burst as it is always specified in bytes. (Per ipfw(8)).
- Add column for internal port on UPnP status page.
- Make listening on interface rather than IP optional for UPnP.
- Fix highlighting of selected rules. #3646
- Add guiconfig to widgets not including it. #3498
- /etc/version_kernel and /etc/version_base no longer exist, use php_uname to get the version for XMLRPC check instead.
- Fix variable typo. #3669
- Delete all IP Aliases when an interface is disabled. #3650
- Properly handle RRD archive rename during upgrade and squelch errors if it fails.
- Convert protocol ssl:// to https:// when creating HTTP headers for XMLRPC.
- Show disabled interfaces when they were already part of an interface group. This avoids showing a random interface instead and letting the user add it by mistake. #3680
- The client-config-dir directive for OpenVPN is also useful when using OpenVPN's internal DHCP while bridging, so add it in that case also.
- Use curl instead of fetch to download update files. #3691

- Escape variable before passing to shell from stop_service().
- Add some protection to parameters that come through _GET in service management.
- Escape argument on call to is_process_running, also remove some unnecessary mwexec() calls.
- Do not allow interface group name to be bigger than 15 chars. #3208
- Be more precise to match members of a bridge interface, it should fix #3637
- Do not expire already disabled users, it fixes #3644
- Validate starttime and stoptime format on firewall_schedule_edit.php
- Be more careful with host parameter on diag_dns.php and make sure it's escaped when call shell functions
- Escape parameters passed to shell_exec() in diag_smart.php and elsewhere
- Make sure variables are escaped/sanitized on status_rrd_graph_img.php
- Replace exec calls to run rm by unlink_if_exists() on status_rrd_graph_img.php
- Replace all `hostname` calls by php_uname('n') on status_rrd_graph_img.php
- Replace all `date` calls by strftime() on status_rrd_graph_img.php
- Add $_gb to collect possibly garbage from exec return on status_rrd_graph_img.php
- Avoid directory traversal in pkg_edit.php when reading package xml files, also check if file exists before try to read it
- Remove id=0 from miniupnpd menu and shortcut
- Remove . and / from pkg name to avoid directory traversal in pkg_mgr_install.php
- Fix core dump on viewing invalid package log
- Avoid directory traversal on system_firmware_restorefullbackup.php
- Re-generate session ID on a successful login to avoid session fixation
- Protect rssfeed parameters with htmlspecialchars() in rss.widget.php
- Protect servicestatusfilter parameter with htmlspecialchars() in services_status.widget.php
- Always set httponly attribute on cookies
- Set 'Disable webConfigurator login autocomplete' as on by default for new installs
- Simplify logic, add some protection to user input parameters on log.widget.php
- Make sure single quotes are encoded and avoid javascript injection on exec.php
- Add missing NAT protocols on firewall_nat_edit.php
- Remove extra data after space in DSCP and fix pf rule syntax. #3688
- Only include a scheduled rule if it is strictly before the end time. #3558

3.56 2.1.3 New Features and Changes

pfSense® software version 2.1.3 follows very shortly after *2.1.2* and is primarily a security release. Refer to the *2.1.1 release notes* for changes from 2.1 to 2.1.1 and *2.1.2 release notes* for changes from 2.1.1 to 2.1.2.

3.56.1 Security Fixes

- pfSense-SA-14_05.tcp
  - FreeBSD-SA-14:08.tcp
- pfSense-SA-14_06.openssl
  - FreeBSD-SA-14:09.openssl

Packages also had their own independent fixes and need updating. During the firmware update process the packages will be reinstalled properly. Otherwise, uninstall and then reinstall packages to ensure that the latest version of the binaries is in use.

Although these security issues warrant updating as soon as possible, they are of relatively minor impact to the average user. According to the FreeBSD SA, the TCP flaw is mitigated by scrub in pf which is enabled by default in pfSense. The OpenSSL flaw is not used by any daemons in the pfSense base system and only certain packages make use of the affected feature, so the impact there is also minimal.

3.56.2 Other Fixes

- Various fixes to accommodate recent changes/optimizations in the tools repository
- Move clog binary to its proper place in /usr/local/ to respect hier(7)
- Fix remove button on Diagnostics > Tables #3627
- Fix more potential places for interface looping in OpenVPN and with normal interfaces
- Fixes for URL table alias updates (locking, reload)
- Fix IPsec Phase 1 duplication
- Fix 'add rule on top of the list' allowing after param to be -1
- Correct Captive Portal redirection URL to unbreak ones passed through Radius attributes and respect user choices.
- Make miniupnpd listen on interface instead of IP
- Don't refuse to delete a bridge in the GUI just because its bridge interface doesn't exist, just log that it doesn't exist and don't attempt to ifconfig destroy it, delete it from config
- Fixes for DynDNS to allow configurable check host.
- Resolver has no option for remote syslog, remove wrong copy/paste that was adding it when apinger was enabled
- Fix typo for GIF tunnels to work over IPv6
- Fix for dhcrelay target using default GW
- List Gateway Groups in Interface to send update from for custom DynDNS entries

3.57 2.1.2 New Features and Changes

pfSense® software version 2.1.2 follows very shortly after *2.1.1* and is primarily a security release. Refer to the *2.1.1 release notes* for changes from 2.1 to 2.1.1.

3.57.1 Security Fixes

The Heartbleed OpenSSL bug and another OpenSSL bug were both covered by the following security announcements:

- pfSense-SA-14_04.openssl
  - FreeBSD-SA-14:06.openssl
  - CVE-2014-0160 (Heartbleed)
  - CVE-2014-0076 (ECDSA Flaw)

Packages also had their own independent fixes and need updating. During the firmware update process the packages will be reinstalled properly. Otherwise, uninstall and then reinstall packages to ensure that the latest version of the binaries is in use.

3.57.2 Other Fixes

- On packages that use row_helper, when user clicks on add or delete button the page scrolls to top. #3569
- Correct typo on function name in Captive Portal bandwidth allocation
- Make extra sure that the firewall does not start multiple instances of dhcpleases if, for example, the PID is stale/invalid and there is still a running instance.
- Fix CRL editing. Use an alphanumeric test rather than purely is_numericint because the ID is generated by uniqid and is not purely numeric. #3591

3.58 2.1.1 New Features and Changes

3.58.1 Security Fixes

- FreeBSD-SA-14:01.bsnmpd / CVE-2014-1452
- FreeBSD-SA-14:02.ntpd / CVE-2013-5211
- FreeBSD-SA-14:03.openssl / CVE-2013-4353, CVE-2013-6449, CVE-2013-6450
- Use HTTPS to get updates. #2952
- Escape necessary chars to avoid XSS injection. #2952
- Add escapeshellarg() calls on more exec parameters.
- Replace some exec() calls by php functions like symlink, copy, unlink, etc.
- Use HTTPS for pfsense.org URLs.
- Protect output to browser by using htmlspecialchars. #3461
- Improve checks for params 'id', 'dup' and other similar ones to make sure they are numeric integer, also, pass them through htmlspecialchars before printing.

- Remove special characters that can lead to shell/XSS compromises from submitted input when installing packages. #3461
- Ask for validation when real package operation will be done and ask for the operation with POST to get protection from CSRF. #3460
- Use HTTPS for fetching packages.

3.58.2 Interfaces

- Updated em/igb/ixgb/ixgbe drivers that add support for i210 and i354 NICs and fix issues with `ix(4)` cards.
- Prevent assigned vlans from having their tag changed.
- Fix ifconfig error on gif in certain cases.
- If rc.newwanip is run on an interface that should not have an IP address, do not take any action. This could lead to certain interfaces bouncing link if they had no IP address.
- In rc.newwanip, if the interface is configured and not enabled, bail. The firewall does not need to change settings for disabled interfaces. #3313
- Skip processing in rc.newwanip if the interface has no IP address.
- Fix pkg_edit.php to show interface description instead of interface name
- Make sure vlan interface exist when they are configured #3270
- Limit CIDR choices for IPv4 on GRE interface. #3277
- Do not destroy an interface when it's being disabled #3350
- Prevent network or broadcast address to be set on interface (console, GUI and wizard). #3196
- Reduce unnecessary operations and other fixes to MTU code. This fixes slow boot times and proper handling of mtu for VLANs.
- Provide a dynamic gateway for GIF and GRE v6 tunnels so it can be used on firewall rules etc. #3484
- Bring up appropriate interface for GRE/GIF. #3281
- Prevent removing the IP from the underlying GRE interface in the OS when assigning GRE interface and configuring an IP address. #3280
- When an interface goes down try to shut the RAs and dhcpd6 service on that interface. #2627
- Sync up ALTQ-capable interfaces list
- Trigger rc.newwaipv6 from pppoe when it gets an inet6 configuration
- Update list of mobile service providers.
- Correct check to enable ieee8021x.

3.58.3 Gateways/Routing

- Respect default gateway option when adding a gateway from interfaces page. #3230
- Use a more accurate error message when attempting to add/edit a gateway that does not have an appropriate IP address for the type. #3282
- Make return_gateways_array() return all disabled gateways when $disabled is true. #3291
- Don't flush interface cache on each call of the function when looping through all gateways.
- Fix an issue that changes wrong gateway entry when items are hidden
- Delete static route when monitor IP is removed, also save monitor IP even when it's disabled
- Return Gateway Group IP protocol version even when no gateway IP can be located.
- Remove broken 'dynamic6' gateway, we already have ipprotocol to tell us the IP version, leave it more simple using only 'dynamic'

3.58.4 NAT/Firewall Rules/Aliases

- Reload filter rules when activate or deactivate dhcpdv6 #3218
- Make sure no extra spaces end up in the parsed IP in the filter logs as it can lead to issues in other places (Easy Rule, etc)
- Use (self) rather than any as the destination for the lockout rules
- Use (self) instead of any for web lockout
- Avoid pf table names conflict. #3268
- Fix display of full URL in URL table listing as seen in an Alias popup. #3242
- Make it more explicit that 'update freq.' for URL table aliases unit is days
- Fix situation where removing an alias entry and then adding a new one resulted in an entry box with broken formatting. #3283
- Make sure pf rule labels never have more than 63 chars. #3208
- Rewrite the display_host_results() function to use spaces instead of tabs. It does a much better job of aligning the fields in each column and works in all the browsers, particularly chrome which doesn't support the tab character.
- Handle comma-separated list of remote networks when making vpn_networks table
- Fix rules that pass out traffic for Proxy ARP VIP entries which had incorrect destination #3331
- Load only the options rather than clearing the whole ruleset.
- Validate IP address ranges correctly on Alias Bulk Import
- Fix display of CIDR/Update Freq in Alias Edit
- In the filter log, the protocol might also say "icmpv6" so account for that when making a rule using Easy Rule.
- Move 'allow dhcpv6 client' rules above block bogonsv6 ones. #3395
- Only add dhcpv6 client allow rules if ipv6allow is set
- Add all advanced options to rule table hover text.
- Open up Firewall Rules Advanced Options section if any values have been set.
- Validate rule Advanced Options numeric entries properly

- Disable default allow incoming rules for 6to4 and 6rd interfaces. This rule unintentionally allows all services on the interface.
- Skip OpenVPN interfaces when creating the first set of manual rules to be consistent with the behavior of Automatic Outbound NAT. #3528
- Try to restore last working ruleset rather than staying without configuration at all if an invalid ruleset is encountered.
- Fix days and weeks selection on schedules
- Prevent putting an subnet in the IPv6 address field since it breaks the filter generation process.
- Put a timeout of 30 seconds on the bogon update download. #3412
- Before downloading file to process urltable, there is a random wait time between 5 and 60 seconds. Because of this, the difference between file mtime and current time can be less than $freq * 86400 and it'll be skipped. Add 90 seconds (60 of max random wait + 30 to be sure) to avoid skipping a file that should be updated. #3469
- Validate if src OR dst have IP address set when protocol is IPv4+v6. #3499
- Improve data validation to avoid save a host/subnet or a IPv4 with invalid mask. The reported error is on javascript and only happen on IE8, but this fix will prevent the same issue happening in the future on a different browser. #3449

3.58.5 Traffic Shaping

- Fixed typo in CoDel wiki link
- Fix codel not being applied on non-priq queue types
- Fix saving and range checking of 'Packet loss rate' and 'Bucket Size' in limiters.
- Add previously missing DSCP VA.
- Clarify note on limiter queue weight to state that higher values get a larger share.

3.58.6 Dashboard & General GUI

- Convert mac address to lowercase when saving to avoid duplicates. It fixes #3195
- Include the CP zone in the form parameters if one is defined. Fixes access to concurrent graph on zones other than the first/default.
- Miscellaneous HTML cleanup
- Fix interface names shown in the traffic graphs widget. #3245
- Send the help links to HTTPS destinations on web servers that support HTTPS.
- Specify favicon in pages directly
- Add some missing privileges to the list. #3279
- Many fixes on privileges. #3216
- Allow setting a default scale type preference for the traffic graphs widget
- Account for a widget being null/not defined, and not just closed/open when deciding if a widget function should be called. This allows the system information dashboard widgets to update properly.
- Avoid dashboard divide by zero errors
- Detect Zones and Cores for thermal sensors using regex. #3337

- Do not sort users when adding privileges. It's unnecessary and lead to unintentional edits to the wrong account.
- Add specific privilege for easyrule.
- Return all stats when all or remote is selected on Traffic Graph and make the default query return "Local" traffic.
- Update year, links for 2.1.1.

3.58.7 Captive Portal

- Fix CP stats generation for concurrent users. #3225
- Remove redundant copies of getNasIP() #3234
- Set default captive portal RADIUS authentication value to radius_protocol during upgrade #3226
- Add Captive Portal Zones privileges definition. #3216
- Prevent a possible division by zero in Captive Portal. #3212
- Fix saving of voucher sync settings
- Reduce the total minutes by the remote minutes used, do not use the value directly. Otherwise the voucher will be cut short or listed invalid when it otherwise should have time left over.
- Make sure to give the Captive Portal zone a name during the upgrade, or else it comes through with a blank/null name.
- Properly set zone dedicated rules in the rules/pipes DBs to properly release when a zone is deactivated
- Don't generate rules for disabled captive portal instances
- Do some more error checking and put secondary radius attributes only if configured on a Captive Portal instance.
- If set use the default bandwidth setting on the Captive Portal even for MAC passthrough.
- Fix various problems with Captive Portal voucher synchronization introduced during conversion to zones.
- Properly compile the Captive Portal database query to insert the values.
- Fix deletion of IPFW rules and pipes for passthru MAC. #3538
- Use the 11th column for the radius context rather than overriding the interim interval field with it. #3447
- Use descr as the field name for voucher description so it gets CDATA protection. #3441
- Consider setting of noconcurrent login for passthrough expiration of users. #3340
- Use the default bandwidth specification if configured even for allowed IP address and hostname.
- Properly detect when there are issues with communicating with syncip and to use the local DB for this. Otherwise detect if the remote says the voucher is not valid say its not valid.

3.58.8 VPN

- Fix find_service_by_openvpn_vpnid() on OpenVPN Status
- Allow special characters to be used on IPsec mobile login banner. #3247
- Fix cisco-avpair processing for IPsec and OpenVPN, and route processing from avpair replies.
- Fix logic in detecting if OpenVPN resync needed
- Fix vpn_pppoe_get_id and stop duplicating pppoeid for multiple servers. #2286
- Use env var provided by openvpn to determine if it's tun or tap. #3475

- Add an option to verify IPsec peers_identifier when it's ASN.1 distinguished name. #2904

3.58.9 Certificates

- Certificate Manager, for 'Create an internal Certificate' use the correct 'Digest Algorithm'
- OpenSSL does not like country codes longer than two letters, so remove entries that are not actually country codes.
- Perform a much more accurate comparison between two certificates to determine if they are identical when checking their revocation status. #3237
- Allow an "empty" CRL to be exported, since this is still a valid action.
- Fixes for "Alternative Names" on certificates.
- Fix issue with CSR generation. #2820
- Increase default openssl to bits 2048.

3.58.10 DHCP

- Optimize DHCPv4 lease display online status for static leases. Do not re-parse complete ARP table for each lease, as it can be slow with large ARP tables.
- Add upgrade code to change the DHCP next-server value to nextserver since it was renamed sometime in 2.1 but upgrade code didn't follow.
- Give clients the IPV6 address of the DNS server via DHCPv6 Server
- Check if dhcp start and end addresses are inside interface subnet. #3196
- Remove 'deny unknown clients' option from DHCPv6 since it's not supported. #3364
- Fix DHCP lease time display, strftime already convert it to local timezone, so no need to calc offset
- Use correct parameter (bootfile-url) to configure netboot on DHCPdv6. #3421
- Only use IPv4 DNS servers in IPv4 DHCP configuration. #3483
- Fix PHP error when saving DHCP settings if no manually configured DNS servers exist.
- Send a HUP to dhcp6 to signal a reload. #3514

3.58.11 Load Balancing

- Prevent a Fall Back Pool from being selected when the DNS protocol is in use. If one is present in the config, ignore it. #3300
- Fix display of pools in the LB status widget and on the LB Virtual Server status.

3.58.12 Time

- Allow multiple valid time servers to be entered in the wizard, as they are allowed under System > General
- Update time zone data to 2013i
- Teach system_timezone_configure() to deal with symlinks to avoid having timezone misconfigured. #3293
- Add 'limited' to ntpd restrict list to workaround FreeBSD-SA-14:02.ntpd/CVE-2013-5211. #3384
- Use "disable monitor" in NTP config to mitigate FreeBSD-SA-14:02.ntpd/CVE-2013-5211.
- Update ntp to ntp-devel for FreeBSD-SA-14:02.ntpd/CVE-2013-5211.
- Avoid placing an empty "interface listen" directive in ntpd.conf.

3.58.13 Misc

- Fix ALIX upgrade crash during RRD processing
- Fix "Could not open shared memory for read 1000" issue on Diagnostics > NanoBSD. #3235
- Fix ufslabels.sh logic to avoid trying to convert slices which are already using appropriate labels. Fixes #3207
- Fix removal of the first cron job entry in the list.
- Remove unused newsyslog cron job from the default configuration and on upgrade.
- Split SSL/TLS into separate checkboxes so that plaintext connections can be made secured by using STARTTLS. Support for SMTPS connections should probably be done away with in future. #3180
- Add source address selection to syslog settings, so it can work more effectively over a VPN. #355
- Rework the usage of the shell i/o during stop_packages(), fixes the "Syntax error: bad fd number" for the remaining people who still saw it on shutdown
- Switch to rw mode before file operations on RFC2136 cache. Fixes #3201
- Make the RADIUS settings respect the description of the timeout field. If the timeout value is left blank, use 5 seconds, don't print an error.
- Call conf_mount_rw before deleting a user. #3294
- Handle the reinstallall case with confirmation. #3548
- Do not list the same CARP ip as an option for its own Interface.
- Accept adding an IP Aliases on top of CARP VIP when the parent interface does have a valid IP address in the alias subnet.
- Simplify log filtering logic calling grep less times, as done on mail_reports.inc on 2c6efc9.
- Fix console recent config restore, allow restoration of the last backup listed. #3438
- Enhanced validation of general DNS servers and gateways
- Add a mechanism by which the serial port can be forced on always regardless of the config setting. (useful for nano+vga setups)
- Add a knob to let the user select which console (video or serial) is preferred in cases where there are multiple consoles present.
- Skip input validation when choosing an existing certificate in the User Manager. #3505
- pfSense_interface_deladdress() only knows how to delete an ip address, not a subnet. #3513

- Make is_linklocal case-insensitive. #3433
- Errors in RRD graph calculations
- Delete /var/crash content when the user clicks 'No'. #3486
- Make sure filesystem is read-write when operating on groups. #3492
- Fix OpenVPN XML section name for selective configuration backup.
- Remove TRIM_set and TRIM_unset support. This method isn't very elegant and isn't necessary in the long run. It's better handled during the install process or while booted off other media (e.g. CD or Memstick).

3.59 2.1.0 New Features and Changes

3.59.1 Security Fixes

Three FreeBSD security advisories are applicable to prior pfSense® software releases. These aren't remotely exploitable in and of themselves, but anyone who can execute arbitrary code on the firewall could use one or more of these to escalate privileges.

- FreeBSD-SA-13:13.nullfs
- FreeBSD-SA-13:12.ifioctl.asc
- FreeBSD-SA-13:09.ip_multicast.asc

3.59.2 IPv6 Support

IPv6 Added to many areas of the GUI. At least the following areas/features are IPv6-enabled. Others may work as well

- Aliases (Firewall) - Aliases can contain both IPv4 and IPv6, only addresses relevant to a given rule will be used
- CARP RA
- CARP Failover
- DHCP Server w/Prefix Delegation
- SLAAC WAN
- 6to4 WAN
- 6to4 WAN w/Prefix Delegation
- 6rd WAN
- 6rd WAN w/Prefix Delegation
- DHCP6 WAN
- DHCP6 WAN w/Prefix Delegation
- DHCPv6 Relay
- DNS Forwarder
- Firewall Rules
- Gateway Groups/Multi-WAN - See *Configuring Multi-WAN for IPv6*
- Gateway Status (apinger)
- GIF Tunnels

- GRE Tunnels
- GUI Access
- IPsec
- L2TP
- NPt
- NTP
- OpenVPN
- Packet Capture
- PPPoE WAN
- Router Advertisements
- Routing
- Server LB
- Static IP
- Syslog (remote)
- Limiters (dummynet pipes)
- Virtual IPs - IP Alias
- Virtual IPs - CARP
- DNS from RA
- Accept RA when forwarding
- Auth via RADIUS
- Auth via LDAP
- XMLRPC Sync
- RRD Graphs
- DHCP Static Mapping - Works by DUID
- DynDNS (HE.net hosted DNS, RFC2136, custom)
- MAC OUI database lookup support for NDP and DHCPv6. (Was already present for DHCP leases and ARP table) requires the nmap package to be installed to activate

Note: Unlike earlier snapshots, BETA, etc, currently the upgrade does NOT flip the "Allow IPv6" checkbox on upgrade, to preserve existing behavior. To activate IPv6 traffic, a user will have to flip this setting manually

3.59.3 Packages

- PBI (push button installer) packages - all of a package's files and dependencies are kept in an isolated location so packages cannot interfere with one another in the way that was possible on 2.0.x and before using tbz packages
- RIP (routed) moved to a package
- OLSRD moved to a package
- Unbound moved back to a package (Will try integration again for 2.2)
- Increase the verboseness of the package reinstallation process in the system logs for a post-firmware-update package reinstallation operation

3.59.4 OS/Binary/Supporting Program Updates

- Based on FreeBSD 8.3
- Updated Atheros drivers
- OpenSSL 1.0.1e (or later) used by OpenVPN, PHP, IPsec, etc
- PHP to 5.3.x
- OpenVPN to 2.3.x
- Added mps kernel module
- Added ahci kernel module
- Updated ixgbe driver
- Many other supporting packages have been updated

3.59.5 Dashboard & General GUI

- Switch from Prototype to jQuery
- Improved navigation and service status in the GUI (shortcut icons in each section to quickly access config, logs, status, control services, etc)
- Multiple language support, a mostly-complete translation for Brazilian Portuguese is included
- Read-only privilege to create a user that cannot modify config.xml
- Dashboard update check can be disabled
- Fixed theme inconsistencies between the login form and other parts of the GUI
- Various fixes to pages to reduce potential exposure to certain CSRF/XSS vectors
- Updated CSRF Magic
- Set CSRF Magic token timeout to be the same as the login expiration
- Added IE Mobile for WP8 to list of browsers that get an alternate theme at login
- Truncate service status so long package descriptions cannot break formatting of the status table
- Many fixes to HTML/XHTML to improve rendering and validation
- Added a note to the setup wizard letting the user know that it can be canceled at any time by clicking the logo image
- Make dashboard update check respect nanobsd-vga #3078

- Firewall Logs Widget filtering and column changes
- Added totals for some dashboard widget meters (memory, swap, disk usage)
- Changed dashboard display for states and mbufs to be meters, and to show usage as a percentage
- Update dashboard mbuf count via AJAX
- Show a count and layout of CPUs in the dashboard if multiple CPUs are detected

3.59.6 Captive Portal

- Multi instance Captive Portal
- Multiple Captive Portal RADIUS authentication sources (e.g. one for users, one for cards)
- Logic fixes for voucher encryption
- Many optimizations to Captive Portal processing, including a database backend and moving functions to a php module to improve speed Optional Captive Portal user privilege
- Add checks to make sure CP hard timeout is less than or equal DHCP server default lease time, to avoid issues with CP sessions being valid for incorrect IPs, and users switching IPs while they should still be connected to the portal
- Fixes for captive portal voucher syncing on HTTPS with a custom port #3001
- Fixes for custom Captive Portal files leaving symlinks on the filesystem after files were removed
- Added MAC OUI database lookup support to CP status (requires nmap package to be installed)

3.59.7 OS/System Management

- Ability to select serial port speed
- Added a manual way to enable TRIM if someone needs it
- Added a manual way to trigger a fsck on reboot
- AES-NI support (Cryptographic Accelerator feature on new Intel/AMD CPUs)
- Support for certain thermal sensors via ACPI, coretemp, and amdtemp
- System startup beep can be disabled
- Separate powerd setting for when on battery
- Add optional ability to change the size of RAM disks for /var/ and /tmp/ for systems that have RAM to spare
- Add optional ability for full installs to use RAM disks for /var/ and /tmp/ as is done on NanoBSD. Reduces overall writes to the media, should be more SSD-friendly
- Use a custom sysDescr for snmp similar to m0n0wall's format. Fixes #2893
- Added tunable to allow disabling net.inet.udp.checksum - disabling UDP checksums can improve performance, but can also have negative side effects
- Added an mtree database with the correct default permissions, owner, sha256 sum, and some other information that is used to verify file permissions post-install and post-upgrade
- APC is not started for PHP unless the system has over 512MB RAM, to reduce memory usage on systems with low RAM

3.59.8 Multi-WAN

- DynDNS multi-WAN failover
- IPsec multi-WAN failover
- OpenVPN multi-WAN failover
- Changed descriptions of the values for gateway monitoring
- Display apinger (gateway monitoring daemon) as a service when it is enabled
- Fixes for apinger to reload via SIGHUP properly, to avoid unnecessary restarts and loss of gateway status data
- "State Killing on Gateway Failure" now kills ALL states when a gateway has been detected as down, not just states on the failing WAN. This is done because otherwise the LAN-side states were not killed before, and thus some connections would be in limbo, especially SIP.
- Due to the change in its behavior, "State Killing on Gateway Failure" is now disabled by default in new configurations and is disabled during upgrade. If the feature is desired, it must be manually re-enabled post-upgrade.

3.59.9 NTP

- NTP daemon now has GPS support

3.59.10 IPsec

- More IPsec hash algorithms and DH key groups added, "base" negotiation mode added
- Mobile IPsec supports separate "split dns" field and doesn't assume the default domain for split DNS domains
- Properly ignore disabled IPsec phase 2 entries
- NAT before IPsec (1:1 or many:1) outbound
- Set default Proposal Check setting to Obey for mobile IPsec
- LDAP and RADIUS are now possible authentication sources for IPsec mobile xauth
- Delete the SPDs for an old IPsec entry when it is disabled or removed #2719
- Manage active SPDs on CARP secondary during sync #2303
- Add an option to force IPsec to reload on failover, which is needed in some cases for IPsec to fail from one interface to another. #2896

3.59.11 OpenVPN

- OpenVPN can accept attributes from RADIUS via avpairs for things like inacl, outacl, dns-server, routes
- OpenVPN checkbox for "topology subnet" to use one IP per client in tun mode
- OpenVPN local/remote network boxes can accept multiple comma-separated networks
- OpenVPN status for SSL/TLS server instances can now display the routing table for the VPN instance
- OpenVPN now allows selecting "localhost" as the interface
- Gateways are created for assigned OpenVPN server instances as well as clients
- OpenVPN instances can run on the same port on different interfaces

- OpenVPN status page now has service controls to show the status of the daemon running each instance, and allow for stop/start/restart from that page
- Changed wording of the error displayed when a daemon is not running or the management interface of OpenVPN cannot be reached for an instance
- OpenVPN client-specific Override cleanup fixes
- Fixed double-click to edit of OpenVPN Client-Specific Overrides

3.59.12 NAT/Firewall Rules/Alias

- Aliases separated into tabs for Hosts, Ports, and URLs to improve manageability
- NAT reflection options re-worded to be less confusing
- Adjustable source tracking timeout for Sticky connections
- Firewall rules now support matching on ECE and CWR TCP flags
- Filtering on ECE and CWR TCP flags is now possible
- Added ICMP to protocol list when creating rdr (port forward) rules
- Keep proper positioning of duplicated outbound NAT rules #1118
- When using the + at the top of Outbound NAT rules, add the rule to the top of the list and not the bottom
- Fix ordering of interface group rules in the ruleset #2837
- Track time and user@host which created or updated a firewall, NAT port forward, or outbound NAT rule. If timestamp records are present, display them at the bottom of the rule page when editing. Have the created time/user pre-filled for automated rules such as NAT port forward associated rules and the switch from automatic to manual outbound NAT
- Fix generation of manual outbound NAT rules so that localhost and VPN rules are not unnecessarily duplicated
- Prevent using "block" for an alias name, as it is a pf reserved keyword
- Allow TCP flags to be used on block or reject rules, since they are also valid there
- Updates/fixes to DSCP handling
- Allow advanced options state-related parameters to be used for TCP, UDP, and ICMP – Formerly only allowed on TCP
- Respect ports found in rules when policy route negation rules are made, #3173
- Do not include disabled OpenVPN networks in generated policy route negation rules

3.59.13 Certificates

- Improved denoting of certificate purposes in the certificate list
- Imported CRLs can be edited and replaced
- Can set digest algorithm for CA/Certs (sha1, sha256, etc)
- Default digest algorithm is now SHA256
- Show CA and certificate start and end dates in the their listings
- Correct tooltip description when adding a certificate #3017

- Relax input validation on a CA/Cert description since it is only used cosmetically in pfSense and not in the actual CA/cert subject
- Allow removing blank/empty CA and Cert entries

3.59.14 Logging

- More system log separation, Gateways, Routing, Resolver split into their own tabs
- Firewall logs can now be filtered by many different criteria
- Firewall logs can be sorted by any column
- Firewall logs can optionally show the matching rule description in a separate column or in between rows
- Firewall logs now show an indicator icon if the direction of a log entry is OUT rather than IN
- Add popup DNS resolution method to firewall log view
- Reduced logging output from IGMP proxy
- Reduced logging output from DynDNS
- Relocated filterdns logs to the resolver log file/tab
- Relocated DHCP client logs to the DHCP tab
- Fix system script logging so the correct script filename is printed in the log, rather than omitting the script name entirely
- Add independent logging choices to disable logging of bogon network rules and private network rules. Add upgrade code to obey the existing behavior for users (if default block logging was disabled, so is bogon/private rule blocking)
- Add a checkbox to disable the lighttpd log for people who don't want their system log full of messages from lighttpd in some cases where they are filling the log unnecessarily

3.59.15 Notifications

- Add the ability to disable Growl or SMTP notifications but keep their settings intact, so the mail settings can be used for other purposes (packages, etc)
- Add a test button to selectively test Growl or SMTP notifications without re-saving settings
- Do not automatically generate a test notification on saving notification settings, as there are now individual test buttons

3.59.16 High Availability (CARP, pfSync, XML-RPC)

- High Availability Synchronization options (Formerly known as "CARP Settings" under Virtual IPs Promoted to its own menu entry, System > High Avail. Sync
- Ensure that the user does not remove only the last IP alias needed for a CARP VIP in an additional subnet
- Disable pfsync interface when state synchronization is not in use
- Fixed issues with DHCP server config synchronization ordering on secondary nodes #2600
- Restart OpenVPN servers when CARP transitions to master (clients were already restarted), otherwise if CARP was disabled, the servers would never recover

- Removed the automatic pfsync rule, since the documentation always recommends adding it manually, and to add it behind the scenes with no way to block it can be counter-productive (and potentially insecure). **If the documentation was not followed and a pfsync or allow all rule was not added on the sync interface, then state synchronization may break after this upgrade. Add an appropriate rule to the sync interface and it will work again.**
- Allow XMLRPC to sync IP Alias VIPs set to Localhost for their interface
- In DHCP leases view, use the internal interface name (lan/opt1/etc) for the failover pool name, rather than a number. In certain cases the number can get out of sync between the two nodes, but the interface names will always match
- Print the user-configured interface description next to the DHCP failover pool name, rather than only the internal name (lan/opt1/etc)
- Add option to synchronize authentication servers (RADIUS, LDAP) via XMLRPC

3.59.17 NanoBSD

- Fixes for conf_mount_ro/conf_mount_rw reference checking/locking
- Diag > NanoBSD now has button to switch media between read/write and read-only
- Diag > NanoBSD now has a checkbox option to keep the media read/write
- Fixed an issue with NanoBSD time zones not being properly respected by all processes the first reboot after a firmware upgrade

3.59.18 DHCP Server

- DHCP can support multiple pools inside a single subnet, with distinct options per pool
- DHCP can allow/deny access to a DHCP pool by partial (or full) MAC address
- DHCP static mappings can have custom settings for gateway, DNS, etc
- DHCP static mappings can optionally have a static ARP entry created
- Fix Dynamic DNS updates from DHCP (ISC changed the config layout and requires zone declarations)
- When crafting DHCP Dynamic DNS zones, do not use invalid DNS servers for the IP type (e.g. skip IPv6 DNS servers, because the DHCP daemon rejects them)
- Added a config backup section choice for DHCPv6

3.59.19 Traffic Shaper

- Schedules can now be used with limiters
- Traffic shaper queues view updated
- CoDel AQM Shaper Discipline
- Allow PRIQ queues to be deleted. #3037
- Limiters now allow the user to set the mask they want to use, rather than assuming masking will always be per-IP. This allows per-subnet limits and similar
- Limiters now allow setting masking for IPv6

- Limiters now allow setting a burst size. This will pass X amount of data (TOTAL, NOT a rate) after an idle period before enforcing the limit

3.59.20 DNS Forwarder

- In DNS forwarder, DNS query forwarding section with options for sequential and require domain
- Allow a null forwarding server in DNS Forwarder domain overrides to ensure that queries stay local and never go outside the firewall
- Add DNS Forwarder option to not forward private reverse lookups
- DNS Forwarder domain overrides can now specify a source address for the query, to help resolve hostnames over VPN tunnels
- DNS Forwarder now can change the port upon which it listens, for better cohabitation with other DNS software such as tinydns or unbound, if both are needed
- DNS Forwarder now has an option to select the interfaces/IP Addresses upon which it will respond to queries
- DNS Forwarder can now be set to only bind to specific IPv4 IPs (the underlying software, dnsmasq, does not support selectively binding to IPv6 IPs)
- Improved handling of some dnsmasq custom config options

3.59.21 User Manager

- Configurable RADIUS authentication timeout in User Manager
- Print the error message from LDAP in the log for a bind failure. Helps track down reasons for authentication failures
- Re-enable admin user if it's disabled when 'Reset webConfigurator password' option is used. Fixes #2877
- Restrict maximum group name length to 16 characters or less (OS restriction)
- Added option to UTF-8 encode LDAP parameters to improve handling of international characters
- CDATA protected LDAP fields in config to avoid invalid XML with international characters

3.59.22 DynDNS

- Fixed handling of DynDNS 25-day update and add ability to configure update interval
- Added DynDNS No-IP Free Account Support
- Add AAAA support to RFC2136 updates
- Add cached IP support to RFC2136, add GUI button to force update for single host
- Fix double click row to edit for RFC2136
- Add option to RFC2136 to find/use the public IP if the interface IP is private. (Off by default to preserve existing behavior on upgrade)
- Add server IP column and cached IP display to RFC2136 host list
- Include RFC2136 hosts in DNS rebinding checks
- Include both dyndns and RFC2136 hosts in referer check

3.59.23 Graphs

- Add ability to reverse-resolve IPs on Status > Traffic Graph in the rate table
- Add ability to filter local or remote IPs on Status > Traffic Graph in the rate table
- Change maximum values for RRD throughput to account for 10G links. Previous maximums would have caused blank spots on the graph during periods of high throughput
- Fixes to RRD data resolution/retention
- Added RRD Graph for mbuf clusters
- Changed default RRD graph colors to be more visually distinct to help avoid ambiguity between multiple values on the same graph

3.59.24 Misc

- Add option to the packet capture page to control whether or not promiscuous mode is used on the NIC. Certain drivers have issues with promiscuous mode
- Make parent interface and all VLANs share MTU #2786
- Fix cellular signal strength indicator
- Fix PPP config cleanup when removing an interface #2758
- Disallow adding IP Alias or CARP VIP that would be the network or broadcast address of a subnet
- Diagnostics > Sockets page to show open network sockets on the firewall
- Diagnostics > Test Port page to perform a simple TCP connection test to see if a port is open
- The pftop page has additional options to display more detailed information and sort it
- Fixed conflict between static IP and static route in the same subnet #2039
- Do not apply static ARP entries to disabled interfaces #1988
- Do not allow bridge members to be assigned to itself #1153
- Changed Diag > Ping to use more available source addresses (CARP VIPs, IP Alias VIPs, OpenVPN interfaces, IPv6 Link-Local IPs)
- Changed Diag > Traceroute to use more available source addresses (CARP VIPs, IP Alias VIPs, OpenVPN interfaces, IPv6 Link-Local IPs)
- Changed shell prompt to not force background color, to be kinder to those not using black as a background in their terminal
- Add a field to allow rejecting DHCP leases from a specific upstream DHCP server. #2704
- Updated the help system to handle some recent added files for 2.x and clean out some old/obsolete files
- Allow selecting “Localhost” as an interface for IP Alias VIPs - this way IP Alias VIPs may be used for binding firewall services (e.g. Proxy, VPN, etc) in routed subnets without burning IPs for CARP unnecessarily
- Updated list of mobile service providers
- Fix max length for wpa passphrase. A 64-char passphrase would be rejected by hostapd and leave an AP in an open state #3034
- Added MSS clamping to the setup wizard
- Add a setting to configure the filterdns hostname resolution interval (defaults to 300s, 5 minutes)

- Omit IP mismatch warnings (e.g. behind a port forward, VPN IP, etc) if HTTP_REFERER protection is disabled
- Fixes for selecting/detecting PPP devices such as 3G/4G modems
- Rather than doing auto-detection to find serial PPP devices, use a glob when listing potential PPP serial devices
- Prevent sshlockout from a crash/coredump if a format string like %s is present in the buffer
- Fix SMART to see adaX devices
- Fix SMART interpretation of output from SCSI devices
- Fixed display of user SSH keys when present
- Updated p0f database from FreeBSD
- Fix UPnP Interface name selection to show the configured description entered by the user
- Allow setting the external UPnP interface (must be default route WAN)
- Fix Diag > Tables AJAX fadeOut after deletion for rows with CIDR mask format
- Improve Diagnostics > Routes to fetch output via AJAX and have configurable filtering and sizes. Improves handling of large routing tables, such as a full BGP feed
- When deleting or renaming a virtual server from the Load Balancer (relayd) manually clean up the NAT rules it leaves behind to avoid conflicts
- Many, many bug fixes
- Various fixes for typos, formatting, input validation, etc

3.59.25 SH/PHP Shell Scripts

- Git package for gitsync is now pulled in as a pfSense-style PBI package
- Shell scripts added to enable/disable CARP:

```
pfSsh.php playback enablecarp
pfSsh.php playback disablecarp
```

- Shell scripts to add and remove packages from the command line:

```
pfSsh.php playback installpkg "Some Package"
pfSsh.php playback uninstallpkg "Some Package"
pfSsh.php playback listpkg
```

- Added shell script to remove shaper settings:

```
pfSsh.php playback removeshaper
```

- Add shell script to control services from the command line:

```
pfSsh.php playback svc start <service name>
pfSsh.php playback svc restart <service name>
pfSsh.php playback svc stop <service name>
```

- Add a simple CLI mail script capable of sending an SMTP message using echo/piped input (uses SMTP notification settings for server details):

```
ifconfig -a | mail.php -s"ifconfig output"
```

- Added a script to convert a user's filesystem from device names to UFS labels, for easier portability in case the disk device changes names (e.g. adX to adY, adX to daY, or adX to adaX). ONLY FOR FULL INSTALLS. NanoBSD already uses labels.

```
/usr/local/sbin/ufslabels.sh
```

3.60 2.0.3 New Features and Changes

pfSense® software 2.0.3 is a maintenance release with some bug fixes since 2.0.2 release. It is possible to upgrade from any previous release to 2.0.3.

Because this release shortly followed after 2.0.2, review the *2.0.2 New Features and Changes* document as well.

The changelog for pfSense 2.0.3-RELEASE follows.

3.60.1 Security Advisories

- Updated to OpenSSL 0.9.8y to address FreeBSD-SA-13:03

3.60.2 PPP

- Fix obtaining DNS servers from PPP type WANs (PPP, PPPoE, PPTP, L2TP)

3.60.3 Captive Portal

- Fix Captive Portal Redirect URL trimming
- Voucher sync fixes
- Captive portal pruning/locking fixes
- Fix problem with fastcgi crashing

3.60.4 OpenVPN

- Clear the route for an OpenVPN endpoint IP when restarting the VPN, to avoid a situation where a learned route from OSPF or elsewhere could prevent an instance from restarting properly
- Always clear the OpenVPN route when using shared key, no matter how the tunnel network "CIDR" is set
- Use the actual OpenVPN restart routine when starting/stopping from services rather than killing/restarting manually
- Allow editing an imported CRL, and refresh OpenVPN CRLs when saving. #2652
- Fix interface assignment descriptions when using > 10 OpenVPN instances

3.60.5 Logging

- Put syslogd into secure mode so it refuses remote syslog messages
- If syslog messages are in the log, and the hostname does not match the firewall, display the supplied hostname
- Fix PPP log display to use the correct log handling method
- Run IPsec logs through htmlspecialchars before display to avoid a potential persistent XSS from racoon log output (e.g. username)

3.60.6 Traffic Shaper

- Fix editing of traffic shaper default queues. #1995
- Fix wording for VoIP address option in the shaper. Add rule going the other direction to catch connections initiated both ways

3.60.7 Dashboard & General GUI

- Use some tweaks to PHP session management to prevent the GUI from blocking additional requests while others are active
- Remove cmd_chain.inc and preload.php to fix some issues with lighttpd, fastcgi, and resource usage
- Firmware settings manifest (Site list) now bolds and denotes entries that match the current architecture, to help avoid accidental cross-architecture upgrades
- Add header to DHCP static mappings table
- When performing a factory reset in the GUI, change output style to follow halt.php and reboot.php so the shutdown output appears in the correct location on the page
- Better validation of parameters passed during S.M.A.R.T. operations for testing HDDs
- Fixed SNMP interface binding glitch (Setting was active but not reflected when viewed in GUI)
- Add a new class called addgatewaybox to make it easier to respect custom themes #2900

3.60.8 Console Menu Changes

- Correct accidental interface assignment changes when changing settings on the console menu
- Console menu option 11 now kills all active PHP processes, kills lighttpd, and then restarts the GUI. This is a more effective way to restart the GUI since if a PHP process is hung, restarting lighttpd alone will not recover from that
- Fix port display after LAN IP reset

3.60.9 Misc Changes

- Change how the listening address is passed to miniupnpd, the old method was resulting in errors for some users
- Fix "out" packet count reporting
- Be a little smarter about the default kernel in rare cases where it cannot determine what was in use
- Pass -S to tcpdump to avoid an increase in memory consumption over time in certain cases
- Minimise rewriting of /etc/gettytab (https://forum.netgate.com/post/51581)
- Make is_pid_running function return more conistent results by using isvalidpid
- Fix ataidle error on systems that have no ATA HDD. #2739
- Update Time Zone database zoneinfo to 2012.j to pick up on recent zone/DST/etc changes
- Fix handling of LDAP certificates, the library no longer properly handles files with spaces in the CA certificate filename
- Bring in the RCFILEPREFIX as constant fixes from HEAD, since otherwise rc.stop_packages was globbing in the wrong dir and executing the wrong scripts. Also seems to have fixed the "bad fd" error
- NTP restart fixes
- Gitsync now pulls in git package from pfSense package repository rather than FreeBSD
- Fixed handling of RRD data in config.xml backups when exporting an encrypted config #2836
- Moved apinger status to /var/run instead of /tmp
- Fixes for FTP proxy on non-default gateway WANs
- Fixes for OVA images
- Use new pfSense repository location (http://github.com/pfsense/pfsense/)
- Add patch to compensate apinger calculation for down gateways by time taken from other tasks like rrd/status file/etc

3.60.10 lighttpd changes

- Improve tuning of lighttpd and php processes
- Use separate paths for GUI and Captive Portal fastcgi sockets
- Always make sure php has its own process manager to make lighttpd happy
- Make mod_fastcgi last to have url.rewrite work properly
- Enable mod_evasive if needed for Captive Portal
- Simplify lighttpd config
- Send all lighttpd logs to syslog

3.60.11 Binary changes

- dnsmasq to 2.65
- rsync to 3.0.9
- links 2.7
- rrdtool to 1.2.30
- PHP to 5.2.17_13
- OpenVPN 2.2 stock again (Removed IPv6 patches since those are only needed on 2.1 now)
- Fix missing “beep” binary on amd64
- Fix potential issue with IPsec routing of client traffic
- Remove lighttpd spawnfcgi dependency
- Add splash device to wrap_vga kernels (It’s in GENERIC so full installs already have it). #2723

filterdns

- Correct an issue with unallocated structure
- Avoid issues with pidfiles being overwritten, lock the file during modifications
- Make filterdns restartable and properly cleanup its tables upon exit or during a reconfiguration

dhcpleases

- Correct use after free and also support hostnames with other DNS suffix
- Reinit on any error rather than just forgetting. Also the difftime checks are done after having complete view, no need to do them every time
- Typo fixes
- Log that a HUP signal is being sent to the pid file submitted by argument
- Prevent bad parsing of empty hostnames in lease file. Add an f option to run dhcplease in foreground. The only option needed while in foreground is h parameter and the only usable one as well

3.61 2.0.2 New Features and Changes

pfSense® software 2.0.2 is a maintenance release with some bug and security fixes since 2.0.1 release. It is possible to upgrade from any previous release to 2.0.2.

What follows is a mostly-complete changelog for pfSense 2.0.2-RELEASE

3.61.1 FreeBSD Security Advisories

Base OS updated to 8.1-RELEASE-p13 to address the following FreeBSD Security Advisories:

- FreeBSD-SA-12:01.openssl (v1.0/v1.1) http://security.freebsd.org/advisories/FreeBSD-SA-12:01.openssl.asc
- FreeBSD-SA-12:02.crypt http://security.FreeBSD.org/advisories/FreeBSD-SA-12:02.crypt.asc
- FreeBSD-SA-12:04.sysret (v1.0/v1.1) http://security.FreeBSD.org/advisories/FreeBSD-SA-12:04.sysret.asc
- FreeBSD-SA-12:07.hostapd https://www.freebsd.org/security/advisories/FreeBSD-SA-12:07.hostapd.asc

3.61.2 PPTP

- Added a warning to PPTP VPN configuration page

> **Warning: PPTP is no longer considered a secure VPN technology because it relies upon MS-CHAPv2 which has been compromised. Be aware that intercepted traffic can be decrypted by a third party, so it should be considered unencrypted. Migrate to another VPN type such as OpenVPN or IPsec.**
>
> More information on this can be found at https://isc.sans.edu/diary/End+of+Days+for+MS-CHAPv2/13807 and https://www.cloudcracker.com/blog/2012/07/29/cracking-ms-chap-v2/

- Fix reference to PPTP secondary RADIUS server shared secret.
- PPTP upgrade fixes.

3.61.3 NTP Changes

- OpenNTPD was dropped in favor of the ntp.org NTP daemon, used by FreeBSD.
- Status page added (Status > NTP) to show status of clock sync
- NTP logging fixed.

Note: `ntpd` will bind/listen to all interfaces by default, and it has to in order to receive replies. Selective interface binding may still be used to control which IP addresses will accept traffic, but be aware that the default behavior has changed.

3.61.4 Dashboard & General GUI Fixes

- Various fixes for typos, wording, and so on.
- Do not redirect on saving services status widget.
- Don't use $pconfig in widgets, it has unintended side effects.
- Fix display of widgets with configuration controls in IE.
- Changed some padding/margin in the CSS in order to avoid wrapping the menu.
- #2165 Change to embed to prevent IE9 from misbehaving when loading the Traffic Graph page

3.61.5 OpenVPN Fixes

- Safer for 1.2.3 upgrades to assume OpenVPN interface == any, since 1.2.3 didn't have a way to bind to an interface. Otherwise people accepting connections on opt interfaces on 1.2.3 will break on upgrade until the proper interface is selected in the GUI
- Don't ignore when multiple OpenVPN DNS, NTP, WINS, etc servers were specified in 1.2.3 when upgrading. 1.2.3 separated by ;, 2.x uses separate vars.
- Fix upgrade code for 1.2.3 with assigned OpenVPN interface.
- Fix LZO setting for Upgraded OpenVPN (was turning compression on even if old config had it disabled.)
- Be more intelligent when managing OpenVPN client connections bound to CARP VIPs. If the interface is in BACKUP status, do not start the client. Add a section to rc.carpmaster and rc.carpbackup to trigger this start/stop. If an OpenVPN client is active on both the master and backup system, they will cause conflicting connections to the server. Servers do not care as they only accept, not initiate.

3.61.6 IPsec fixes

- Only do foreach on IPsec p2's if it's actually an array.
- #2201 Don't let an empty subnet into racoon.conf, it can cause parse errors.
- #2201 Reject an interface without a subnet as a network source in the IPsec Phase 2 GUI.
- Add routes even when IPsec is on WAN, as WAN may not be the default gateway.
- #1986 Revamped IPsec status display and widget to properly account for mobile clients.
- Fixed a bug that caused the IPsec status and widget to display slowly when mobile clients were enabled.

3.61.7 User Manager Fixes

- #2066 Improve adding/removing of users accounts to the underlying OS, especially accounts with a numeric username.
- Include admin user in bootup account sync
- Fix permission and certificate display for the admin user
- Fix ssh key note to refer to DSA not just RSA since both work.
- ":" chars are invalid in a comment field, filter them out.
- When renaming a user, make sure to remove the previous user or it gets left in /etc/passwd.
- #2326 Do not allow empty passwords since this might cause problems for some authentication servers like LDAP.

3.61.8 Captive Portal Fixes

- Take routing table into account when figuring out which IP address to use for talking to CP clients.
- Prevent browser auto-fill username and password on voucher config, as it can interfere with the settings being properly saved if sync isn't fully configured, which this can make happen accidentally.
- Correct the Called-Station-Id attribute setting to be the same on STOP/START packets
- Correct the Called-Station-Id attribute setting to be consistent on the data sent
- #2082 Correct the log to display the correct information about an existing session

- #2052 Remove duplicate rule
- Fix which roll to write when writing the active voucher db
- Always load ipfw when enabling CP to ensure the pfil hooks are setup right
- #2378 Fix selection of CP interfaces when using more than 10 opt interfaces.
- Strengthen voucher randomization.

3.61.9 NAT/Firewall Rules/Alias Fixes

- #2327 Respect the value of the per-rule "disable reply-to" checkbox.
- #1882 Fix an invalid pf rule generated from a port forward with dest=any on an interface with ip=none
- #2163 1:1 Reflection fixes for static route subnets and multiple subnets on the same interface.
- Better validation on URL table alias input from downloaded files.
- #2293 Don't put an extra space after "pass" when assuming it as the default action or later tests will fail to match this as a pass rule.
- Update help text for Host aliases to indicate FQDNs are allowed.
- #2210 Go back to scrub rather than "scrub in", the latter breaks MSS clamping for egress traffic
- Fix preservation of the selection of interfaces on input errors for floating rules.
- Fix URL table update frequency box.
- Fix input validation for port forwards, Local Port must be specified.
- Added a setting to increase the maximum number of pf tables, and increased the default to 3000.
- Properly determine active GUI and redirect ports for anti-lockout rule, for display and in the actual rule.
- Handle loading pf limits (timers, states, table/entry limits, etc) in a separate file to avoid a chicken-and-egg scenario where the limits would never be increased properly.

3.61.10 Interface/Bridging Fixes

- Correct checking if a gif is part of bridge so that it actually works correctly adding a gif after having created it on bootup
- Use the latest functions from pfSense module for getting interface list
- Use the latest functions from pfSense module for creating bridges
- Implement is_jumbo_capable in a more performant way. This should help with large number of interfaces
- Since the CARP interface name changed to "vipN" from "carpN", devd needs to follow that change as well.
- #2242 Show lagg protocol and member interfaces on Status > Interfaces.
- #2212 Correctly stop dhclient process when an interface is changed away from DHCP.
- Fixed 3G SIM PIN usage for Huawei devices
- Properly obey MTU set on Interface page for PPP type WANs.

3.61.11 Other Misc. Fixes

- #2057 Add a checkbox that disables automatically generating negate rules for directly connected networks and VPNs.
- Mark "Destination server" as a required field for DHCP Relay
- Clarify the potential pitfalls when setting the Frequency Probe and Down parameters.
- Add a PHP Shell shortcut to disable referer check (playback disablerefercheck)
- #2040 Make Wireless Status tables sortable
- #2068 Fix multiple keys in a file for RFC2136 dyndns updates.
- Check to see if the pid file exists before trying to kill a process
- #2144 Be smarter about how to split a Namecheap hostname into host/domain.
- Add a small script to disable APM on ATA drives if they claim to support it. Leaving this on will kill drives long-term, especially laptop drives, by generating excessive Load Cycles. The APM bit set will persist until the drive is power cycled, so it's necessary to run on each boot to be sure.
- #2158 Change SNMP binding option to work on any eligible interface/VIP. If the old bindlan option is there, assume the lan interface for binding.
- Fix reference to PPTP secondary RADIUS server shared secret.
- PPTP upgrade fixes.
- #2147 Add button to download a .p12 of a cert+key.
- #2233 Carry over the key length on input errors when creating a certificate signing request.
- #2207 Use PHP's built-in RFC 2822 date format.
- Allow specifying the branch name after the repository URL for gitsync command-line arguments and remove an unnecessary use of the backtick operator.
- Correct send_multiple_events to conform with new check_reload_status behaviour
- Do not wipe logs on reboot on full install
- Set FCGI_CHILDREN to 0 since it does not make sense for php to manage itself when lighttpd is doing so. This makes it possible to recover from 550-Internal… error.
- Support for xmlrpcauthuser and xmlrpcauthpass in $g.
- Fix Layer 7 pattern upload, button text check was incorrect.
- Correct building of traffic shaping queue to not depend on parent mask
- #2239 Add alias support to static routes
- Use !empty instead of isset to prevent accidental deletion of the last used repository URL when firmware update gitsync settings have been saved without a repository URL.
- Better error handling for crypt_data and also better password argument handling
- Stop service needs to wait for the process to be stopped before trying to restart it.
- Use a better default update url
- Fix missing description in rowhelper for packages.
- #2402, #1564 Move the stop_packages code to a function, and call the function from the shell script, and call the function directly for a reboot.

- #1917 Fix DHCP domain search list
- Update Time Zone zoneinfo database using latest zones from FreeBSD
- Handle HTTPOnly and Secure flags on cookies
- Fixed notifications for firmware upgrade progress
- Removed an invalid declaration that considered 99.0.0.0/8 a private address.
- Fixed redirect request for IE8/9
- #1049 Fix crashes on NanoBSD during package removal/reinstall. Could result in the GUI being inaccessible after a firmware update.
- Fix some issues with upgrading NanoBSD+VGA and NanoBSD+VGA Image Generation
- Fix issues upgrading from systems with the old "Uniprocessor" kernel which no longer exists.
- Fix a few potential XSS/CSRF vectors.
- Fixed issue with login page not showing the correct selected theme in certain configurations.
- Fix limiters+multi-wan

3.61.12 Binary/Supporting Program Updates

- Some cleanup to reduce overall image size
- Fixes to ipfw-classifyd file reading and handling
- Updated miniupnpd
- ISC DHCPD 4.2.4-P1
- mdp5 upgraded to 5.6
- pftop updated
- lighttpd updated to 1.4.32, for CVE-2011-4362 and CVE-2012-5533.

3.62 2.0.1 New Features and Changes

This is a maintenance release with bug and security fixes since 2.0 release. It is possible to upgrade from any previous release to 2.0.1.

For those who use the built in certificate manager, pay close attention to the notes below on a potential security issue with those certificates.

3.62.1 Change Log

The following changes were made after 2.0-RELEASE and were included in 2.0.1-RELEASE.

- Improved accuracy of automated state killing in various cases (#1421)
- Various fixes and improvements to relayd
  - Added to Status > Services and widget
  - Added ability to kill relayd when restarting (#1913)
  - Added DNS load balancing

  - Moved relayd logs to their own tab
  - Fixed default SMTP monitor syntax and other send/expect syntax
- Fixed path to FreeBSD packages repo for 8.1
- Various fixes to syslog:
  - Fixed syslogd killing/restarting to improve handling on some systems that were seeing GUI hangs resetting logs
  - Added more options for remote syslog server areas
  - Fixed handling of 'everything' checkbox
  - Moved wireless to its own log file and tab
- Removed/silenced some irrelevant log entries
- Fixed various typos
- Fixes for RRD upgrade/migration and backup (#1758)
- Prevent users from applying NAT to CARP which would break CARP in various ways (#1954)
- Fixed policy route negation for VPN networks (#1950)
- Fixed "Bypass firewall rules for traffic on the same interface" (#1950)
- Fixed VoIP rules produced by the traffic shaper wizard (#1948)
- Fixed uname display in System Info widget (#1960)
- Fixed LDAP custom port handling
- Fixed Status > Gateways to show RTT and loss like the widget
- Improved certificate handling in OpenVPN to restrict certificate chaining to a specified depth – CVE-2011-4197
- Improved certificate generation to specify/enforce type of certificate (CA, Server, Client) – CVE-2011-4197
- Clarified text of serial field when importing a CA (#2031)
- Fixed MTU setting on upgrade from 1.2.3, now upgrades properly as MSS adjustment (#1886)
- Fixed Captive Portal MAC passthrough rules (#1976)
- Added tab under Diagnostics > States to view/clear the source tracking table if sticky is enabled
- Fixed CARP status widget to properly show "disabled" status.
- Fixed end time of custom timespan RRD graphs (#1990)
- Fixed situation where certain NICs would constantly cycle link with MAC spoofing and DHCP (#1572)
- Fixed OpenVPN ordering of client/server IPs in Client-Specific Override entries (#2004)
- Fixed handling of OpenVPN client bandwidth limit option
- Fixed handling of LDAP certificates (#2018, #1052, #1927)
- Enforce validity of RRD graph style
- Fixed crash/panic handling so it will do textdumps and reboot for all, and not drop to a db> prompt.
- Fixed handling of hostnames in DHCP that start with a number (#2020)
- Fixed saving of multiple dynamic gateways (#1993)
- Fixed handling of routing with unmonitored gateways

- Fixed Firewall > Shaper, By Queues view
- Fixed handling of spd.conf with no phase 2's defined
- Fixed synchronization of various sections that were leaving the last item on the slave (IPsec phase 1, Aliases, VIPs, etc)
- Fixed use of quick on internal DHCP rules so DHCP traffic is allowed properly (#2041)
- Updated ISC DHCP server to 4.2.3 (#1888) – this fixes a denial of service vulnerability in dhcpd.
- Added patch to mpd to allow multiple PPPoE connections with the same remote gateway
- Lowered size of CF images again fix newer and ever-shrinking CF cards.
- Clarified text for media selection (#1910)

3.62.2 Notes for certificate generation vulnerability

Certificates generated with the built-in certificate manager in all 2.0 versions prior to 2.0.1 are excessively permissive for non-CA certificates. These certificates can be used as a certificate authority, meaning a user can use their own certificate to create chained certificates. The firewall defaults OpenVPN on 2.0.1 and newer versions to not accept chained certificates, which mitigates this. However, if untrusted users have certificates generated from 2.0 release, the best practice is to regenerate all certificates and issuing new ones. Certificates generated by easy-rsa and imported into 2.0 are not affected. If using certificates generated on pfSense® for other purposes, revoke those and issue new certificates generated on 2.0.1. A CRL must be utilized in that case. To be on the safe side, start from scratch with a new CA and certificates after deleting all existing ones.

3.62.3 Upgrade considerations

It is very important to read the *Upgrade Guide* before performing an upgrade for those still on 1.2.x versions.

3.63 2.0 New Features and Changes

This is a partial list of the new features and major changes in the pfSense® software 2.0 release.

3.63.1 Operating System

- Based on FreeBSD 8.1 release.
- i386 and amd64 variants for all install types (full install, nanobsd/embedded, etc.)
- USB memstick installer images available

3.63.2 Interfaces

- GRE tunnels
- GIF tunnels
- 3G support
- Dial up modem support
- Multi-Link PPP (MLPPP) for bonding PPP connections (ISP/upstream must also support MLPPP)
- *LAGG Interfaces*
- Interface groups
- IP Alias type Virtual IPs
- IP Alias VIPs can be stacked on CARP VIPs to go beyond the 255 VHID limit in deployments that need very large numbers of CARP VIPs.
- QinQ VLANs
- Can use Block Private Networks / Block Bogon Networks on any interface
- All interfaces are optional except WAN
- All interfaces can be renamed, even LAN/WAN
- Bridging enhancements - can now control all options of if_bridge, and assign bridge interfaces

3.63.3 Gateways/Multi-WAN

- Gateways, including dynamic gateways, are specified under System > Routing
- Gateways can have custom monitor IPs
- Gateways can have a custom weight, allowing load balancing to have ratios between WANs of different speeds
- Gateways can have custom latency, loss, and downtime trigger levels.
- Gateway monitoring via icmp is now configurable.
- Multiple gateways may exist per interface
- Multi-WAN is now handled via gateway groups
- Gateway groups can include multiple tiers with any number of gateways on each, for complex failover and load balancing scenarios.

3.63.4 General Web GUI

- Set to HTTPS by default, HTTP redirects to HTTPS port
- Dashboard and widgets added
- System > Advanced screen split into multiple tabs, more options available.
- SMTP email alerts and growl alerts
- New default theme - pfsense_ng
- Some community-contributed themes added

- Contextual help available on every page in the web interface, linking to a webpage containing help and documentation specific to that page.
- Help menu for quick access to online resources (forum, docs, paid support, etc.)

3.63.5 Aliases

- Aliases may be nested (aliases in aliases)
- Alias autocomplete is no longer case sensitive
- IP Ranges in Aliases
- More Alias entries supported
- Bulk Alias importing
- URL Aliases
- URL Table Aliases - uses a pf persist table for large (40,000+) entry lists

3.63.6 Firewall

- Traffic shaper rewritten - now handles any combination of multi-WAN and multi-LAN interfaces. New wizards added.
- Layer7 protocol filtering
- *EasyRule - add firewall rules from log view (and from console!)*
- Floating rules allow adding non-interface specific rules
- Dynamically sized state table based on amount of RAM in the system
- More Advanced firewall rule options
- FTP helper now in kernel
- TFTP proxy
- Schedule rules are handled in pf, so they can use all the rule options.
- State summary view, report shows states grouped by originating IP, destination IP, etc.

3.63.7 NAT

- All of the NAT screens were updated with additional functionality
- Port forwards can now handle create/update associated firewall rules automatically, instead of just creating unrelated entries.
- Port forwards can optionally use “rdr pass” so no firewall rule is needed.
- Port forwards can be disabled
- Port forwards can be negated (“no rdr”)
- Port forwards can have source and destination filters
- NAT reflection improvements, including NAT reflection for 1:1 NAT
- Per-entry NAT reflection overrides

- 1:1 NAT rules can specify a source and destination address
- 1:1 NAT page redesigned
- Outbound NAT can now translate to an address pool (Subnet of IPs or an alias of IPs) of multiple external addresses
- Outbound NAT rules can be specified by protocol
- Outbound NAT rules can use aliases
- Improved generation of outbound NAT rules when switching from automatic to manual.

3.63.8 IPsec

- Multiple IPsec p2's per p1 (multiple subnets)
- IPsec xauth support
- IPsec transport mode added
- IPsec NAT-T
- Option to push settings such as IP, DNS, etc, to mobile IPsec clients (mod_cfg)
- Mobile IPsec works with iOS and Android (Certain versions, see *IPsec Remote Access VPN Example Using IKEv1 with Xauth*)
- More Phase 1/2 options can be configured, including the cipher type/strength
- ipsec-tools version 0.8

3.63.9 User Manager

- New user manager, centralizing the various user configuration screens previously available.
- Per-page user access permissions for administrative users
- Three built-in authentication types - local users, LDAP and RADIUS.
- Authentication diagnostics page

3.63.10 Certificate Manager

- Certificate manager added, for handling of IPsec, web interface, user, and OpenVPN certificates.
- Handles creation/import of Certificate Authorities, Certificates, Certificate Revocation lists.
- Eliminates the need for using command line tools such as EasyRSA for managing certificates.

3.63.11 OpenVPN

- OpenVPN wizard guides through making a CA/Cert and OpenVPN server, sets up firewall rules, and so on. Greatly simplifies the process of creating a remote access OpenVPN server.
- OpenVPN filtering - an OpenVPN rules tab is available, so OpenVPN interfaces don't have to be assigned to perform filtering.
- OpenVPN client export package - provides a bundled Windows installer with certificates, Viscosity export, and export of a zip file containing the user's certificate and configuration files.
- OpenVPN status page with connected client list – can also kill client connections
- User authentication and certificate management
- RADIUS and LDAP authentication support

3.63.12 Captive Portal

- Voucher support added
- Multi-interface capable
- Pass-through MAC bandwidth restrictions
- Custom logout page contents can be uploaded
- Allowed IP addresses bandwidth restrictions
- Allowed IP addresses supports IP subnets
- "Both" direction added to Allowed IP addresses
- Pass-through MAC Auto Entry - upon successful authentication, a pass-through MAC entry can be automatically added.
- Ability to configure calling station RADIUS attributes

3.63.13 Wireless

- Virtual AP (VAP) support added
- more wireless cards supported with the FreeBSD 8.1 base

3.63.14 Server Load Balancing

- relayd and its more advanced capabilities replace slbd.

3.63.15 Other

- L2TP VPN added
- DNS lookup page added
- PFTop and Top in GUI - realtime updates
- Config History now includes a diff feature
- Config History has download buttons for prior versions
- Config History has mouseover descriptions
- CLI filter log parser (/usr/local/bin/filterparser)
- Switched to PHP 5.2.x
- IGMP proxy added
- Multiple Dynamic DNS account support, including full multi-WAN support and multi-accounts on each interface.
  - DynDNS Account Types supported are:

    DNS-O-Matic, DynDNS (dynamic), DynDNS (static), DynDNS (custom), DHS, DyNS, easyDNS, No-IP, ODS.org, ZoneEdit, Loopia, freeDNS, DNSexit, OpenDNS, Namecheap.com
- More interface types (VPNs, etc) available for packet capture
- DNS Forwarder is used by the firewall itself for DNS resolution (configurable) so the firewall benefits from faster resolution via multiple concurrent queries, sees all DNS overrides/DHCP registrations, etc.
- DHCP Server can now handle arbitrary numbered options, rather than only options present in the GUI.
- Automatic update now also works for NanoBSD as well as full installs
- More configuration sections can be synchronized via XMLRPC between CARP nodes.

3.64 General Release Information

3.64.1 Versions of pfSense software and FreeBSD

The tables in this document contain detailed information on pfSense® software releases.

Versions are grouped up by major/minor number changes so they are easier to locate. The most recent versions are listed first, and the rest are in descending order by release date.

- *pfSense Plus software*
- *pfSense CE software*
- *Legend*
- *Understanding pfSense Plus and CE software version numbers*

pfSense Plus software

25.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
26.03		TBD	24.1	16.0-CURRENT@3f5f52216f7e	plus-RELENG_26_03
25.11		2025-12-11	24.1	16.0-CURRENT@44f3e9f7f6c9	plus-RELENG_25_11
25.07.1		2025-08-18	24.0	15.0-CURRENT@bf06074106cf	plus-RELENG_25_07_1
25.07		2025-08-04	24.0	15.0-CURRENT@bf06074106cf	plus-RELENG_25_07

24.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
24.11		2024-11-25	23.6	15.0-CURRENT@f8a46de2dd48	plus-RELENG_24_11
24.03		2024-04-23	23.3	15.0-CURRENT@a5a965d75934	plus-RELENG_24_03

23.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
23.09.1		2023-12-07	23.3	14.0-CURRENT@0c783a37d5d5	plus-RELENG_23_09
23.05.1		2023-06-29	22.9	14.0-CURRENT@0c59e0b4e581	plus-RELENG_23_05_1
23.05		2023-05-22	22.9	14.0-CURRENT@0c59e0b4e581	plus-RELENG_23_05
23.01		2023-02-15	22.8	14.0-CURRENT@aec9453fec7	plus-RELENG_23_01

22.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
22.05.1	✕	2022-12-06	22.7	12.3-STABLE@5f81a4619dcf	plus-RELENG_22_05_1
22.05	✕	2022-06-26	22.7	12.3-STABLE@5f81a4619dcf	plus-RELENG_22_05
22.01	✕	2022-02-14	22.2	12.3-STABLE@ef1e43df92c6	plus-RELENG_22_01

21.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
21.05.2	✕	2021-10-26	21.7	12.2-STABLE@424f6363927	plus-RELENG_21_05_2
21.05.1	✕	2021-08-05	21.7	12.2-STABLE@424f6363927	plus-RELENG_21_05_1
21.05	✕	2021-06-02	21.7	12.2-STABLE@424f6363927	plus-RELENG_21_05
21.02.2	✕	2021-04-13	21.5	12.2-STABLE@f4d0bc6aa6b	plus-RELENG_21_02_2
21.02-p1	✕	2021-02-25	21.4	12.2-STABLE@f4d0bc6aa6b	plus-RELENG_21_02
21.02	✕	2021-02-17	21.4	12.2-STABLE@f4d0bc6aa6b	plus-RELENG_21_02

pfSense CE software

2.9.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.9.0	🕓	TBD	24.1	16.0-CURRENT@3f5f52216f7e	RELENG_2_9_0

2.8.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.8.1	✓	2025-09-04	24.0	15.0-CURRENT@bf06074106cf	RELENG_2_8_1
2.8.0	✓	2025-05-28	24.0	15.0-CURRENT@bf06074106cf	RELENG_2_8_0

2.7.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.7.2	✗	2023-12-07	23.3	14.0-CURRENT@0c783a37d5d5	RELENG_2_7_1
2.7.1	✗	2023-11-16	23.3	14.0-CURRENT@0c783a37d5d5	RELENG_2_7_1
2.7.0	✗	2023-06-29	22.9	14.0-CURRENT@0c59e0b4e581	RELENG_2_7_0

2.6.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.6.0	✗	2022-02-14	22.2	12.3-STABLE@ef1e43df92c6	RELENG_2_6_0

2.5.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.5.2	✗	2021-07-07	21.7	12.2-STABLE@f4d0bc6aa6b	RELENG_2_5_2
2.5.1	✗	2021-04-13	21.5	12.2-STABLE@f4d0bc6aa6b	RELENG_2_5_1
2.5.0	✗	2021-02-17	21.4	12.2-STABLE@f4d0bc6aa6b	RELENG_2_5_0

2.4.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.4.5-p1	✗	2020-06-09	19.1	11.3-STABLE@r357046	RELENG_2_4_5
2.4.5	✗	2020-03-26	19.1	11.3-STABLE@r357046	RELENG_2_4_5
2.4.4-p3	✗	2019-05-20	19.1	11.2-RELEASE-p10	RELENG_2_4_4
2.4.4-p2	✗	2019-01-07	18.9	11.2-RELEASE-p4	RELENG_2_4_4
2.4.4-p1	✗	2018-12-03	18.9	11.2-RELEASE-p4	RELENG_2_4_4
2.4.4	✗	2018-09-24	18.8	11.2-RELEASE-p3	RELENG_2_4_4
2.4.3-p1	✗	2018-05-14	18.0	11.1-RELEASE-p10	RELENG_2_4_3
2.4.3	✗	2018-03-29	17.9	11.1-RELEASE-p7	RELENG_2_4_3
2.4.2-p1	✗	2017-12-14	17.3	11.1-RELEASE-p6	RELENG_2_4_2
2.4.2	✗	2017-11-20	17.3	11.1-RELEASE-p4	RELENG_2_4_2
2.4.1	✗	2017-10-24	17.3	11.1-RELEASE-p2	RELENG_2_4_1
2.4	✗	2017-10-12	17.0	11.1-RELEASE-p1	RELENG_2_4_0

2.3.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.3.5-p2	✕	2018-05-14	15.8	10.3-RELEASE-p26	RELENG_2_3_5
2.3.5-p1	✕	2017-12-14	15.8	10.3-RELEASE-p26	RELENG_2_3_5
2.3.5	✕	2017-10-31	15.8	10.3-RELEASE-p20	RELENG_2_3_5
2.3.4-p1	✕	2017-07-20	15.8	10.3-RELEASE-p19	RELENG_2_3_4
2.3.4	✕	2017-05-04	15.8	10.3-RELEASE-p19	RELENG_2_3_4
2.3.3-p1	✕	2017-03-09	15.8	10.3-RELEASE-p17	RELENG_2_3_3
2.3.3	✕	2017-02-20	15.8	10.3-RELEASE-p16	RELENG_2_3_3
2.3.2	✕	2016-07-19	15.5	10.3-RELEASE-p5	RELENG_2_3_2
2.3.1	✕	2016-05-18	15.4	10.3-RELEASE-p3	RELENG_2_3_1
2.3	✕	2016-04-12	15.0	10.3-RELEASE	RELENG_2_3_0

2.2.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.2.6	✕	2015-12-21	12.0	10.1-RELEASE-p25	RELENG_2_2
2.2.5	✕	2015-11-05	12.0	10.1-RELEASE-p24	RELENG_2_2
2.2.4	✕	2015-07-26	11.9	10.1-RELEASE-p15	RELENG_2_2
2.2.3	✕	2015-06-24	11.7	10.1-RELEASE-p13	RELENG_2_2
2.2.2	✕	2015-04-15	11.7	10.1-RELEASE-p9	RELENG_2_2
2.2.1	✕	2015-03-17	11.7	10.1-RELEASE-p6	RELENG_2_2
2.2	✕	2015-01-23	11.6	10.1-RELEASE-p4	RELENG_2_2

2.1.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.1.5		2014-08-27	10.1	8.3-RELEASE-p16	RELENG_2_1
2.1.4		2014-06-25	10.1	8.3-RELEASE-p16	RELENG_2_1
2.1.3		2014-05-02	10.1	8.3-RELEASE-p16	RELENG_2_1
2.1.2		2014-04-10	10.1	8.3-RELEASE-p14	RELENG_2_1
2.1.1		2014-04-04	10.1	8.3-RELEASE-p14	RELENG_2_1
2.1		2013-09-15	9.8	8.3-RELEASE-p11	RELENG_2_1

2.0.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
2.0.3		2013-04-15	8.0	8.1-RELEASE-p13	RELENG_2_0
2.0.2		2012-12-21	8.0	8.1-RELEASE-p13	RELENG_2_0
2.0.1		2011-12-20	8.0	8.1-RELEASE-p6	RELENG_2_0
2.0		2011-09-17	8.0	8.1-RELEASE-p4	RELENG_2_0

1.2.x

Version	Support	Released	Config Rev	FreeBSD Version	Branch
1.2.3		2009-12-10	3.0	7.2-RELEASE-p5	RELENG_1_2
1.2.2		2009-01-09	3.0	7.0-RELEASE-p8	RELENG_1_2
1.2.1		2008-12-26	3.0	7.0-RELEASE-p7	RELENG_1_2
1.2		2008-02-25	3.0	6.2-RELEASE-p11	RELENG_1_2

Legend

Version
The pfSense Plus or CE software version number. When possible, the version number links to the release notes detailing what was changed in that particular release.

See also:

See *Understanding pfSense Plus and CE software version numbers* later in this document for an explanation of the version number formats.

Support
The support status.



Current supported release



Previous unsupported release



Future release

TBD
To Be Determined, not yet known.

Released
The date a specific version of pfSense software was released to the public.

Config Rev
The internal `config.xml` revision number, which indicates changes to the configuration format that may make a configuration file incompatible with older versions.

FreeBSD Version
Each version of pfSense software is based on a specific version of FreeBSD. The underlying FreeBSD version is listed for each corresponding version of pfSense software.

Branch
A link to the pfSense software source code branch used to build a specific release.

Understanding pfSense Plus and CE software version numbers

pfSense Plus and CE software utilize different version number formats. This makes it easier to distinguish between them and also makes it clear that the releases do not necessarily happen at the same time, even if they share a common code base.

pfSense Plus software version numbers use the format `<year>.<month>.<patch>` where the `<patch>` suffix is omitted when the value is `0`. This version numbering scheme follows the format used by TNSR software, also produced by Netgate, which in turn is modeled after the version format used by the Linux Foundation FD.io project. This change happened at the start of 2021 when the name changed from "pfSense Factory Edition" to "pfSense Plus".

pfSense CE software version numbers use the format `<major>.<minor>.<patch>`, and each component is present even if the value is `0`. This version numbering scheme is similar to the format used by FreeBSD software. In the past, this format was also used for releases of pfSense Factory Edition software before it was renamed to pfSense Plus.

3.65 Current and Upcoming Supported Releases

3.65.1 pfSense Plus Software

- *25.11*
- *25.07.1*
- *25.07*

3.65.2 pfSense CE Software

- *2.8.1*
- *2.8.0*

3.66 Older/Unsupported Releases

3.66.1 pfSense Plus Software

- *24.11*
- *24.03*
- *23.09.1*
- *23.09*
- *23.05.1*
- *23.05*
- *23.01*
- *22.05.1*
- *22.05*
- *22.01*
- *21.05.2*
- *21.05.1*
- *21.05*
- *21.02.2*
- *21.02-p1*
- *21.02*

3.66.2 pfSense CE Software

- *2.7.2*
- *2.7.1*
- *2.7.0*
- *2.6.0*
- *2.5.2*
- *2.5.1*
- *2.5.0*
- *2.4.5.p1*
- *2.4.5*
- *2.4.4.p3*
- *2.4.4.p2*
- *2.4.4.p1*
- *2.4.4*
- *2.4.3.p1*
- *2.3.5.p2*
- *2.4.3*
- *2.4.2.p1*
- *2.4.2*
- *2.4.1*
- *2.4*
- *2.3.5.p1*
- *2.3.5*
- *2.3.4.p1*
- *2.3.4*
- *2.3.3.p1*
- *2.3.3*
- *2.3.2.p1*
- *2.3.2*
- *2.3.1*
- *2.3*
- *2.2.6*
- *2.2.5*
- *2.2.4*
- *2.2.3*
- *2.2.2*

- *2.2.1*
- *2.2*
- *2.1.5*
- *2.1.4*
- *2.1.3*
- *2.1.2*
- *2.1.1*
- *2.1.0*
- *2.0.3*
- *2.0.2*
- *2.0.1*
- *2.0*

CHAPTER

FOUR

PRODUCT MANUALS

The pfSense Security Gateway Manuals help those who purchased appliances from Netgate get started with a new device running pfSense® software, or help get it back up and running in the case that something breaks.

Below is a list of active appliances:

- All Manuals
- Amazon AWS
- Microsoft Azure
- Netgate 1100
- Netgate 2100
- Netgate 4100
- Netgate 4200
- Netgate 6100
- Netgate 8200
- Netgate 8300
- Netgate 1537
- Netgate 1541

CHAPTER

FIVE

NETWORKING CONCEPTS

5.1 Understanding Public and Private IP Addresses

5.1.1 Private IP Addresses

The network standard RFC 1918 defines reserved IPv4 subnets for use only in private networks (Table *RFC 1918 Private IP Address Space*). RFC 4193 defines Unique Local Addresses (ULA) for IPv6 (Table *RFC 4193 Unique Local Address Space*). In most environments, a private IP subnet from RFC 1918 is chosen and used on all internal network devices. The devices are then connected to the Internet through a firewall or router implementing Network Address Translation (NAT) software, such as pfSense® software. IPv6 is fully routed from the internal network without NAT by Global Unicast Addresses (GUA). NAT will be explained further in *Network Address Translation*.

Table 1: RFC 1918 Private IP Address Space

CIDR Range	IP Address Range
10.0.0.0/8	10.0.0.0 - 10.255.255.255
172.16.0.0/12	172.16.0.0 - 172.31.255.255
192.168.0.0/16	192.168.0.0 - 192.168.255.255

Table 2: RFC 4193 Unique Local Address Space

Prefix	IP Address Range
fc00::/7	fc00:: - fdff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

A complete list of special-use IPv4 networks may be found in RFC 3330. There are private IPv4 addresses, such as `1.0.0.0/8` and `2.0.0.0/8`, that have since been allocated to the dwindling IPv4 pool. Use of these addresses are problematic and not recommended. Also, avoid using `169.254.0.0/16`, which according to RFC 3927 is reserved for "Link-Local" autoconfiguration. It should not be assigned by DHCP or set manually and routers will not allow packets from that subnet to traverse outside a specific broadcast domain. There is sufficient address space set aside by RFC 1918, so there is no need to deviate from the list shown in Table *RFC 1918 Private IP Address Space*. Improper addressing will result in network failure and should be corrected.

5.1.2 Public IP Addresses

Except for the largest networks, public IP addresses are assigned by Internet Service Providers. Networks requiring hundreds or thousands of public IP addresses commonly have address space assigned directly from their Regional Internet Registry (RIR). An RIR is an organization that oversees allocation and registration of public IP addresses in designated regions of the world.

Most residential Internet connections are assigned a single public IPv4 address. Most business class connections are assigned multiple public IP addresses. A single public IP address is adequate in many circumstances and can be used in conjunction with NAT to connect hundreds of privately addressed systems to the Internet. This documentation will assist in determining the number of public IP addresses required.

Most IPv6 deployments will give the end user at least a /64 prefix network to use as a routed internal network. For each site, this is roughly 2^{64} IPv6 addresses, or 18 *quintillion* addresses, fully routed from the Internet with no need for NAT.

5.1.3 Reserved and Documentation Addresses

In addition to blocks defined in RFC 1918, RFC 6890 describes blocks reserved for other special purposes such as documentation, testing, and benchmarking, including address space for Carrier-grade NAT allocated in RFC 6598. These special networks include:

Table 3: RFC 6890 IPv4 Reserved Address Space

CIDR Range	Purpose
192.0.2.0/24	Documentation and example code
198.51.100.0/24	Documentation and example code
203.0.113.0/24	Documentation and example code
198.18.0.0/25	Benchmarking network devices
169.254.0.0/16	Link Local
100.64.0.0/10	Carrier-grade NAT space
192.0.0.0/24	IETF Protocol Assignments
192.0.0.0/29	DS-Lite
192.88.99.0/24	6to4 Relay Anycast
240.0.0.0/4	Reserved

See also:

For a similar list for IPv6 prefixes, see *Special IPv6 Subnets*.

The documentation uses examples with addresses from the above documentation ranges as well as RFC 1918 networks since they are more familiar to users.

Some find these addresses tempting to use for VPNs or even local networks. Though the best practice is to only use them for their intended purposes, they are much less likely to be seen “in the wild” than RFC 1918 networks.

5.2 IP Subnet Concepts

When configuring TCP/IP settings on a device, a subnet mask (Or prefix length for IPv6) must be specified. This mask enables the device to determine which IP addresses are on the local network, and which must be reached by a gateway in the routing table. The default LAN IP address of `192.168.1.1` with a mask of `255.255.255.0`, or `/24` in CIDR notation has a network address of `192.168.1.0/24`. CIDR is discussed in *Understanding CIDR Subnet Mask Notation*.

5.3 IP Address, Subnet and Gateway Configuration

The TCP/IP configuration of a host consists of the address, subnet mask (or prefix length for IPv6) and gateway. A host identifies which IP addresses are on its local network by using the IP address combined with the subnet mask. A host sends packets for addresses outside the local network to the host's configured default gateway which it assumes will pass the traffic on to the desired destination. An exception to this rule is a static route which instructs a device to contact specific non-local subnets reachable via locally connected routers. This list of gateways and static routes is kept on the routing table of each host. To see the routing table used by pfSense® software, see *Route Table Contents*.

See also:

More information about routing can be found in *Routing*.

In a typical deployment of pfSense software hosts on the LAN are assigned an IP address, subnet mask and gateway within the LAN range of the firewall running pfSense software. The LAN IP address on the firewall becomes the default gateway for hosts on the LAN. For hosts connecting by an interface other than LAN, use the appropriate configuration for the interface to which the device is connected.

Hosts within a single network communicate directly with each other without involvement from the gateway. This means that no firewall, including one running pfSense software, can control host-to-host communication within a network segment. If this functionality is a requirement, hosts must be segmented via the use of multiple switches, VLANs, or employ equivalent switch functionality like PVLAN.

See also:

VLANs are covered in *Virtual LANs (VLANs)*.

5.4 Understanding CIDR Subnet Mask Notation

pfSense® software uses CIDR (Classless Inter-Domain Routing) notation rather than the common subnet mask `255.x.x.x` when configuring addresses and networks. Refer to the *CIDR Subnet Table* to find the CIDR equivalent of a decimal subnet mask.

Table 4: CIDR Subnet Table

Subnet Mask	CIDR Prefix	Total IP Addresses	Usable IP Addresses	Number of /24 networks
255.255.255.255	/32	1	1	1/256th
255.255.255.254	/31	2	2*	1/128th
255.255.255.252	/30	4	2	1/64th
255.255.255.248	/29	8	6	1/32nd
255.255.255.240	/28	16	14	1/16th
255.255.255.224	/27	32	30	1/8th
255.255.255.192	/26	64	62	1/4th
255.255.255.128	/25	128	126	1 half

continues on next page

Table 4 – continued from previous page

Subnet Mask	CIDR Prefix	Total IP Addresses	Usable IP Addresses	Number of /24 networks
255.255.255.0	/24	256	254	1
255.255.254.0	/23	512	510	2
255.255.252.0	/22	1024	1022	4
255.255.248.0	/21	2048	2046	8
255.255.240.0	/20	4096	4094	16
255.255.224.0	/19	8192	8190	32
255.255.192.0	/18	16,384	16,382	64
255.255.128.0	/17	32,768	32,766	128
255.255.0.0	/16	65,536	65,534	256
255.254.0.0	/15	131,072	131,070	512
255.252.0.0	/14	262,144	262,142	1024
255.248.0.0	/13	524,288	524,286	2048
255.240.0.0	/12	1,048,576	1,048,574	4096
255.224.0.0	/11	2,097,152	2,097,150	8192
255.192.0.0	/10	4,194,304	4,194,302	16,384
255.128.0.0	/9	8,388,608	8,388,606	32,768
255.0.0.0	/8	16,777,216	16,777,214	65,536
254.0.0.0	/7	33,554,432	33,554,430	131,072
252.0.0.0	/6	67,108,864	67,108,862	262,144
248.0.0.0	/5	134,217,728	134,217,726	524,288
240.0.0.0	/4	268,435,456	268,435,454	1,048,576
224.0.0.0	/3	536,870,912	536,870,910	2,097,152
192.0.0.0	/2	1,073,741,824	1,073,741,822	4,194,304
128.0.0.0	/1	2,147,483,648	2,147,483,646	8,388,608
0.0.0.0	/0	4,294,967,296	4,294,967,294	16,777,216

Note: The use of `/31` networks is a special case defined by RFC 3021 where the two IP addresses in the subnet are usable for point-to-point links to conserve IPv4 address space. Not all operating systems are compatible with RFC 3021, so use it with caution. On systems that are not compatible with RFC 3021, the subnet is unusable because the only two addresses defined by the subnet mask are the null route and broadcast and no usable host addresses.

pfSense software is capable of using `/31` networks for interfaces and Virtual IP addresses.

5.4.1 Where do CIDR numbers come from?

The CIDR number comes from the number of ones in the subnet mask when converted to binary.

The subnet mask `255.255.255.0` is `11111111.11111111.11111111.00000000` in binary. This adds up to `24` consecutive ones, or `/24` (pronounced "slash twenty-four").

A subnet mask of `255.255.255.192` is `11111111.11111111.11111111.11000000` in binary, or `26` ones, hence `/26`.

5.5 CIDR Summarization

In addition to specifying subnet masks, CIDR can also be employed for IP or network summarization purposes. The "Total IP Addresses" column in *CIDR Subnet Table* indicates how many addresses are summarized by a given CIDR mask. For network summarization purposes, the "Number of /24 networks" column is useful. CIDR summarization can be used in several parts of the pfSense® web interface, including firewall rules, NAT, virtual IPs, IPsec, and static routes.

IP addresses or networks that can be contained within a single CIDR mask are known as "CIDR summarizable".

When designing a network, ensure all private IP subnets in use at a particular location are CIDR summarizable. For example, if three `/24` subnets are required at one location, a `/22` network subnetted into four `/24` networks should be used. The following table shows the four `/24` subnets used with the subnet `10.70.64.0/22`.

Table 5: CIDR Route Summarization

10.70.64.0/22 split into /24 networks
10.70.64.0/24
10.70.65.0/24
10.70.66.0/24
10.70.67.0/24

This keeps routing more manageable for multi-site networks connected to another physical location via the use of a private WAN circuit or VPN. With CIDR summarizable subnets, one route destination covers all the networks at each location. Without it, there are several different destination networks per location.

The previous table was developed using a network calculator found at the subnetmask.info website.

The calculator converts from dotted decimal to CIDR mask, and vice versa, as shown in Figure *Subnet Mask Converter*. If the *CIDR Subnet Table* provided in this chapter is not available, this tool can be used to convert a CIDR prefix to dotted decimal notation. Enter a CIDR prefix or a dotted decimal mask and click the appropriate **Calculate** button to find the conversion.



Fig. 1: Subnet Mask Converter

Enter the dotted decimal mask into the Network/Node Calculator section along with one of the `/24` networks. Click **Calculate** to populate the bottom boxes with the range covered by that particular `/24` as demonstrated in Figure *Network/Node Calculator*. In this example, the network address is `10.70.64.0/22`, and the usable `/24` networks are `64` through `67`. The term "Broadcast address" in this table refers the highest address within the range.



Fig. 2: Network/Node Calculator

5.5.1 Finding a matching CIDR network

IPv4 Ranges in the format of `x.x.x.x-y.y.y.y` can be used in Aliases. For Network type aliases, an IPv4 range is automatically converted to the equivalent set of CIDR blocks. For Host type aliases, a range is converted to a list of IPv4 addresses. See *Aliases* for more information.

If an exact match isn't necessary, numbers can be entered into the Network/Node Calculator to approximate the desired summarization.

5.6 Broadcast Domains

A broadcast domain is the portion of a network sharing the same layer 2 segment. Broadcast messages from hosts are sent to every port in their broadcast domain, thus hosts inside a broadcast domain can reach each other directly. For example hosts can use ARP or NDP to locate neighbors within a broadcast domain and communicate directly at layer 2 without involving an intermediate gateway router.

In a network with a single switch without VLANs, the broadcast domain is that entire switch. In a network with multiple interconnected switches without the use of VLANs, the broadcast domain includes all of those switches. When using VLANs, each VLAN is typically its own broadcast domain. The exact size of the broadcast domain in that case varies depending on how many access ports are in the VLAN, along with interconnected switches (trunked, stacked, etc.).

Some switches have special modes which segment a broadcast domain into multiple smaller isolated broadcast domains. This is sometimes called "Private VLANs", and they are typically used for security purposes. In these modes, hosts can only directly communicate between a specific set of ports, commonly limited to the host and the gateway for the segment, even if they are a part of a subnet with many other hosts. This is similar in concept to wireless AP client isolation.

Since broadcast messages are sent to every port in the broadcast domain, large broadcast domains should be avoided as they are "noisy" and do not scale well. Depending on the type of broadcast messages, some switches can optimize this behavior, but it's best to plan for the worst case. For example in a network with thousands of ports on a single broadcast domain, thousands of hosts communicating among each other generate large amounts of broadcast traffic which is copied everywhere in the broadcast domain. The best practice is to keep each segment as small as possible, where feasible, to prevent switches and hosts from having to process large amounts of unnecessary broadcast traffic.

A single broadcast domain *can* contain more than one IPv4 or IPv6 subnet, however, that is generally not considered good network design. Though it appears on the surface that multiple subnets in the same broadcast domain are separate, there is no true isolation or security between them. IP subnets should be segregated into different broadcast domains via the use of separate switches or VLANs. The exception to this is running both IPv4 and IPv6 networks within a single broadcast domain. This is called dual stack, and it is a common and useful technique using both IPv4 and IPv6 connectivity for hosts.

Broadcast domains can be combined by *bridging* two network interfaces together. In this scenario care must be taken

to avoid switch loops where a switch ends up with a connection back to itself, creating an infinite traffic loop (*Bridging and Layer 2 Loops*). Another reason to avoid bridging is that by combining broadcast domains, both networks and the bridge between them must carry broadcast traffic for every network on the bridge. The increased load, especially for larger networks, can be significant, especially if broadcast domains are being bridged using a VPN. There are also proxies for certain protocols which do not combine broadcast domains but yield the same net effect, such as a DHCP relay which relays DHCP requests into a broadcast domain on another interface.

See also:

- *Bridging*
- *Bridging and Layer 2 Loops*
- *Virtual LANs (VLANs)*
- Broadcast Domain (Wikipedia)

5.7 IPv6

5.7.1 Basics

IPv6 allows for exponentially more IP address space than IPv4. IPv4 uses a 32-bit address, which allows for 2^{32} or over 4 billion addresses, less if the sizable reserved blocks and IPs burned by subnetting are removed. IPv6 uses a 128-bit address, which is 2^{128} or 3.403 x 10^{38} IP addresses. The standard size IPv6 subnet defined by the IETF is a /64, which contains 2^{64} IPs, or 18.4 *quintillion* addresses. The entire IPv4 space can fit inside a typical IPv6 subnet many times over with room to spare.

One of the more subtle improvements with IPv6 is that no IP addresses are lost to subnetting. With IPv4, two IP addresses are lost per subnet to account for a null route and broadcast IP address. In IPv6, broadcast is handled via the same mechanisms used for multicast involving special addresses sent to the entire network segment. Additional improvements include integrated packet encryption, larger potential packet sizes, and other design elements that make it easier for routers to manage IPv6 at the packet level.

Unlike IPv4, all packets are routed in IPv6 without NAT. Each IP address is directly accessible by another unless stopped by a firewall. This can be a very difficult concept to grasp for people who are used to having their LAN exist with a specific private subnet and then performing NAT to whatever the external address happens to be.

There are fundamental differences in the operation of IPv6 in comparison to IPv4, but mostly they are only that: differences. Some things are simpler than IPv4, others are slightly more complicated, but for the most part it's simply different. Major differences occur at layer 2 (ARP vs. NDP for instance) and layer 3 (IPv4 vs. IPv6 addressing). The protocols used at higher layers are identical; only the transport mechanism for those protocols has changed. HTTP is still HTTP, SMTP is still SMTP, etc.

Firewall and VPN Concerns

IPv6 restores true peer-to-peer connectivity originally in place with IPv4 making proper firewall controls even more important. In IPv4, NAT was misused as an additional firewall control. In IPv6, NAT is removed. Port forwards are no longer required in IPv6 so remote access will be handled by firewall rules. Care must be taken to ensure encrypted VPN LAN to LAN traffic is not routed directly to the remote site. See *IPv6 VPN and Firewall Rules* for a more in-depth discussion on IPv6 firewall concerns with respect to VPN traffic.

5.7.2 Requirements

IPv6 requires an IPv6-enabled network. IPv6 connectivity delivered directly by an ISP is ideal. Some ISPs deploy a dual stack configuration in which IPv4 and IPv6 are delivered simultaneously on the same transport. Other ISPs use tunneling or deployment types to provide IPv6 indirectly. It is also possible to use a third party provider such as Hurricane Electric's tunnelbroker service.

In addition to the service, software must also be capable of using IPv6. pfSense® software has been IPv6-capable since 2.1-RELEASE. Client operating systems and applications must also be capable of using IPv6, but this is common in most current operating systems and applications. To ensure reliability, it is always beneficial to employ the latest operating system and software updates.

Some mobile operating systems have varying levels of IPv6 capabilities. For example, Android cannot obtain a dynamic address using DHCPv6, but it can use stateless autoconfiguration (SLAAC). IPv6 is part of LTE and newer wireless specifications, so any current mobile devices can utilize IPv6 as well if the carrier provides IPv6 connectivity.

5.7.3 IPv6 WAN Types

Details can be found in *IPv6 Configuration Types*, but some of the most common ways of deploying IPv6 are:

Static Addressing
: Native and using IPv6 either on its own or in a dual stack configuration alongside IPv4.

DHCPv6
: Address automatically obtained by DHCPv6 to an upstream server. Prefix delegation may also be used with DHCPv6 to deliver a routed subnet to a DHCPv6 client.

Stateless address autoconfiguration (SLAAC)
: Automatically determines the IPv6 address by consulting router advertisement messages and generating an IP address inside a prefix. This is not very useful for a router, as there is no way to route a network for the "inside" of the firewall. It may be useful for appliance modes.

6RD Tunnel
: A method of tunneling IPv6 traffic inside IPv4. This is used by ISPs for rapid IPv6 deployment.

6to4 Tunnel
: Similar to 6RD but with different mechanisms and limitations.

GIF Tunnel
: Not technically a direct WAN type, but commonly used. Customer builds an IPv4 GIF tunnel to a provider to tunnel IPv6 traffic.

While not technically a WAN type, IPv6 connectivity can also be arranged over a VPN such as IPsec, WireGuard, or OpenVPN. Most VPNs are capable of carrying IPv4 and IPv6 traffic simultaneously, so they can deliver IPv6 over IPv4, though with more overhead than a typical tunnel broker that uses GIF. These are good options for a company that has IPv6 at a datacenter or main office but not at a remote location.

5.7.4 Address Format

An IPv6 address consists of 32 hexadecimal digits, in 8 sections of 4 digits each, separated by colons. It looks something like this: `1234:5678:90ab:cdef:1234:5678:90ab:cdef`

IPv6 addresses have several shortcuts that allow them to be compressed into smaller strings following certain rules.

If there are any leading zeroes in a section, they may be left off. `0001:0001:0001:0001:0001:0001:0001:0001` could be written as `1:1:1:1:1:1:1:1`.

Any number of address parts consisting of only zeroes may be compressed by using `::`, but this can only be done once in an IPv6 address to avoid ambiguity. A good example of this is local host, compressing `0000:0000:0000:0000:0000:0000:0000:0001` to `::1`. Any time `::` appears in an IPv6 address, the values between are all zeroes. An IP address such as `fe80:1111:2222:0000:0000:0000:7777:8888`, can be represented as `fe80:1111:2222::7777:8888`. However, `fe80:1111:0000:0000:4444:0000:0000:8888` cannot be shortened using `::` more than once. It would either be `fe80:1111::4444:0:0:8888` or `fe80:1111:0:0:4444::8888`, but it *cannot* be `fe80:1111::4444::8888` because there is no way to tell how many zeroes have been replaced by either `::` operator.

Determining an IPv6 Addressing Scheme

Because of the increased length of the addresses, the vast space provided in even a basic /64 size prefix, and the ability to use hexadecimal digits, there is more freedom to design device network addresses.

On servers using multiple IP address aliases for virtual hosts, jails, etc., a useful addressing scheme is to use the seventh section of the IPv6 address to denote the server. Then use the eighth section for individual IPv6 aliases. This groups all the IP addresses into a single recognizable host. For example, the server itself would be `2001:db8:1:1::a:1`, and then the first IP alias would be `2001:db8:1:1::a:2`, then `2001:db8:1:1::a:3`, etc. The next server would be `2001:db8:1:1::b:1`, and repeats the same pattern.

Some administrators like to have fun with their IPv6 addresses by using hexadecimal letters and number/letter equivalents to make words out of their IP addresses. Lists of hexadecimal words around the web can be used to create more memorable IP addresses such as `2001:db8:1:1::dead:beef`.

Decimal vs. Hexadecimal Confusion

Creating consecutive IPv6 addresses with a hexadecimal base may cause confusion. Hexadecimal values are base 16 unlike decimal values which are base 10. For example, the IPv6 address `2001:db8:1:1::9` is followed by `2001:db8:1:1::a`, *not* `2001:db8:1:1::10`. By going right to `2001:db8:1:1::10`, the values `a-f` have been skipped, leaving a gap. Consecutive numbering schemes are not required, and their use is left to the discretion of the network designer. For some, it is psychologically easier to avoid using the hexadecimal digits.

Given that all IPv4 addresses can be expressed in IPv6 format, this issue will arise when designing a dual stack network that keeps one section of the IPv6 address the same as its IPv4 counterpart.

5.7.5 IPv6 Subnets

Creating subnets with IPv6 is easier than IPv4. It's also different. Want to divide or combine a subnet? All that is needed is to add or chop off digits and adjust the prefix length by a multiple of four. No longer is there a need to calculate subnet start/end addresses, usable addresses, the null route, or the broadcast address.

IPv4 had a subnet mask (dotted quad notation) that was later replaced by CIDR masking. IPv6 doesn't have a subnet mask but instead calls it a Prefix Length, often shortened to "Prefix". Prefix length and CIDR masking work similarly; The prefix length denotes how many bits of the address define the network in which it exists. Most commonly the prefixes used with IPv6 are multiples of four, as seen in Table *IPv6 Subnet Table*, but they can be any number between `0` and `128`.

Using prefix lengths in multiples of four makes it easier for humans to distinguish IPv6 subnets. All that is required to design a larger or smaller subnet is to adjust the prefix by multiple of four. For reference, see Table *IPv6 Subnet Table* listing the possible IPv6 addresses, as well as how many IP addresses are contained inside each subnet.

Table 6: IPv6 Subnet Table

Prefix	Subnet Example	Total IP Addresses	# of /64 nets
4	x::	2^{124}	2^{60}
8	xx::	2^{120}	2^{56}
12	xxx::	2^{116}	2^{52}
16	xxxx::	2^{112}	2^{48}
20	xxxx:x::	2^{108}	2^{44}
24	xxxx:xx::	2^{104}	2^{40}
28	xxxx:xxx::	2^{100}	2^{36}
32	xxxx:xxxx::	2^{96}	4,294,967,296
36	xxxx:xxxx:x::	2^{92}	268,435,456
40	xxxx:xxxx:xx::	2^{88}	16,777,216
44	xxxx:xxxx:xxx::	2^{84}	1,048,576
48	xxxx:xxxx:xxxx::	2^{80}	65,536
52	xxxx:xxxx:xxxx:x::	2^{76}	4,096
56	xxxx:xxxx:xxxx:xx::	2^{72}	256
60	xxxx:xxxx:xxxx:xxx::	2^{68}	16
64	xxxx:xxxx:xxxx:xxxx::	2^{64} (18,446,744,073,709,551,616)	1
68	xxxx:xxxx:xxxx:xxxx:x::	2^{60} (1,152,921,504,606,846,976)	0
72	xxxx:xxxx:xxxx:xxxx:xx::	2^{56} (72,057,594,037,927,936)	0
76	xxxx:xxxx:xxxx:xxxx:xxx::	2^{52} (4,503,599,627,370,496)	0
80	xxxx:xxxx:xxxx:xxxx:xxxx::	2^{48} (281,474,976,710,656)	0
84	xxxx:xxxx:xxxx:xxxx:xxxx:x::	2^{44} (17,592,186,044,416)	0
88	xxxx:xxxx:xxxx:xxxx:xxxx:xx::	2^{40} (1,099,511,627,776)	0
92	xxxx:xxxx:xxxx:xxxx:xxxx:xxx::	2^{36} (68,719,476,736)	0
96	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx::	2^{32} (4,294,967,296)	0
100	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:x::	2^{28} (268,435,456)	0
104	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xx::	2^{24} (16,777,216)	0
108	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxx::	2^{20} (1,048,576)	0
112	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx::	2^{16} (65,536)	0
116	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:x::	2^{12} (4,096)	0
120	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xx::	2^{8} (256)	0
124	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxx::	2^{4} (16)	0
128	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx	2^{0} (1)	0

A `/64` is a standard size IPv6 subnet as defined by the IETF. It is the smallest subnet that can used locally if autoconfiguration is desired.

Typically, an ISP assigns a `/64` or smaller subnet to establish service on the WAN. An additional network is routed for LAN use. The size of the allocation depends upon the ISP, but it's not uncommon to see end users receive at least a `/64` and even up to a `/48`.

A tunnel service provider such as tunnelbroker.net run by Hurricane Electric will allocate a `/48` in addition to a routed `/64` subnet and a `/64` interconnect.

Assignments larger than `/64` usually adopt the first `/64` for LAN and subdivide the rest for requirements such as VPN tunnel, DMZ, or a guest network.

Special IPv6 Subnets

Special use networks are reserved in IPv6. A full list of these can be found in the Wikipedia IPv6 article. Six examples of IPv6 special networks and their addresses are shown below in *IPv6 Special Networks and Addresses*.

Table 7: IPv6 Special Networks and Addresses

Network	Purpose
::1	Localhost
2001:db8::/32	Documentation prefix used for examples (RFC 3849)
3fff::/20	Documentation prefix used for examples (RFC 9637)
64:ff9b::/96	*NAT64*
fc00::/7	Unique Local Addresses (ULA) - also known as "Private" IPv6 addresses.
fe80::/10	Link Local addresses, only valid inside a single broadcast domain.
ff00::0/8	Multicast addresses
::ffff:0:0/96	IPv4 Mapped Addresses

Neighbor Discovery

IPv4 hosts find each other on a local segment using ARP broadcast messages, but IPv6 hosts find each other by sending Neighbor Discovery Protocol (NDP) messages. Like ARP, NDP works inside a given broadcast domain to find other hosts inside a specific subnet.

By sending special ICMPv6 packets to reserved multicast addresses, NDP handles the tasks of neighbor discovery, router solicitations, and route redirects similar to IPv4's ICMP redirects.

pfSense® software automatically adds firewall rules on IPv6 enabled interfaces that permit NDP to function. All current known neighbors on IPv6 can be viewed in the firewall GUI at **Diagnostics > NDP Table**.

Router Advertisements

IPv6 routers are located through their Router Advertisement (RA) messages instead of by DHCP. IPv6-enabled routers that provide dynamic address assignment are expected to announce themselves on the network to all clients and respond to router solicitations. When acting as a client (WAN interfaces), pfSense software accepts RA messages from upstream routers. When acting as a router, pfSense software provides RA messages to clients on its internal networks. See *Router Advertisements (Or: "Where is the DHCPv6 gateway option?")* for more details.

Address Allocation

Client addresses can be allocated by static addressing through SLAAC (*Router Advertisements (Or: "Where is the DHCPv6 gateway option?")*), DHCP6 (*IPv6 Router Advertisements*), or other tunneling methods such as OpenVPN.

DHCP6 Prefix Delegation

DHCP6 Prefix Delegation delivers a routed IPv6 subnet to a DHCP6 client. A WAN type interface can be set to receive a prefix over DHCP6 (*DHCP6*, *Track Interface*). A router functioning at the edge of a large network can provide prefix delegation to other routers inside the network (*DHCPv6 Prefix Delegation*).

5.7.6 IPv6 and NAT

Though IPv6 removes almost any need for NAT, there are rare situations that call for the use of NAT with IPv6 such as Multi-WAN for IPv6 on residential or small business networks.

IPv6 all but eliminates the need for traditional port translated NAT (PAT) where internal addresses are translated using ports on a single external IP address.

Outbound NAT

While it is possible to perform *Outbound NAT* on IPv6 traffic, the best practice is to allow IPv6 traffic to pass without performing any address or port translation.

Prefix Translation (1:1 NAT)

It is possible to translate one IPv6 prefix to another, which is Network Prefix Translation (NPt). This is available in the pfSense® software GUI under **Firewall > NAT** on the **NPt** tab. For example, NPt can translate `2001:db8:1111:2222::/64` to `2001:db8:3333:4444::/64` while maintaining the host portion of the address. For more on NPt, see *IPv6 Network Prefix Translation (NPt)*.

NAT64

pfSense software is capable of performing NAT64 which is useful as a transition mechanism to allow IPv6-only hosts access to IPv4-only resources. Several different functions must be configured for a full NAT64 configuration. See *NAT64* and *Configuring NAT64 for IPv6-only Clients* for details.

IPv4 Mapped Addresses

There is also a mechanism built into IPv6 to access IPv4 hosts using a special address notation, such as `::ffff:192.168.1.1`. The behavior of these addresses can vary between OS and application and can be unreliable.

5.7.7 NAT64

NAT64 is a form of *NAT* which enables clients with only IPv6 addresses to reach remote hosts using IPv4 addresses. NAT64 accomplishes this by mapping IPv4 addresses into a special IPv6 prefix dedicated to this purpose, such as the default NAT64 prefix, `64:ff9b::/96`.

Note: Though NAT64-related settings allow this prefix to be set to a custom value, in practice the value rarely if ever deviates from the default prefix of `64:ff9b::/96`.

See also:

- *NAT64 Configuration Recipe*

Requirements

To perform NAT64 there are a few prerequisites:

- The device performing NAT64 must have an external IPv4 address.
- If using this firewall for DNS, it must use the DNS Resolver.
- The internal interface with IPv6 clients **does not** need to have IPv4 configured.
- The local IPv6 interface must be properly configured, have appropriate firewall rules, etc. This can be a tracked WAN configuration or manually configured.

NAT64 Components

Several components come together to allow a fully-functioning NAT64 environment:

- *NAT64 Firewall Rules on the local interface*
- *PREF64 in Router Advertisements*
- *DNS64 in the DNS Resolver*
- Optionally configure DHCPv4 option `108` / `v6-only-preferred` if the segment has local/private IPv4.
- Configure DHCPv6 to at least advertise DNS servers to clients as not all clients are compatible with RDNSS/DNSSL and NAT64 requires working DNS.

See also:

See *NAT64 Configuration Recipe* for a complete walk-through of configuring each component in NAT64.

How NAT64 Works

This is a basic summary of the process for an IPv6 client to successfully make NAT64 requests:

- IPv6 client discovers the NAT64 prefix via *PREF64*.
- IPv6 client performs a DNS lookup for a host.
- The DNS Resolver uses *DNS64* to include mapped IPv4 addresses in its responses using the configured NAT64 prefix.

  For example, the client requests `www.example.com` which resolves to `203.0.113.5` and has no IPv6 address. The DNS response to the client will include both `203.0.113.5` and the IPv6 mapped equivalent, `64:ff9b::cb00:7105`.
- IPv6 client contacts the IPv4 host using the mapped address returned via DNS64.
- The firewall translates this request using NAT64 so the incoming IPv4 packet has both its source and destination translated to appropriate IPv4 addresses. This is a stateful mapping which allows the firewall to appropriately handle return packets for ongoing connections.

This entire process is transparent to the client software, such as a web browser. It will act as though it is communicating directly to IPv6 hosts even when using hostnames for sites with no IPv6 connectivity.

Contacting Remote IPv4 Hosts

IPv6 clients can contact IPv4 hosts using mapped addresses without DNS64 by manually appending the IPv4 address to the NAT64 prefix. For example, to contact the IPv4 address `203.0.113.5` the client would instead contact `64:ff9b::203.0.113.5`. The client could also use the fully-IPv6 equivalent address, `64:ff9b::cb00:7105`, but that is much more difficult for humans.

NAT64 **does not** enable IPv6 clients to talk to IPv4 hosts directly using IPv4 address notation. However, IPv6 clients may be able to do so if they have a customer-side translator (CLAT). This is completely dependent on the client, routers/firewalls are not involved in that functionality. The availability of a CLAT is up to the operating system and its enabled features, installed software, etc.

Note: NAT64 on pfSense software only allows IPv6 clients to contact IPv4 servers and exchange two-way traffic. NAT64 **does not** allow IPv4 remote hosts to reach local IPv6 hosts.

NAT64 and Policy Routing

NAT64 is compatible with policy routing. For example, if an IPv4 destination has a static route out through an alternate path, a NAT64 rule can be configured to match. The NAT64 rule should match the equivalent mapped destination with a gateway configured to ensure it takes the expected path.

NAT64 and other NAT

NAT64 translates traffic when it reaches a rule **inbound** on an interface. This happens **before** the firewall processes Outbound NAT rules. As a result, while the packets *can* be translated again by Outbound NAT, the packets would not likely match in a way that makes performing additional Outbound NAT practical as the packet would already have an external address and may not be distinguishable from other connections.

5.7.8 IPv6 and pfSense Software

Unless noted otherwise, it is safe to assume that a given area or feature of pfSense® software is capable of using IPv6.

Some noteworthy areas of pfSense software that *cannot* use IPv6 are Captive Portal and most DynDNS providers.

Note: On installations upgraded from versions of pfSense software prior to 2.1, the firewall blocks IPv6 traffic by default.

To allow IPv6:

- Navigate to **System > Advanced** on the **Networking** tab
- Check **Allow IPv6**
- Click **Save**

Packages

Some packages are maintained by the community, so IPv6 capabilities vary. In most cases, IPv6 availability depends upon the capabilities of the underlying software. It is safer to assume a package does not work with IPv6 unless otherwise noted. Packages are updated periodically, the best practice is to test a package to determine its IPv6 capabilities.

5.7.9 Controlling IPv6 Preference for traffic from the firewall itself

By default, pfSense® software prefers IPv6 when possible. If IPv6 routing is not functional, but the system believes it is, pfSense software may fail to check updates or download packages properly.

To change this behavior, pfSense software provides a method in the GUI to control whether services on the firewall prefer IPv4 over IPv6:

- Navigate to **System > Advanced** on the **Networking** tab
- Check **Prefer to use IPv4 even if IPv6 is available**
- Click **Save**

Once the settings have been saved, the firewall itself will prefer IPv4 for outbound communication.

See also:

- *Configuring IPv6 Through A Tunnel Broker Service*

Around the world, the availability of new IPv4 addresses is declining. The amount of free space varies by region, but some have already run out of allocations and others are rapidly approaching their limits. As of January 31, 2011, IANA allocated all of its space to regional internet registries (RIRs). In turn, these RIR allocations have run out in some locations such as APNIC (Asia/Pacific), RIPE (Europe), and LACNIC (Latin America and Caribbean) for /8 networks. Though some smaller allocations are still available, it is increasingly difficult to obtain new IPv4 address space in these regions. ARIN (North America) ran out on September 24th, 2015.

To account for this, IPv6 was created as a replacement for IPv4. Available in some forms since the 1990s, factors like inertia, complexity, and the cost of developing or purchasing compatible routers and software has slowed its uptake until the last few years. Even then, it's been rather slow with only 8% of Google users having IPv6 connectivity by July 2015 and slightly over 40% of users in 2022.

Over the years, IPv6 implementations in software, operating systems, and routers has improved, so the situation is primed to get better. Still it is up to ISPs to start delivering IPv6 connectivity to users. It's a catch-22 situation: Content providers are slow to provide IPv6 because few users have it. Meanwhile, users don't have it because there isn't a lot of IPv6 content and even less content available only over IPv6. Users don't know they need it so they don't demand the service from their ISPs.

Some providers are experimenting with *Carrier Grade NAT (CGN)* to stretch their IPv4 networks farther. CGN places their IPv4 residential customers behind another layer of NAT further breaking protocols that already don't deal with one layer of NAT. Mobile data providers have been doing this for some time, but the applications typically found on mobile devices aren't affected since they work as if they're behind a typical SOHO router style NAT. While solving one problem, it creates others as observed when CGN is used as a firewall's WAN, when tethering on a PC, or in some cases attempting to use a traditional IPsec VPN without NAT-T, or PPTP. ISPs employing CGN should be used only if there is no other choice.

There are many books and websites available with volumes of in-depth information on IPv6. The Wikipedia article on IPv6 is a great resource for additional information and links to other sources. It's worth using as a starting point for more information on IPv6. There are also many good books on IPv6 available, but be careful to purchase books with recent revisions. There have been changes to the IPv6 specification over the years, and it's possible that the material could have changed since the book's printing.

See also:

Hangouts Archive to view the July 2015 Hangout on IPv6 Basics

This documentation is not an introduction to networks, but there are certain networking concepts that need to be addressed.

Note: Readers without basic fundamental networking knowledge should locate additional introductory material as this chapter will not adequately provide all necessary information.

IPv6 concepts are introduced later in *IPv6*. For clarity, traditional IP addresses are referred to as IPv4 addresses. Except where otherwise noted, most functions will work with either IPv4 or IPv6 addresses. The general term IP address refers to either IPv4 or IPv6.

5.8 Brief introduction to OSI Model Layers

The OSI model has a network framework consisting of seven layers. These layers are listed in hierarchy from lowest to highest. A brief overview of each level is outlined below. More information can be found in many networking texts and on Wikipedia (http://en.wikipedia.org/wiki/OSI_model).

Layer 1 - Physical
Refers to either electrical or optical cabling that transports raw data to all the higher layers.

Layer 2 - Data Link
Typically refers to Ethernet or another similar protocol that is being spoken on the wire. This documentation often refers to layer 2 as meaning the Ethernet switches or other related topics such as ARP and MAC addresses.

Layer 3 - Network Layer
The protocols used to move data along a path from one host to another, such as IPv4, IPv6, routing, subnets etc.

Layer 4 - Transport Layer
Data transfer between users, typically refers to TCP or UDP or other similar protocols.

Layer 5 - Session Layer
Manages connections and sessions (typically referred to as “dialogs”) between users, and how they connect and disconnect gracefully.

Layer 6 - Presentation Layer
Handles any conversions between data formats required by users such as different character sets, encodings, compression, encryption, etc.

Layer 7 - Application Layer
Interacts with the user or software application, includes familiar protocols such as HTTP, SMTP, SIP, etc.

CHAPTER

SIX

HARDWARE

6.1 Minimum Hardware Requirements

The minimum hardware requirements for pfSense® software on hardware not sold by Netgate are:

- 64-bit amd64 (x86-64) compatible CPU
- 1GB or more RAM
- 8 GB or larger disk drive (SSD, HDD, etc.)
- One or more compatible network interface cards
- Bootable USB drive or high capacity optical drive (DVD or BD) for initial installation

Note: The minimum requirements are not suitable for all environments; see *Hardware Sizing Guidance* for details.

6.2 Hardware Selection

The use of open source operating systems with untested hardware may create hardware/software conflicts. *Hardware Tuning and Troubleshooting* offers tips on resolving various issues.

6.2.1 Preventing hardware headaches

Use Genuine Netgate Hardware

The best practice is to use hardware from the Netgate Store. Netgate hardware has been developed to assure that specific hardware platforms have been thoroughly tested and validated.

Search for the experiences of others

The experiences of others are a valuable source of knowledge which can be found by researching pfSense software and hardware compatibility online, especially on the Netgate Forum. Reports of failure are not necessarily considered definitive because problems can arise from a number of issues other than hardware incompatibility.

If the hardware in question is from a major manufacturer, an internet search by make, model, and `site:netgate.com` will search the Netgate website for relevant user experiences. Searching for the make, model, and `pfSense` will find user experiences on other websites. Repeating the same search with `FreeBSD` instead of `pfSense` can also turn up useful experiences.

6.2.2 Naming Conventions

This documentation refers to the 64-bit hardware architecture as `amd64`, the architecture designation used by FreeBSD. Intel adopted the architecture created by AMD for x86-64, thus the name amd64 refers to all x86 64-bit CPUs.

Netgate sells ARM appliances compatible with the Plus edition of pfSense software. This hardware is based on the `armv6` and `armv7` architectures (also called `arm`) and `aarch64` (also called `arm64`). Items specific to those unique architectures will be called out as necessary. The generic term ARM may be considered to apply to all of these, but **only** for the specific ARM-based appliances sold by Netgate, such as the 2100 and 3100.

6.3 Hardware Sizing Guidance

When sizing hardware for pfSense® software, required throughput and necessary features are the primary factors that govern hardware selection.

The information on Netgate Store now contains up-to-date specifications and performance data on all hardware sold by Netgate. The data on the Netgate Store is updated as needed, and it is always the most accurate and current source of performance data.

Tip: Contact Netgate Sales for personalized help in selecting the most suitable model for any implementation.

Estimating throughput of third party / white box hardware is difficult and inaccurate. In some cases, ballpark estimates may be made by comparing hardware specifications with those found on the Netgate Store for comparable models.

6.3.1 Throughput Considerations

In real networks the traffic flow will likely contain packets of varying size, not all maximum size packets, but it completely depends on the environment and the type of traffic involved. IMIX testing attempts to approximate a mixture of traffic that more closely resembles real-world environments. Simple IMIX traffic is sets of 7 (40) byte packets, (4) 576 byte packets, 1 (1500) byte packets, plus Ethernet framing overhead.

Note: The Netgate Store entries for hardware include data for maximum size packets ("IPERF3") and results for IMIX traffic patterns.

As a general reference, table *500,000 PPS Throughput at Various Frame Sizes* lists a few common packet sizes and the throughput achieved at an example rate of 500,000 packets per second.

Table 1: 500,000 PPS Throughput at Various Frame Sizes

Frame size	Throughput at 500 Kpps
64 bytes	244 Mbps
500 bytes	1.87 Gbps
1000 bytes	3.73 Gbps
1500 bytes	5.59 Gbps

Performance difference by network adapter type

The choice of NIC has a significant impact on performance. Inexpensive, low-end cards consume significantly more CPU than better quality cards such as Intel. The first bottleneck with firewall throughput is the CPU. Throughput improves significantly by using a better quality NIC with slower CPUs. By contrast, increasing the speed of the CPU will not proportionally increase the throughput when coupled with a low quality NIC.

6.3.2 Feature Considerations

Features, services and packages enabled on the firewall can lower the total potential throughput as they consume hardware resources that could otherwise be used to transfer network traffic. This is especially true for packages that intercept or inspect network traffic, such as Snort or Suricata.

Most base system features do not significantly factor into hardware sizing, but a few can potentially have a considerable impact on hardware utilization.

Large State Tables

Active network connections through the firewall are tracked in the firewall state table. Each connection through the firewall consumes two states: One entering the firewall and one leaving the firewall. For example, if a firewall must handle 100,000 simultaneous web server client connections the state table must be able to hold 200,000 states.

See also:

States are covered further in *Firewall*.

Firewalls in environments which require large numbers of simultaneous states must have sufficient RAM to contain the state table. Each state takes approximately 1 KB of RAM, which makes calculating the memory requirements relatively easy. Table *Large State Table RAM Consumption* provides a guideline for the amount of memory required for larger state table sizes. This is solely the memory used for the state tracking. The operating system itself along with other services will require at least 175-256 MB additional RAM and possibly more depending on the features used.

Table 2: Large State Table RAM Consumption

States	Connections	RAM Required
100,000	50,000	~97 MB
500,000	250,000	~488 MB
1,000,000	500,000	~976 MB
3,000,000	1,500,000	~2900 MB
8,000,000	4,000,000	~7800 MB

It is safer to overestimate the requirements. Based on the information above, a good estimate would be that 100,000 states consume about 100 MB of RAM, or that 1,000,000 states would consume about 1 GB of RAM.

VPN (all types)

The question customers typically ask about VPNs is "How many connections can my hardware handle?" That is a secondary factor in most deployments and is of lesser consideration. That metric is a relic of how other vendors have licensed VPN capabilities in the past and has no specific direct equivalent in pfSense software. The primary consideration in hardware sizing for VPN is the potential throughput of VPN traffic.

Encrypting and decrypting network traffic with all types of VPNs is CPU intensive. pfSense software offers several cipher options for use with IPsec. The various ciphers perform differently and the maximum throughput of a firewall is dependent on the cipher used and whether that cipher can be accelerated by the hardware.

See also:

The Netgate Store contains VPN performance data for each device sold by Netgate using the most optimal cipher for each device based on its capabilities.

Hardware *cryptographic accelerators*, such as those found on most Netgate hardware, greatly increase maximum VPN throughput and largely eliminate the performance difference between accelerated ciphers. For IPsec, ciphers may be accelerated by onboard cryptographic accelerators. For example, AES-GCM is accelerated by AES-NI. IPsec is faster not only due to that, but because it also does not require a separate authentication algorithm. IPsec also has less per-packet operating system processing overhead than OpenVPN, so for the time being IPsec will nearly always be faster than OpenVPN.

Where high VPN throughput is a requirement for a firewall, hardware cryptographic acceleration is of utmost importance to ensure not only fast transmission speeds but also reduced CPU overhead. The reduction in CPU overhead means the VPN will not lower the performance of other services on the firewall.

The current best available acceleration is available by using pfSense Plus software on hardware with a QAT device, followed by a CPU which includes instructions for IPsec-MB (SSE, AVX2, AVX512), or failing that, a CPU which includes AES-NI instructions combined with AES-GCM in IPsec.

Packages

Certain packages have a significant impact on hardware requirements, and their use must be taken into consideration when selecting hardware.

Snort/Suricata

Snort and Suricata are pfSense software packages for network intrusion detection. Depending on their configuration, they can require a significant amount of RAM. 1 GB should be considered a minimum, but some configurations may need 2 GB or more, not counting RAM used by the operating system, firewall states, and other packages.

Suricata is multithreaded and can potentially take advantage of NETMAP for inline IPS on compatible hardware.

6.4 Memory Management

FreeBSD manages memory by dividing it into multiple areas or "pools" and moving memory pages between these areas as needed when allocating memory and taking other actions. These areas are visible in places such as the memory graphs (*Memory Graph*), system utilities such as `top` (*System Activity (Top)*), in `sysctl` OIDs, and elsewhere.

See also:

For more details on FreeBSD memory management and structure, see https://wiki.freebsd.org/Memory.

6.4.1 Memory Pools

The different memory pools in FreeBSD currently are:

Active
: Active (in use) memory pages referenced by userland (non-kernel). If this memory is unreferenced, after a while it will eventually be marked inactive.

Inactive
: Memory pages which were in use but have not been referenced recently are considered Inactive. Inactive pages can return to active status if they are referenced again. If the pages are "clean" (not in need of swapping or writing to disk), they may be freed directly. Otherwise, they are considered "dirty" and may be moved to the Laundry queue.

Free
: Memory available for immediate use.

Cache
: Memory used by the operating system for caching. For ZFS, this is the ZFS ARC cache (23.05+). For UFS, this is the UFS directory hash.

Note: On current versions this value is only present in the memory graphs, and not reported in other places such as `top`.

The OS will attempt to use RAM when possible for caching rather than allowing it to sit idle and free, so the percentage of free RAM will often appear lower than expected.

Typically, the OS will release cache memory such as the ZFS ARC when there is a shortage of free memory for other processes. However, this may not happen fast enough if a process attempts to allocate a large chunk of memory all at once. In environments where that is a frequent concern, a good practice is to limit the ZFS ARC size as described in *ZFS Tuning*.

Wire
: Memory allocated by the kernel, including the kernel itself, which cannot be paged/swapped and cannot be freed until explicitly released.

Note: In the OS and system utilities such as `top`, the ZFS ARC cache and UFS buffers sizes are included in wired memory. The `top` utility also prints a separate line breaking down ZFS ARC usage.

In the graphs on pfSense Plus software version 23.05 and later, however, these values are removed from the Wired total and graphed separately. ZFS ARC usage is graphed under Cache and UFS buffers are graphed under Buffers.

UserWire
: Similar to Wired, but memory wired by user processes, not the kernel.

Laundry
: Memory pages which are considered "dirty" and are due to be "cleaned". This means they may be eligible for swap (still in use but not recently referenced) or written to disk (e.g. file cache). After cleaning, the pages may be returned to active memory if referenced recently, or to an inactive state otherwise.

Buffers
: Memory used for UFS buffers.

6.4.2 Memory Usage Notes

Free RAM is Wasted RAM

A common mantra among FreeBSD developers is "Free RAM is Wasted RAM". While free RAM is available for use by new/large processes, letting that RAM sit idle and unused is wasting the potential of the system. Allowing the OS to use that free RAM for caching will result in overall higher performance with minimal, if any, consequences.

If a system appears to have low amounts of free memory but is not experiencing any memory allocation failures, the best practice is usually to leave it alone and let the OS use what it can.

On the other hand, if there are problems running scheduled tasks or certain daemons because of low memory, then it may be best to reduce the amount of RAM available for caching. On ZFS systems, see *ZFS Tuning* for advice.

ZFS Disk Activity Increases Memory Usage

On systems using ZFS, large volumes of filesystem activity can lead to an increase in ZFS ARC usage which consumes memory temporarily to boost filesystem performance.

The most common scenario where this becomes a factor is during the first boot after upgrading to a newer version of pfSense software. During an upgrade, large portions of the operating system are rewritten on disk, which can lead to a large chunk of wired memory being allocated to the ZFS ARC to speed up the process.

A similar effect can be observed from other tasks which scan or check large portions of the disk, such as from packages or certain periodic jobs.

In many cases this is harmless because the operating system will yield memory as needed if other processes require more memory. However, that does not always happen quickly. See *ZFS Tuning* for advice on limiting the memory available for the ZFS ARC and tuning the thresholds at which it will yield memory.

Rebooting the firewall after one of these events will return the reported memory usage to a normal level.

Not All Swap Usage is Bad

Installations containing disk partitions dedicated to swap can expand the amount of memory usable by the operating system. Using swap allows idle or unused areas of allocated memory to be relocated to disk instead of RAM. If those areas are requested again, they can be read back into RAM on demand. This is useful for ensuring processes do not run out of memory, but since it involves using disks, it is significantly slower than RAM.

Given that knowledge, one might assume that any usage of swap is bad and should be avoided at all costs, however, that is not universally true.

For example, there are daemons which run at all times but only periodically require the resources they have allocated. While the process is idle, there is little harm in allowing it to be swapped in favor of something that needs RAM at that moment, or even for caching.

Swap consumption generally only becomes a concern when there is significant usage, such as near 50% or higher. At that point it becomes necessary to inspect everything consuming memory to see what can be reduced or tuned.

Additionally, allowing the operating system to have swap space enables the ability to save kernel panic text dumps which facilitate debugging kernel problems. Without swap space to use as a dump device, the operating system will print the crash dump to the console and cannot store it anywhere.

6.5 Hardware Tuning and Troubleshooting

The underlying operating system beneath pfSense® software can be fine-tuned in several ways. A few of these tunables are available under **Advanced Options** (See *System Tunables*). Others are outlined in the FreeBSD main page tuning(7).

The default installation includes a well-rounded set of values tuned for good performance without being overly aggressive. There are cases where hardware or drivers necessitate changing values or a specific network workload requires changes to perform optimally.

The hardware sold in the Netgate Store is tuned further since Netgate has detailed knowledge of the hardware, removing the need to rely on more general assumptions.

Note: Refer to *Managing Loader Tunables* for making edits to define loader tunables persistently.

Changes in loader tunables require a firewall reboot to take effect.

- *General Issues*
  - *File System Tuning*
  - *Mbuf Exhaustion*
  - *Disable MSI-X*
  - *PPPoE with Multi-Queue NICs*
  - *TSO/LRO*
  - *IP Input Queue (`intr_queue`)*
- *Card-Specific Issues*
  - *Broadcom `bce(4)` Cards*
  - *Broadcom `bge(4)` Cards*
  - *Chelsio `cxgbe(4)` Cards*
  - *Intel igb(4) and em(4) Cards*
  - *Intel ix(4) Cards*
  - *VMware `vmx(4)` Interfaces*
  - *Flow Control*

6.5.1 General Issues

File System Tuning

For information on tuning ZFS memory usage, see *ZFS Tuning*.

Mbuf Exhaustion

A common problem encountered by users of commodity hardware is mbuf exhaustion. To oversimplify, "mbufs" are network memory buffers; portions of RAM set aside for use by networking for moving data around.

The count of active mbufs is shown on the dashboard and is tracked by a graph under **Status > Monitoring**.

See also:

For details on mbufs and monitoring mbuf usage, see *Mbuf Clusters*.

If the firewall runs out of mbufs, it can lead to a kernel panic and reboot under certain network loads that exhaust all available network memory buffers. In certain cases this condition can also result in expected interfaces not being initialized and made available by the operating system. This is more common with NICs that use multiple queues or are otherwise optimized for performance over resource usage.

Additionally, mbuf usage increases when the firewall is using certain features such as *Limiters*.

To increase the amount of mbufs available, add the following as a *Loader Tunable*:

```
kern.ipc.nmbclusters="1000000"
```

On 64-bit systems with multiple GB of RAM, 1 million (1000000) mbuf clusters is a safe starting point. Should mbuf clusters become fully allocated, that would consume about 2.3 GB of physical memory:

```
1000000 memory buffer clusters available × (2048 KB per cluster + 256 bytes per
memory buffer)
```

The amount of available clusters can be reduced for systems with low amounts of physical RAM, or increased further as needed, as long as the value does not exceed available kernel memory.

Some network interfaces may need other similar values raised such as `kern.ipc.nmbjumbop`. In addition to the graphs mentioned above, check the output of the command `netstat -m` to verify if any areas are near exhaustion.

Disable MSI-X

Message Signaled Interrupts are an alternative to classic style Interrupts for retrieving data from hardware. Some cards behave better with MSI, MSI-X, or classic style Interrupts, but the card will try the best available choice (MSI-X, then MSI, then Interrupts).

MSI-X and MSI can be disabled via loader tunables. Add the following as *Loader Tunables*:

```
hw.pci.enable_msix="0"
hw.pci.enable_msi="0"
```

To nudge the card to use MSI, disable only MSI-X. To nudge the card to use regular Interrupts, disable both MSI and MSI-X.

PPPoE with Multi-Queue NICs

Tip: Before attempting this tuning, try the `if_pppoe` backend for PPPoE WANs. It is much faster and more efficient and may eliminate the need for additional tuning. For details, see *Use if_pppoe Kernel Module*.

Network cards capable of using multiple queues rely on hashing to assign traffic to a particular queue. This works well with IPv4/IPv6 TCP and UDP traffic, for example, but fails with other protocols such as those used for PPPoE.

This can lead to a network card under performing with the default network settings, as noted on #4821 and FreeBSD PR 203856. This problem primarily affects systems with multiple CPUs and/or CPU cores, as those are the systems which benefit most from multiple NIC queues.

Adding a **System Tunable** or *Loader Tunable* entry for `net.isr.dispatch=deferred` can lead to performance gains on affected hardware.

Tuning the values of `net.isr.maxthreads` and `net.isr.numthreads` may yield additional performance gains. Generally these are best left at default values matching the number of CPU cores, but depending on the workload may work better at lower values.

Warning: In the past, `deferred` mode has led to issues on 32-bit platforms, such as crashes/panics, especially with ALTQ. There have been no recent reports, however, so it should be safe on current releases.

TSO/LRO

The settings for **Hardware TCP Segmentation Offload** (TSO) and **Hardware Large Receive Offload** (LRO) under **System > Advanced** on the **Networking** tab default to **checked** (disabled) for good reason. Nearly all hardware/drivers have issues with these settings, and they can lead to throughput issues. Ensure the options are checked. Sometimes disabling via `sysctl` is also necessary.

Add the following as a *Loader Tunable*:

```
net.inet.tcp.tso="0"
```

IP Input Queue (`intr_queue`)

This will show the current setting:

```
sysctl net.inet.ip.intr_queue_maxlen
```

However, in largely loaded installations this may not be enough. Here is how to check:

```
sysctl net.inet.ip.intr_queue_drops
```

If the above shows values above `0`, try doubling the current value of `net.inet.ip.intr_queue_maxlen`.

For example:

```
sysctl net.inet.ip.intr_queue_maxlen="3000"
```

Keep performing the above until the point is found where drops are eliminated without any adverse effects.

Afterward, add an entry under **System > Advanced**, **System Tunables** tab to set `net.inet.ip.intr_queue_maxlen` to `3000`

6.5.2 Card-Specific Issues

Broadcom `bce(4)` Cards

Several users have noted issues with certain Broadcom network cards, especially those built into Dell hardware. If `bce` interfaces are behaving erratically, dropping packets, or causing crashes, then the following tweaks may help.

Add the following as *Loader Tunables*:

```
kern.ipc.nmbclusters="1000000"
hw.bce.tso_enable="0"
hw.pci.enable_msix="0"
```

That will increase the amount of network memory buffers, disable TSO directly, and disable MSI-X.

Packet loss with many (small) UDP packets

If a lot of packet loss is observed with UDP on `bce(4)` cards, try changing the `netisr` settings. These can be set as system tunables under **System > Advanced**, on the **System Tunables** tab. On that page, add two new tunables:

```
net.isr.direct_force="1"
net.isr.direct="1"
```

Broadcom `bge(4)` Cards

See above, but change `bce` to `bge` in the setting names.

Chelsio `cxgbe(4)` Cards

Rate Limiting

By default, the Chelsio driver implements rate limiting capabilities which can be undesirable for routing performance. To disable this rate limiting capability, add the following as a *Loader Tunable*:

```
hw.cxgbe.niccaps_allowed="1"
```

Resource Allocation

It is possible to disable the allocation of resources that are not related to the router so that the network adapter can use its entire set of resources for the corresponding functions:

Add the following as *Loader Tunables*:

```
hw.cxgbe.toecaps_allowed="0"
hw.cxgbe.rdmacaps_allowed="0"
hw.cxgbe.iscsicaps_allowed="0"
hw.cxgbe.fcoecaps_allowed="0"
```

Chelsio TCP Offload Engine (TOE)

There is an experimental driver for the Chelsio TCP Offload Engine (TOE) via the `t4_tom` kernel module. TOE offloads the entire TCP connection to hardware, but this also can cause problems with connection handling.

> **Warning:** According to analysis of Chelsio TCP Offload Engine behavior by Calomel, under certain circumstances traffic handled by TOE bypasses the FreeBSD TCP stack and will not be filtered nor logged by PF.
>
> Users have observed this behavior particularly when using reroot instead of a full reboot, which left services exposed.

While this feature could, in theory, reduce CPU usage for handling TCP connections, it should not be used in any role where security is a primary concern. For internal routing-only or private endpoint roles it may be acceptable.

Intel igb(4) and em(4) Cards

Certain intel igb cards, especially multi-port cards, can easily exhaust mbufs and cause kernel panics. The following tweak will prevent this from being an issue. Add the following as a *Loader Tunable*:

```
kern.ipc.nmbclusters="1000000"
```

That will increase the amount of network memory buffers, allowing the driver enough headroom for its optimal operation.

Intel ix(4) Cards

Auto-Negotiation of Non-default Speeds

The `ix(4)` hardware may be capable of linking at non-default speeds such as 2.5G and 5G, depending on the chipset, installed SFP module, and so on. However, the driver will not auto-negotiate at these speeds by default.

To negotiate at these speeds the driver must be configured to advertise these rates.

This is managed using the *Runtime Tunable* OID `dev.ix.<X>.advertise_speed` where `<X>` is the interface number. For example, to adjust the advertised speeds of `ix3`, use `dev.ix.3.advertise_speed`. Use `sysctl -x <oid>` at the CLI to view the values in hexadecimal instead of decimal.

Table 3: `ix(4)` Driver Speed Values

Speed	Hex	Decimal
100M	0x1	1
1G	0x2	2
10G	0x4	4
10M	0x8	8
2.5G	0x10	16
5G	0x20	32

The default value of the `sysctl` OID varies by hardware and SFP modules. Common defaults include `0x6` (hex) or `6` (decimal) to advertise 1G and 10G, and `0xb` (hex) or `11` (decimal) to advertise 10M, 100M, and 1G.

To auto-negotiate at the non-default speeds, add the value of the desired speeds to the existing value.

For example, to advertise speeds of 1G, 2.5G, 5G, and 10G, add their corresponding values: `0x2+0x10+0x20+0x4` (hex) or `2+16+32+4` (decimal) for a total of `0x36` (hex) or `54` (decimal). Values in hexadecimal must start with `0x`.

When setting the value at the CLI, the command will print an error message if the value is not valid for the interface. The tunables GUI does not currently validate the content.

General Tuning

Add the following as *Loader Tunables* if the values are not already this large, or larger:

```
kern.ipc.nmbclusters="1000000"
kern.ipc.nmbjumbop="524288"
```

Add a (*Runtime Tunable*) if the value is not already this large, or larger, or set to `0` (automatic):

```
hw.intr_storm_threshold="10000"
```

VMware `vmx(4)` Interfaces

VMware VMXNET interfaces can utilize multiple queues when using MSI-X. Multiple queues enable network performance to scale with the number of vCPUs and allows for parallel packet processing. Transmit and Receive descriptors may also be increased to help with throughput.

Add the following *Loader Tunables*:

Note: Some options have a separate set of tunables for each individual network interface. In these cases, replace <id> replace with the device ID such as `0`, `1`, etc. where the ID number matches the interface number. For example, tunables for `vmx3` are under `dev.vmx.3`.

```
hw.pci.honor_msi_blacklist="0"
dev.vmx.<id>.iflib.override_ntxds="0,4096"
dev.vmx.<id>.iflib.override_nrxds="0,2048,0"
```

Save the file, then reboot and check the change with `dmesg | grep -Eiw 'descriptors|queues'` at a command prompt.

Flow Control

In some circumstances, flow control may need to be disabled. The exact method depends on the hardware involved, as in the following examples:

These example entries are *Loader Tunables*:

cxgbe(4)

```
hw.cxgbe.pause_settings="0"
```

ixgbe(4) (aka ix)

```
hw.ix.flow_control="0"
```

These example entries go in **System > Advanced**, on the **System Tunables** tab (*System Tunables*):

Note: Some options have a separate set of tunables for each individual network interface. In these cases, replace <id> replace with the device ID such as 0, 1, etc. where the ID number matches the interface number. For example, tunables for igc3 are under dev.igc.3.

igc(4)

```
dev.igc.<id>.fc="0"
```

igb(4)

```
dev.igb.<id>.fc="0"
```

em(4)

```
dev.em.<id>.fc="0"
```

For ix and others, the flow control value can be further tuned:

0
No Flow Control

1
Receive Pause

2
Transmit Pause

3
Full Flow Control, Default

6.6 ZFS Tuning

ZFS on pfSense® software is much more robust than UFS and enables a variety of handy ZFS features, such as ZFS Boot Environments. Some ZFS features are not relevant on systems acting as a firewall, however. For example, ZFS can be tuned to use much less memory and sacrifice performance because filesystem performance is not as critical in a firewall role as compared to others.

Note: This is **not** a general purpose ZFS tuning guide. The suggestions in this document primarily apply to ZFS on a device acting in a firewall role. The tuning requirements for other use cases, such as for file or database servers, can be vastly different.

6.6.1 ZFS Memory Tuning

Installations utilizing ZFS may encounter higher than expected memory usage. This memory may appear as "Wired", "Cache", or "ARC" depending on how the memory usage is being checked (*Memory Management*). In each case the usage is the same, but may be reported differently. This memory is considered "Wired" because it cannot be paged out (e.g. swap) so it has the appearance of consuming large portions of memory that other things may need.

ZFS attempts to optimize performance with its Adaptive Replacement Cache (ARC) which can be aggressive in its consumption of RAM. The classic phrase *"Free RAM is wasted RAM"* also applies here, and ZFS will make good use of the memory when allowed to do so. ZFS *will* yield this RAM if other processes require more memory, but it may not give up memory fast enough for every use case. For example, if a process starts and attempts to allocate a large block of memory rapidly, and there is insufficient free RAM or swap, then it may fail.

Tip: It's normal and OK to see a little swap usage when ARC usage is high, as items swapped are generally idle and not actively used, whereas the RAM can be put to better use caching until such time as those pages are needed.

If the swap usage is high or nearly full, however, that is likely a sign that there is too much contention for memory and if the ARC usage is also high, it likely needs to be tuned to a much lower maximum.

ARC Maximum/Minimum

The ARC size has default limits which allow it to consume large portions of the system RAM:

Default Maximum
: The default maximum ARC size (`vfs.zfs.arc.max`) is automatic (`0`) and uses 1/2 RAM or the total RAM minus 1GB, whichever is **greater**.

  As the default processes on pfSense software can easily consume 1GB of RAM given the opportunity, this situation can become problematic in some cases!

Default Minimum
: The default minimum ARC size (`vfs.zfs.arc.min`) is automatic (`0`) and is 1/32 of all RAM or 64M, whichever is **greater**.

To tune the **maximum**, configure a system tunable (*System Tunables*) for `vfs.zfs.arc.max` and set it to a size expressed in **bytes**.

Tip: For `X` MB RAM, use `X * 1024 * 1024`. For example, 1GB is `1073741824` bytes, 512M is `536870912`, 256M is `268435456`, 128M is `134217728`, and 64M is `67108864`.

The exact amount depends on the available resources, but to be conservative on systems with low RAM, set it to 25% of the total RAM or less.

The **minimum** rarely requires tuning, but can be lowered if necessary.

Note: The ZFS ARC limit will take effect immediately, but the OS may not free the wired memory previously used by the ARC right away. Rebooting after configuring the limits will ensure they are fully respected.

Prefetch

ZFS prefetch attempts to read more blocks than initially requested into the ARC in case they are needed in the near future. This can also cause increased ARC usage.

To disable prefetch, configure a system tunable (*System Tunables*) and set `vfs.zfs.prefetch.disable` to `1`

Free Target

The ARC Free Target is the amount of **free** RAM under which the ARC will start to give up memory. This can be awkward to calculate since it is dealing with a remaining amount, not a total, and the fact that this must be given in pages (chunks of 4096 bytes).

For example, to trigger ARC cleanup when the system hits 60% RAM usage, this number should be 40% of total RAM bytes divided by 4096.

On a system with 4GB RAM, this would be when the system has 1.6GB free. To find the equivalent pages value, calculate it with `(<percent> * <total bytes>)/<page size>`.

Tip: The page size is almost always `4096`, but check the value of `sysctl hw.pagesize` to be certain.

In this case that would be `(0.40 * (1024 * 1024 * 1024 * 4))/4096` which results in `419430` pages.

Configure a system tunable (*System Tunables*) to set `vfs.zfs.arc.free_target` to the desired number of pages (e.g. `419430`).

6.7 Console Types

There are two console types available with pfSense® software, VGA and Serial. The active default console depends on the image/installer used and configuration settings. The difference between the two console types is explained in more detail below.

6.7.1 VGA Console

The VGA (video) console is a console with a monitor and keyboard. The video console requires hardware with a connection for a monitor (e.g. HDMI, VGA) and keyboard (USB, PS/2). In some cases a serial BIOS that does VGA redirection may work.

The VGA console is active by default using the normal memstick installer or ISO.

6.7.2 Serial Console

The serial console uses a serial/COM port to communicate with a serial client. It is primarily intended for systems without a monitor or keyboard. The serial console can also be used on systems where those are either not available or not wanted, so long as the hardware has at attached (non-USB) serial port.

The serial console is active by default when installing using the *serial* memstick and may be enabled under **System > Advanced** on VGA images.

Accessing the serial console requires a null modem serial cable attached between the `COM1` port on the firewall and a serial client. A hardware serial port is required on the firewall, but the client may use a USB serial adapter if needed.

Serial clients are quite common, often pre-installed on an operating system or easily available. The free PuTTY client is the most popular GUI choice. Other choices include GNU `screen`, `tip`, `cu` and `minicom`.

See also:

See *Connect to the Console* for details on how to connect to a serial console.

The default speed of the serial port is `115200/8/N/1`. The serial port speed may be changed under **System > Advanced**.

If the device has a BIOS accessible over serial console, it is also possible that it will not be using the same serial speed that the OS is using.

The most common serial speeds to try would be: `115200`, `38400`, and `9600`.

If the BIOS serial speed does not match the OS serial speed, the best practice is to adjust one or the other to match, so that POST messages may be viewed as well as the OS messages without having to adjust the client

6.8 Connect to the Console

A connection to the console on the target hardware is a requirement to run the installer.

6.8.1 Connecting to a VGA Console

For hardware with a VGA console, this is as simple as connecting a monitor and keyboard.

6.8.2 Connecting to a Serial Console

For hardware with a serial console, the process is more involved and requires a client PC with an appropriate port and terminal software. Follow the instructions below to connect using a serial console.

The instructions in this section cover general serial console topics. Some devices, such as firewalls from the Netgate Store, require slightly different methods to connect to the serial console. For devices from the Netgate Store, visit the Netgate Documentation for model-specific serial console instructions.

Serial Console Requirements

Connecting to a serial console on most firewalls requires the correct hardware on every part of the link, including:

- The client PC must have a physical serial port or a USB-to-Serial adapter
- The firewall must have a physical serial port
- A null modem serial cable and/or adapter, or a device-specific serial cable
- A terminal program on the client, such as PuTTY
- The correct serial settings for the client software

For most of the firewalls purchased from the Netgate Store, the only hardware requirement is a USB A to Mini-B cable. See Netgate Documentation for specifics.

In addition to the proper hardware connection, a serial console client program must also be available on the client PC, and the serial speed and other settings must be available.

Locating a Serial Port (Server/Firewall)

First, ensure the firewall hardware has a serial port. To use the serial console, the hardware must have a physical serial port at `COM1`. Embedded units typically have a DB9 (9-pin) serial port, but some have an RJ45 style console connector with an adapter cable that ends with a DB9 connector.

Connect a Serial Cable

First, a null modem serial cable must be connected between the firewall and a client PC. Depending on the serial port and cable being used, a serial cable gender changer may also be necessary to match the available ports.

If a real null modem serial cable is unavailable, a null modem adapter can be used to convert a standard serial cable into a null modem cable.

If the client PC does not have a physical serial port, use a USB-to-Serial adapter.

Locate the Client Serial Port

On the client PC, the serial port device name must be determined so that the client software can be used on the correct port.

Windows

On Windows clients, a physical serial port is typically `COM1`. With a USB-to-Serial adapter, it may be `COM3`. Open **Device Manager** in Windows and expand **Ports (COM & LPT)** to find the port assignment.

macOS

On macOS, the name can be tricky for a user to determine since it can vary based several factors. On recent versions of macOS, the devices are likely to be named `/dev/cu.usbserial-<id>` where the `<id>` is an identifier for the USB serial adapter, such as a serial number.

When in doubt, run `ls -l /dev/cu.*` from a Terminal prompt to see a list of available USB serial devices and locate the appropriate one for the hardware. If there are multiple devices, the correct device is likely the one with the most recent timestamp or highest ID.

Linux

The device associated with a USB-to-Serial adapter is likely to show up as `/dev/ttyUSB0`. Look for messages about the device attaching in the system log files or by running `dmesg`.

Note: If the device does not appear in `/dev/`, check to see if the device requires additional drivers.

FreeBSD

The device associated with a USB-to-Serial adapter is likely to show up as `/dev/cuaU0`. Look for messages about the device attaching in the system log files or by running `dmesg`.

Determine Serial Console Settings

The settings for the serial port, including the speed, must be known before a client can successfully connect to a serial console.

Whichever serial client is used, ensure that it is set for the proper Speed (`115200`), Data Bits (`8`), Parity (`No`), and Stop Bits (`1`). This is typically written as `115200/8/N/1`.

Note: Some hardware defaults to a slower speed. This is relevant to the BIOS and initial output, not to pfSense® software which defaults to `115200`.

Many serial clients default to `9600/8/N/1`, so adjusting these settings is required to connect. Use `115200/8/N/1` with pfSense software regardless of the setting of the hardware/BIOS.

For hardware using BIOS serial speeds other than `115200`, change the baud rate to `115200` in the BIOS setup so the BIOS and pfSense software are both accessible with the same settings. Refer to the hardware manual for information on setting its baud rate.

`115200` is the default speed pfSense software uses out of the box, but the serial speed used by pfSense software can be changed later. See *Serial Console Speed*.

Locate a Serial Client

A serial client program must be used on the client PC. The most popular client for Windows is PuTTY, which is free and works well. PuTTY is also available for Linux and can be installed on macOS using brew. On UNIX and UNIX-Like operating systems, the `screen` program is readily available or easily installed, and it can also be used to connect to serial ports from a terminal program or system console.

Windows

PuTTY is the most popular free choice for serial communication on Windows. SecureCRT is another client that works well.

Warning: Do not use HyperTerminal. Even if it is already present on the client PC, it is unreliable and prone to formatting incorrectly and losing data.

macOS

On macOS clients, the GNU `screen` utility is the easiest and most common choice. ZTerm and `cu` (similar to FreeBSD) can be used as well.

Linux

On Linux clients, the GNU `screen` utility is the easiest and most common choice. Programs such as PuTTY, `minicom`, or `dterm` can be used as well.

FreeBSD

On FreeBSD clients, the GNU `screen` utility is the easiest and most common choice.

As an alternative, use the built-in program `tip`. Typing `tip com1` (Or `tip ucom1` if using a USB serial adapter) will connect to the first serial port. Disconnect by typing `~.` at the start of a line.

Start a Serial Client

Now that all the requirements have been met, it is time to run the serial client.

If the client software is not covered in this section, consult its documentation to determine how to make a serial connection.

PuTTY

- Start PuTTY
- Select *Serial* for the **Connection Type**
- Enter the serial port device name for **Serial Line**, e.g. `COM3` or `/dev/ttyUSB0`.
- Enter the appropriate **Speed**, e.g. `115200`
- Click **Open**

Minicom

```
$ minicom -D /dev/ttyUSB0 -R 115200
```

GNU screen

- Open a terminal / command prompt
- Invoke the `screen` command using the path to the serial port, for example:

```
$ sudo screen /dev/ttyUSB0 115200
```

In some cases there may be a terminal encoding mismatch. If this happens, run `screen` in UTF-8 mode:

```
$ sudo screen -U /dev/cu.usbserial-1234 115200
```

The standard screen controls apply. Press `Ctrl-A`, \ to quit, or `Ctrl-A`, `Ctrl-\` in some cases.

BSD tip

The `tip` command on FreeBSD consults `/etc/remotes` and connects to serial ports based on the settings there. To set up a connection to a USB-to-serial adapter at `115200`, add a line such as the following to `/etc/remote`:

```
ucom1fast:dv=/dev/cuaU0:br#115200:pa=none:
```

To access the port, invoke `tip`:

```
$ tip ucom1fast
```

To quit, press **Enter**, then type `~.`. If connected through a terminal ssh client, `~~.` may need to be used instead so that the ssh client itself doesn't interpret the keys.

6.9 Cryptographic Accelerators

Cryptographic acceleration is available on some platforms, typically on hardware that has it available in the CPU like AES-NI, or built into the board such as the ones used on Netgate ARM-based systems. Most cryptographic accelerator hardware with FreeBSD drivers will work, provided the drivers are in the kernel or available as loadable modules.

Note: Some modules and hardware are only available on pfSense® Plus software.

6.9.1 Compatible Devices

Cryptographic accelerator devices compatible with pfSense® software include:

AES-NI
: These instructions are present in most modern CPUs.

IPsec Multi-Buffer (IPsec-MB, IIMB) Cryptographic Acceleration [Plus only]
: The Intel® Multi-Buffer Crypto for IPsec Library, often shortened to IPsec-MB or IIMB.

  IPsec-MB assists VPN performance by replacing the cryptographic functions provided by the kernel for AES-CBC, AES-GCM, and ChaCha20-Poly1305 with accelerated functions that utilize the optimal CPU SIMD instruction set (Single Instruction, Multiple Data), such as SSE, AVX, AVX2, and AVX512 (Advanced Vector Extensions). As such, this feature requires a CPU which includes one or more of these features, but they are common on current hardware.

  This offers faster speeds and lower CPU utilization not only for IPsec but for any VPN utilizing the accelerated algorithms in the kernel. In addition to IPsec this also includes OpenVPN DCO and WireGuard. Exact performance varies by hardware, workload, and available CPU instruction sets. IPsec-MB is faster than AES-NI and can even meet or exceed the performance of dedicated acceleration hardware such as QAT on current versions of pfSense software.

  IPsec-MB can be loaded alongside other cryptographic modules without conflicting, so it is separate from the other options. That said, when it is enabled it will take precedence over all algorithms it is capable of accelerating, even if other options could potentially be faster (e.g. QAT).

There are several aspects of IPsec-MB behavior which can be fine-tuned. See *Tuning IPsec-MB* for details.

Intel QuickAssist Technology (QAT) [Plus only]
Intel QuickAssist Technology (QAT) accelerates many types of AES and SHA operations, such as AES-GCM encryption, and QAT is ideal for use with IPsec and OpenVPN DCO. It is currently the fastest acceleration option for compatible algorithms.

QAT devices can be found in certain Intel-based platforms such as select models of c3000 and c2000 SoCs, and also in QAT add-on cards. Several Netgate hardware models include QAT devices, such as the 4100, 5100, 6100, 7100, 8200, and more.

CESA [Plus only]
Present on some ARM platforms such as the Netgate 3100.

SafeXcel [Plus only]
Present on some ARM platforms such as the Netgate 2100 and 1100.

Note: For specifics on which hardware accelerators are available on Netgate hardware, and relevant performance data, visit the Netgate Store.

6.9.2 Activating the Hardware

Some hardware acceleration is active at all times and there is no way to disable it short of removing the crypto card if it is a hardware add-on. For example, CESA acceleration cannot be disabled because it's an integrated feature of the system and the drivers are present the kernel.

Others, such as QAT, IPsec-MB, AES-NI, or SafeXcel require choosing the appropriate module under **System > Advanced** on the **Miscellaneous** tab (See *Cryptographic & Thermal Hardware*). Choose the appropriate module to match the hardware for **Cryptographic Hardware** and then Save. The module will be loaded and available immediately.

To deactivate a loaded module, select *None* for **Cryptographic Hardware**, Save, and then reboot the system.

6.9.3 Confirming Accelerator Use

Confirming that the cryptographic acceleration device is being used by the firewall can be tricky, depending on the hardware in question.

Most often the evidence of cryptographic accelerator use is apparent in one or more of the following observations:

- Increased VPN throughput
- Decreased system load (e.g. CPU utilization) for similar levels of VPN throughput

In cases where it is not clear, some cryptographic accelerators show signs of use by checking for interrupt activity on the device using `vmstat -i | grep <name>`, where `<name>` corresponds to the name of the device:

QAT
Use the shell command `vmstat -i | grep qat`

CESA
Use the shell command `vmstat -i | grep cesa`

SafeXcel
Use the shell command `vmstat -i | grep safexcel`

In each of these cases, first check that there is any output at all. If the device has not been used at all since the firewall last rebooted or loaded the device driver, there will be no output from the command.

Note: To see if the driver is loaded, check `kldstat -v | grep <name>` to ensure the driver is present, and check `dmesg | grep <name>` to see if the device was detected.

If there is output from `vmstat -i` for the device, check the third entry on the line, which is the total number of interrupts observed on the device(s). If this number is increasing with VPN activity, the device is being used by the firewall. For example:

```
# vmstat -i | grep qat
irq300: qat0                       5481147          3
```

In that output the `5481147` number represents the number of interrupts on the `qat0` device. Run the command again after transferring data across the VPN, and compare the number.

Note: If the command produces no output at all, the device is not being used or the device driver is not loaded.

6.9.4 Practical Use

IPsec

IPsec will take advantage of acceleration automatically when an active accelerator is compatible with the cipher chosen for a tunnel. For QAT and AES-NI, the optimal cipher choice is AES-GCM.

OpenVPN

To take advantage of acceleration in OpenVPN, choose a cipher which is accelerated by the available hardware, such as AES-256-GCM.

When using OpenVPN in DCO mode on pfSense Plus software, OpenVPN can use QAT or IPsec-MB to accelerate its encryption automatically, assuming the features are enabled (QAT and/or IPsec-MB modules are loaded and active on compatible hardware). If the hardware cannot accelerate via QAT or IPsec-MB, but it does include AES-NI instructions, then ensure the AES-NI module is **loaded** or DCO mode cannot use AES-NI.

See also:

DCO and Hardware Cryptographic Acceleration

In non-DCO mode, such as on pfSense CE, nothing needs to be selected for OpenVPN to utilize AES-NI. The OpenSSL engine has its own code for handling AES-NI in this mode that works well without using additional modules.

6.9.5 Tuning IPsec-MB

The behavior of IPsec-MB can be tuned by using one of several system tunables configurable on the *System Tunables*:

kern.crypto.iimb.enable_aescbc
Enables handling of AES-CBC by IIMB. IIMB can be slower than QAT for CBC, so this toggle can disable handling for AES-CBC while accelerating other algorithms. This allows IPsec-MB and QAT to coexist in such environments. Compatible only with x86-64.

Default is enabled (`1`). To disable, set a value of `0`.

kern.crypto.iimb.enable_multiq
Uses multiple queues to handle encryption jobs, i.e. each session is bound to a job thread. There are only a few job threads available:

- 1 thread for < 4 CPUs
- 2 threads for < 8 CPUs
- 4 threads for >= 8 CPUs

Default is `1`.

kern.crypto.iimb.use_task (default 0)
Use a separate taskq for running the encryption job completion callbacks. The callbacks are functions in the VPN code that send the packets onto the next step, e.g. `ip_output` or `netisr_queue` for input into the local stack. This option helps on high-performance systems (fast CPU, fast NICs).

Default is disabled (`0`). To enable, set a value of `1`.

Developer Tunables

The following tunables are typically only needed during development or debugging:

kern.crypto.iimb.arch
Used to override the SIMD architecture on x86. By default, it uses the best option available on the CPU. This option allows comparing different CPU feature benchmarks.

Options: `auto` (default), `sse`, `avx`, `avx2`, `avx512`.

kern.crypto.iimb.prefetch
Pre-fetch encryption keys before calling the crypto function, this might help with micro-performance, but thus far has not led to any significant measurable differences.

Default is enabled (`1`). To disable, set a value of `0`.

kern.crypto.iimb.max_jobs
Maximum number of batched jobs. The IIMB thread will collect up to this many jobs and handle them in a batch. Maximum is `256`, but it can be tuned to be smaller. Limited to `256` because thus far developers have not seen larger values lead to any significant measurable differences in performance, because it adds latency to the network stack.

Default value is `256`.

6.10 Disabling Sounds/Beeps

Some hardware has a PC Speaker which can be used as a means of notification. By default, the firewall will play a tone at startup/shutdown and will emit a beep when a user logs into the GUI. Additionally, some packages are capable of producing beeps for events.

6.10.1 Disable Startup/Shutdown Tune

The startup and shutdown tunes may be disabled as follows:

- Navigate to **System > Advanced**, **Notifications** tab
- Check **Disable the startup/shutdown beep**
- Click **Save**

6.10.2 Disable Login Beep

The GUI login beep happens because the GUI login event is recorded by syslog under the LOG_AUTH facility. Messages in this facility trigger the operating system to generate a beep. To disable the beep, the GUI login messages must be suppressed as follows:

- Navigate to **System > Advanced**, **Admin Access** tab
- Check **Disable logging of webConfigurator successful logins**
- Click **Save**

6.10.3 Disable All Sounds

As an alternative, the system bell may be disabled globally:

- Navigate to **System > Advanced**, **System Tunables** tab
- Click  to create a new tunable entry using the following values:

  Tunable
  `kern.vt.enable_bell`

  Description
  `Control system sounds`

  Value
  `0`

- Click **Save**

See also:

- *Halting and Powering Off the Firewall*
- *Rebooting the Firewall*
- *Network Interface Drivers with ALTQ Traffic Shaping Compatibility*
- *Troubleshooting Disk and File System Issues*
- *Troubleshooting Boot Issues*

- *Troubleshooting DMA and LBA Errors*
- *Troubleshooting High CPU Load*
- *Troubleshooting Disk and File System Issues*
- *Troubleshooting Lost Traffic or Disappearing Packets*
- *Troubleshooting Unexpected Reboots*

The pfSense® software distribution is compatible with most hardware supported by FreeBSD.

Current versions of pfSense software are compatible with 64-bit (amd64, x86-64) architecture hardware and Netgate ARM-based firewalls.

Alternate hardware architectures such as Raspberry Pi, other Non-Netgate ARM devices, PowerPC, MIPS, SPARC, etc. are not compatible.

6.11 Hardware Compatibility

The best way to ensure that hardware is compatible with pfSense software is to buy hardware from the Netgate Store that has been tested and known to work well with pfSense software. The hardware in the store is tested with each release of pfSense software and is tuned for optimal performance.

For home-built solutions, the FreeBSD Hardware Notes for the FreeBSD version used in a given build of pfSense software is the best resource for determining hardware compatibility. pfSense software version 2.8.0-RELEASE is based on 15.0-CURRENT@bf06074106cf. Another good resource is the Hardware section of the FreeBSD FAQ.

6.11.1 Network Adapters

A wide variety of wired Ethernet Network Interface Cards (NICs) have FreeBSD drivers, and are thus compatible with pfSense software. However, not all NICs are created equal. The hardware can vary greatly in quality from one manufacturer to another.

The best practice is to use Intel NICs because they have solid drivers in FreeBSD and they perform well. Most hardware sold in the Netgate Store contains Intel NICs.

Of the various other PCIe/PCI cards with FreeBSD drivers, some work fine, others may suffer from instability or poor performance. In some cases, FreeBSD may have drivers for a particular NIC, but specific implementations of the chipset may be lower in quality or have implementation problems. When in doubt, search the Netgate Forum for experiences of others using the same or similar hardware.

When a firewall requires the use of VLANs, select adapters capable of performing VLAN processing in hardware. This is discussed in *Virtual LANs (VLANs)*.

USB Network Adapters

USB network adapters of any make/model should not be used due to their unreliability and poor performance.

Wireless Adapters

Wireless adapters and recommendations are covered in *Wireless*.

CHAPTER

SEVEN

INSTALLING AND UPGRADING

Hardware from the Netgate Store is preloaded with pfSense® Plus software. To reinstall pfSense Plus software or to install pfSense Plus software or pfSense CE software to other hardware, download the Netgate Installer.

Warning: Hardware preloaded with pfSense software from commercial vendors other than the Netgate Store or authorized partners must not be trusted. Third parties may have made unauthorized, unknown alterations or additions to the software. Selling preloaded copies of pfSense software is a violation of the Trademark Usage Guidelines.

If pfSense software was preloaded on third party hardware by a vendor, wipe the system and reinstall it with a genuine copy.

See also:

If something goes wrong during the installation process, see *Troubleshooting Installation Issues*.

This chapter also covers upgrading pfSense software installations (*Upgrade Guide*) which keeps them up-to-date with the latest security, bug fixes, and new features. This includes the new ability to *Migrate from pfSense® CE software to Netgate pfSense Plus software*.

7.1 Netgate Installer

The Netgate Installer is the only available method for installing pfSense® software on all devices.

The Netgate Installer image does not contain installation packages for pfSense software, it fetches them over the Internet. This allows a single installer to offer choices between multiple versions of pfSense software without needing to package them all into a gigantic single disk image or multiple separate images. This also means that when installing, the target device always receives the most up-to-date versions of available components. This offers a consistent experience no matter which type of device is the target.

Warning: This installer checks if a system is eligible to access pfSense Plus software before proceeding. If a system is ineligible, the user can either follow the directions to become eligible or install pfSense CE Software instead.

See also:

For a full walkthrough of the installation process using the Netgate Installer, see *Perform the Installation*.

7.1.1 Limitations

The Netgate Installer has the following known limitations at this time:

- Limited use of PPP-based WANs

  The installer can use PPPoE, but it cannot use other PPP types such as L2TP, PPTP, and PPP (e.g. 4G Cellular).

- Not compatible with 32-bit ARM devices.

7.2 Download Installation Media

Installation images can be downloaded from the Netgate Store at https://shop.netgate.com/products/netgate-installer using a Netgate Store Account.

Note: The installer is free but uses the Netgate Store to handle the download process.

There are three installation images, each for different types of hardware:

- AMD64 Memstick (Serial and VGA) for installing via USB media
- AMD64 ISO image for installing via IPMI or optical drive
- AARCH64 Memstick for installing on 64-bit ARM devices from Netgate, such as the Netgate 1100 and Netgate 2100.

Note: Customers who have purchased firewalls preloaded with pfSense Plus software from the Netgate Store already have a Netgate Store Account and access to the Netgate Installer. The Netgate Product Manuals contain specific instructions for each model.

Some Netgate devices can also run Community Edition, but pfSense Plus software offers the best user experience.

For other hardware, continue reading.

- Navigate to the Netgate Store in a web browser on a client device.
- Login using the Netgate Store Account
- Navigate to https://shop.netgate.com/products/netgate-installer
- Select an **Installation Image**:

  AMD64 Memstick USB
  : For 64-bit x86-64 Intel or AMD hardware to install via USB. This is the correct choice for most Netgate hardware and most third party hardware. This installer works with both serial and VGA consoles. Automatically detects known hardware which uses alternate serial console ports (e.g. Netgate 4200, ADI devices).

  AMD64 ISO
  : For 64-bit x86-64 Intel or AMD hardware to install via physical or virtual optical drive. This image works well for installing via IPMI and for virtual machines. This installer works with both serial and VGA consoles and like the memstick image, automatically adjusts the console settings to match known devices.

  AARCH64 Memstick ARM
  : For 64-bit ARM devices from Netgate, such as the Netgate 1100 and Netgate 2100.

- Click **Add To Cart**

  The installer is free, but uses the checkout process for delivery.

- Click **Enter Cart**
- Complete the checkout process and download the installation image.

> **Warning:** Be aware that Safari on macOS will, by default, automatically decompress downloaded files. This will make validating the integrity of the downloaded file impossible as the hashes are made against the compressed versions of the files. Disable this feature before proceeding or use an alternative browser.

See also:

Troubleshooting Installation Issues.

7.2.1 Verifying the integrity of the download

The integrity of the installer image can be verified by comparing a computed hash value of the original downloaded compressed file against a hash computed by Netgate when the files were created. The current hashes use SHA-256.

The pfSense Plus Installer Checksums page contains the SHA-256 sum of each current installer image. This is isolated from the image download for extra security. This file is plain text and can be opened in a text editor or viewed in a web browser.

Use the accompanying SHA-256 sum from the checksums page to verify that the download successfully completed and is an official release of pfSense software.

> **Warning:** The SHA-256 sums are computed against the compressed versions of the downloaded files. Compare the hash *before* decompressing the file.
>
> Safari on macOS will decompress files when downloading by default, making this validation impossible. Disable this feature before downloading or use an alternative browser.

Hash calculation programs vary by operating system, some common examples include:

Windows

PowerShell has a built-in Cmdlet `Get-FileHash` which can compute file hashes easily without needing to install additional software.

Example:

```
PS> Get-FileHash -Algorithm SHA256 .\netgate-installer-amd64.img.gz
Algorithm       Hash                                                   [...]
---------       ----
SHA256          0FEC506A48DE95A698F4DFE531018E1D7F847E440DA64E5482A6EEAC965F6B40
```

The SHA256 hash in the output can be compared with the value in the **contents** of the pfSense Plus Installer Checksums page.

Note: It is also possible to use the Linux `sha256sum` command within Windows Subsystem for Linux, Cygwin, or similar mechanisms.

Alternately, use a GUI-based hash calculation program such as OpenHashTab to compare the value against the provided hash. With OpenHashTab installed, right-click on the downloaded file to access the **File Hashes** tab containing the **SHA256** hash, among others.

Tip: If a SHA256 hash is not displayed, right click in the hash view and click **Settings**, then check the box for **SHA256** and click **OK**.

---

macOS

Use the `shasum` command line utility to generate a hash of the downloaded file.

Example:

```
$ shasum -a 256 netgate-installer-amd64.img.gz
0fec506a48de95a698f4dfe531018e1d7f847e440da64e5482a6eeac965f6b40  netgate-installer-
↪amd64.img.gz
```

The generated SHA256 hash can be compared with the contents of the pfSense Plus Installer Checksums page.

Linux

Use the `sha256sum` command line utility to generate a hash of the downloaded file.

Example:

```
$ sha256sum netgate-installer-amd64.img.gz
0fec506a48de95a698f4dfe531018e1d7f847e440da64e5482a6eeac965f6b40  netgate-installer-
↪amd64.img.gz
```

The generated SHA256 hash can be compared with the contents of the pfSense Plus Installer Checksums page.

FreeBSD

Use the `sha256` command line utility to generate a hash of the downloaded file.

Example:

```
$ sha256 netgate-installer-amd64.img.gz
SHA256 (netgate-installer-amd64.img.gz) =␣
↪0fec506a48de95a698f4dfe531018e1d7f847e440da64e5482a6eeac965f6b40
```

The generated SHA256 hash can be compared with the contents of the pfSense Plus Installer Checksums page.

7.3 Prepare Installation Media

The installation image downloaded in the previous section must first be transferred to the proper media. The files cannot be copied to media directly, but must be written using appropriate tools.

The primary difference between the USB memstick and ISO image is in how the images are written to an installation disk. Both types of images install pfSense® software to a target disk. Another difference is between the console types for the different USB memstick images. After installation, they each retain their appropriate console settings.

7.3.1 Decompress the Installation Media

The installation disk image is compressed when downloaded to save bandwidth and storage. Decompress the file before writing this image to an installation disk.

The `.gz` extension on the file indicates that the file is compressed with `gzip`. The image can be decompressed on Windows using 7-Zip, or on BSD/Linux/Mac with the `gunzip` or `gzip -d` commands.

7.3.2 Writing the Installation Media

Creating an installation disk requires a different procedure depending on the type of media. Follow the instructions in the appropriate section for the chosen media type.

Prepare a USB Memstick

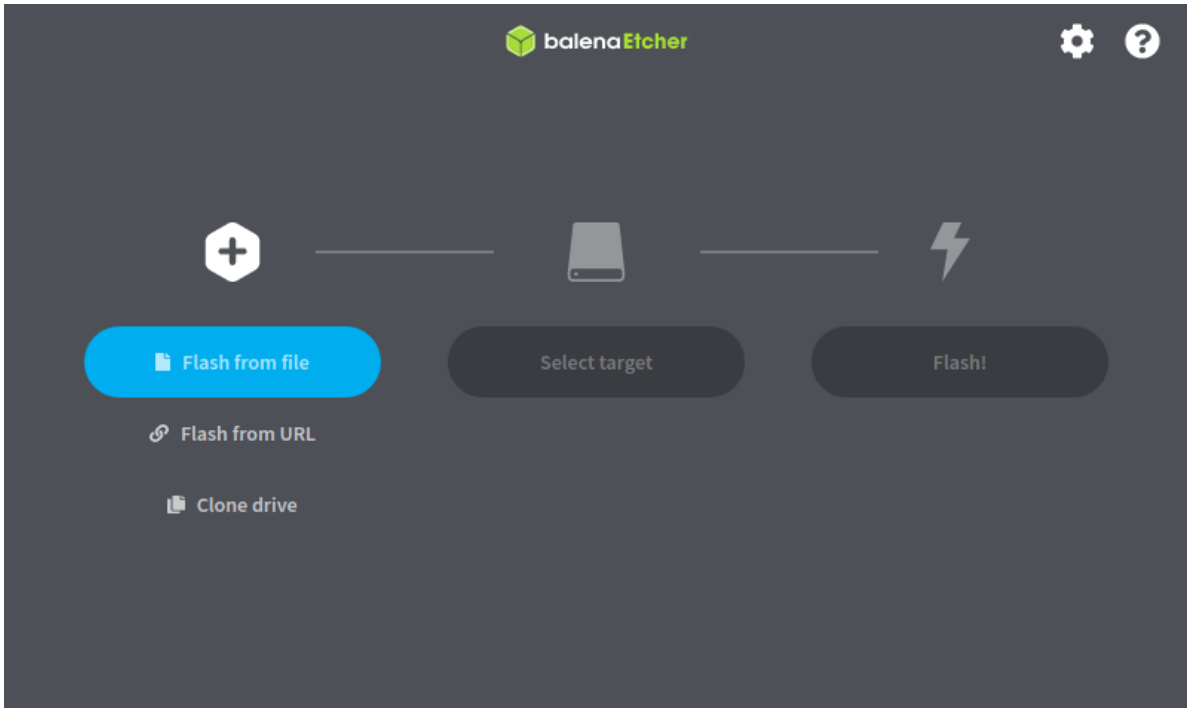
Warning: Be extremely careful when writing pfSense® software installation images! If the client PC contains other hard drives it is possible to select the wrong drive and overwrite a portion of that drive with the installer disk. This renders the disk completely unreadable except to certain disk recovery programs, if at all.

Using Etcher

The easiest way to create bootable installation media is to use Etcher. Etcher is available on Windows, macOS, and Linux so the procedure to write an image is the same across each of those platforms. Etcher is simple to use, can read compressed image files, and has several features which help prevent users from making unintentional mistakes in the process such as selecting the wrong target drive. Additionally, unlike other methods there is no need to perform other steps before writing the image to prepare the image file or disk.

- Download and install Etcher from https://www.balena.io/etcher/
- Insert a USB flash drive into the client computer
- Start Etcher

  Etcher will display its main screen as shown in the following image:

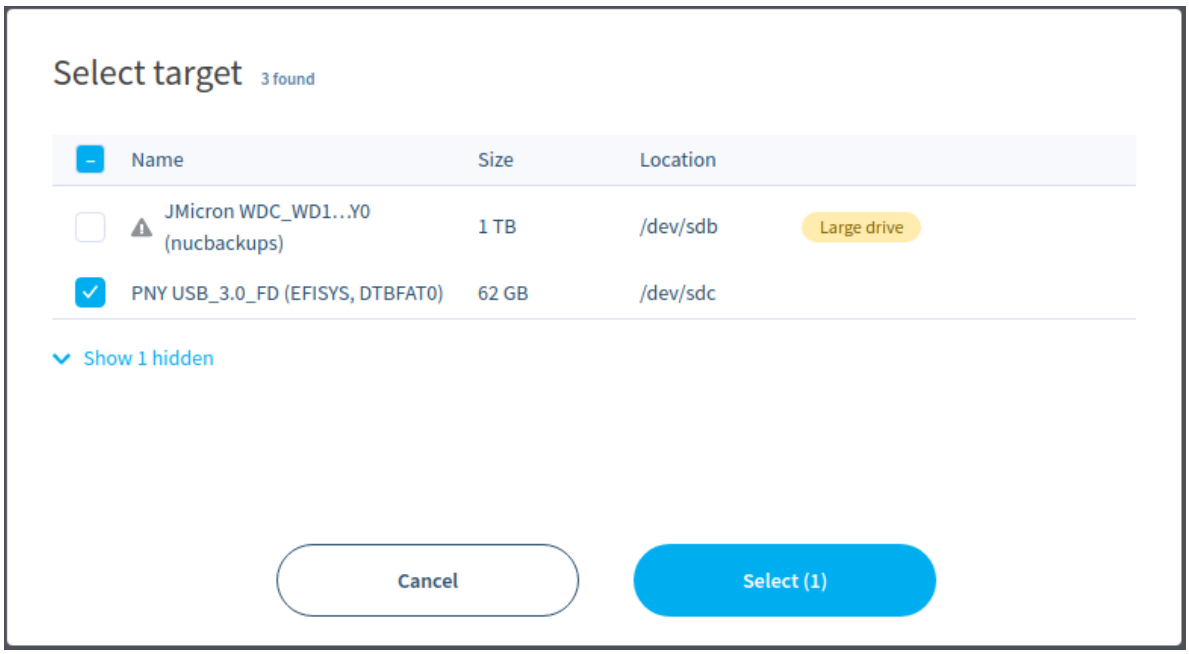


- Click **Flash from file**
- Locate and select the installation image file

Tip: Etcher can use compressed images directly, there is no need to manually decompress the image file first.

- Click **Select target**
- Click the USB flash drive to which Etcher should write the image

Note: Etcher attempts to hide and/or visibly mark potentially dangerous selections such as system drives, the drive containing the source image, and large drives. This makes it easier to identify the correct selection.



- Click **Select (1)** to continue
- Click **Flash!** to write the image to the target USB flash drive

  At this point there may be an authentication or UAC prompt to continue.

Note: Etcher requires elevated privileges to write USB drives. In the majority of cases, Etcher will trigger an operating system prompt for additional privileges as needed. If it does not, re-run Etcher as an administrator explicitly.

- Wait for the flash process to complete

  If there is an error from Etcher, try another USB flash drive or follow the advice given within Etcher to resolve the problem.

Warning: After writing the drive, the installation media will contain partitions which cannot be read by most operating systems. Ignore any operating system warnings about failing to mount the drive or prompts to format the drive.

- Close Etcher when complete
- Remove the USB flash drive from the client system

The installation media is now ready to use. Proceed to the installation instructions for the operating system.

Alternate Methods

For other techniques and additional guidance on writing disk images, see the reference document Writing an Installation Image to Flash Media.

Prepare a DVD

To use an ISO image file containing pfSense® software with an optical disk drive, the ISO image must be burned to a DVD by appropriate writing software.

Since the ISO image is a full-disc image, it must be burned appropriately for image files **not** as a data DVD containing the single ISO file. Burning procedures vary by OS and available software.

Decompress the ISO Image

Before the image can be burned, it must be decompressed. The `.gz` extension on the file indicates that it is compressed with `gzip`. This can be decompressed on Windows using 7-Zip, or on BSD/Linux/Mac with the `gunzip` or `gzip -d` commands.

Burn the DVD

Burning on Windows

Windows 7 and later include the ability to burn ISO images without extra software. On top of that, virtually every major DVD burning software package for Windows includes the ability to burn ISO images. Refer to the documentation for the DVD burning program. A Google search with the name of the burning software and `burn iso` also helps locate instructions.

Burning with Windows

To burn a disc image in Windows 7 or later:

- Open Windows Explorer and locate the decompressed ISO image file
- Right-click the ISO image file
- Click **Burn disc image**
- Select the appropriate **Disc burner** drive from the drop-down list
- Insert a blank DVD
- Click **Burn**

Later versions such as Windows 10 also show a **Disc Image Tools** tab on the ribbon when an ISO image is selected in Windows Explorer. That tab has a **Burn** icon that also invokes the same disc burning interface.

Other Free Burning Software

Other free options for Windows users include ISO Recorder, CDBurnerXP, InfraRecorder and ImgBurn. Before downloading and installing any program, check its feature list to make sure it is capable of burning an ISO image.

Burning on Linux

Linux distributions such as Ubuntu typically include a GUI DVD burning application that can handle ISO images.

If a DVD burning application is integrated with the window manager, try one of the following:

- Right-click on the decompressed ISO image file
- Choose **Open With**
- Choose **Disk image writer**

Or:

- Right-click on the decompressed ISO image file
- Choose **Write disc to**

Other popular applications include K3B and Brasero Disc Burner.

If a GUI burning program is not available, it may be possible to burn from the command line.

First, determine the burning device's SCSI ID/LUN (Logical Unit Number) with the following command:

```
$ cdrecord --scanbus
scsibus6:
     6,0,0   600) 'TSSTcorp' 'CDDVDW SE-S084C ' 'TU00' Removable CD-ROM
```

Note the SCSI ID/LUN is `6,0,0` in this example.

Burn the image as in the following example, replacing `<max speed>` with the speed of the burner (e.g. `24`) and `<lun>` with the SCSI ID/LUN of the recorder:

```
$ sudo cdrecord --dev=<lun> --speed=<max speed> netgate-installer-amd64.iso
```

Burning in FreeBSD

FreeBSD can use the same `cdrecord` options as Linux above by installing `sysutils/cdrtools` from ports or packages, and can also use GUI applications such as K3B or Brasero Disc Burner if they are installed from ports.

See also:

For more information on creating DVDs in FreeBSD, see the DVD burning entry in the FreeBSD Handbook.

Verify the Disc Content

After writing the disc, verify it was burned properly by viewing the files on the disc. More than 20 folders should be visible, including `bin`, `boot`, `cf`, `conf`, and more. If only one large ISO file is visible, the disc was not burned properly. Repeat the burning steps listed earlier and be sure to burn the ISO file as a DVD image and **not** as a data file.

7.4 Perform the Installation

This section describes the process of installing pfSense® software to a target drive, such as an SSD or HDD. In a nutshell, this involves booting from the installation memstick, ISO, or optical disc and then completing the installer.

This procedure uses the *Netgate Installer*.

Note: If the installer encounters an error while trying to boot or install from the installation media, see *Troubleshooting Installation Issues*.

7.4.1 Prerequisites

The following items are requirements to run the installer:

- *Download Installation Media*
- *Prepare Installation Media*
- *Connect to the Console*
- A network connection capable of reaching the Internet

  This installer is an online installer and requires Internet connectivity to download installation data from Netgate servers. The installer can use DHCP, static IP address, and PPPoE configurations for upstream connectivity. Connect the WAN port of the device into a live network connection compatible with one of those connectivity types.

See also:

Virtual environments may have additional requirements, see the following documents for examples:

- *Virtualizing with Proxmox® VE*
- *Virtualizing pfSense Software with Hyper-V*
- *Virtualizing pfSense Software with VMware vSphere / ESXi*

See also:

Hangouts Archive also covers a variety of relevant topics.

7.4.2 Booting the Installation Media

For USB memstick installations, insert the USB memstick and then power on the target system. The BIOS may require the disk to be inserted before the hardware boots.

For DVD installations, power on the hardware then place the CD into an optical drive.

Certain systems may need to be nudged to boot from the installer image in different ways. Typically, this involves pressing a hotkey during boot to bring up a boot menu, going into the BIOS to pick a boot device, or invoking a special command from a BIOS prompt.

Consult the Netgate Product Manuals for information on booting install media on various Netgate hardware. For third party hardware, check with the OEM.

Once the device boots from the installation media, the installer will launch automatically.

Specifying Boot Order in BIOS

If the target system does not boot from the USB memstick or CD, the most likely reason is that the given device was not found early enough in the list of boot media in the BIOS. Many newer motherboards have a one time boot menu invoked by pressing a key during POST, commonly `Esc` or `F12`.

Failing that, change the boot order in the BIOS. First, power on the hardware and enter the BIOS setup. The boot order option is typically found under a **Boot** or **Boot Priority** heading, but it could be anywhere. If booting from a USB or optical drive is not enabled, or has a lower priority than booting from a hard drive containing another OS, the hardware will not boot from the installer media. Consult the motherboard manual for more detailed information on altering the boot order.

7.4.3 Installing to the Target Drive

Serial Console Terminal Type

For installations using a serial console connection, the first prompt will ask for the terminal type to use for the installer. For PuTTY or GNU screen, `xterm` is the best type to use. The following terminal types can be used:

ansi
: Generic terminal with color coding

vt100
: Generic terminal without color, most basic/compatible option, select if no others work

xterm
: X terminal window. For modern terminal clients such as GNU screen, PuTTY, SecureCRT, Tabby, and other similar clients the `xterm` choice is most likely to produce the best looking output.

cons25w
: FreeBSD console style terminal

The installer assumes `cons25w` for VGA consoles.

Navigating the Installer

Once the installer launches, navigating its screens works as follows:

- To select items, use the arrow keys to move the selection focus until the desired item is highlighted.
- For installer screens containing a list, use the `up` and `down` arrow keys to highlight **entries** in the list.
- Use the `left` and `right` arrow keys to highlight an **action** button at the bottom of the screen such as **Select** and **Cancel**.
- Pressing `Enter` activates the selected action on the selected entry.

Performing the Installation

The installer contents are the same for both console types. The following document walks through the installation process in its entirety.

Installation Walkthrough

License Screen

When the installer starts the first screen it presents offers license terms for pfSense® software which the user must accept before installation.

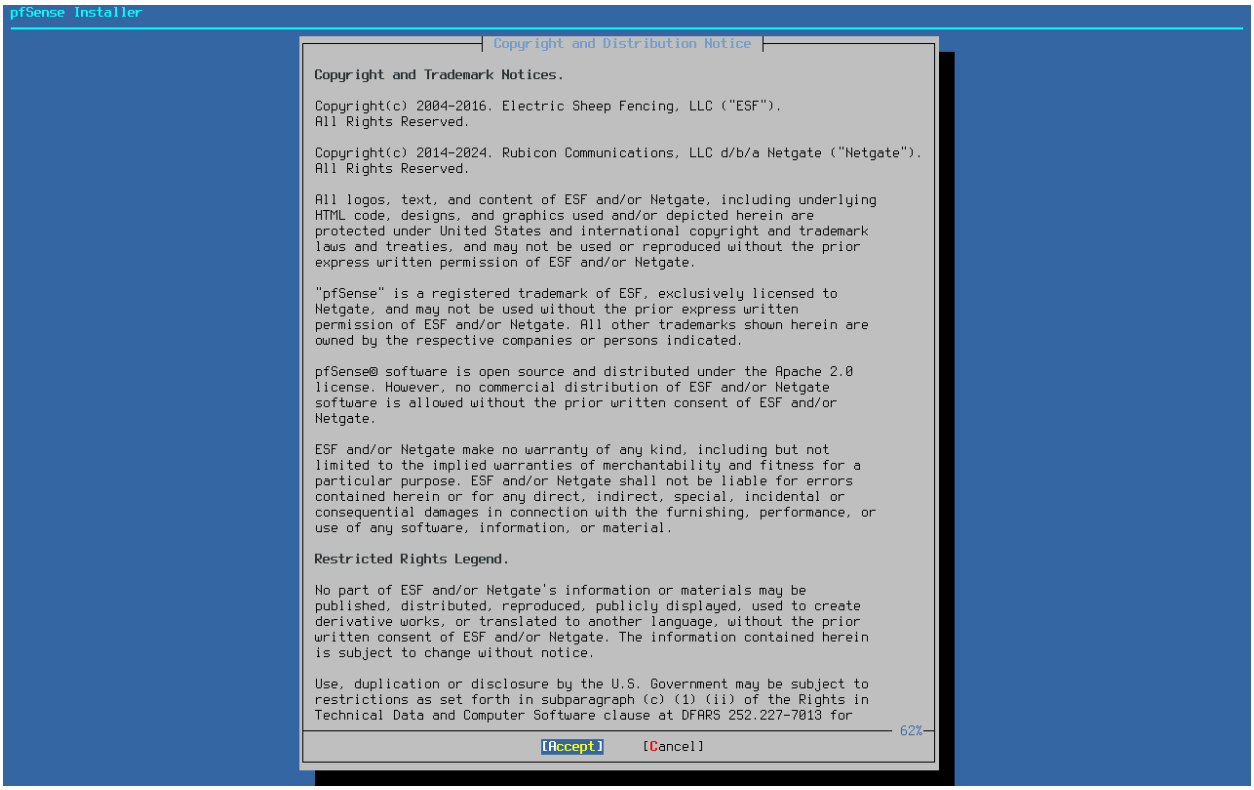


Fig. 1: Installer License

Read the terms carefully. Use the Page Down and Page Up keys to display additional license text. Press `Enter` to **Accept** the terms and proceed.

Welcome Menu

Next, the installer prompts to launch rescue options or start the installation process.

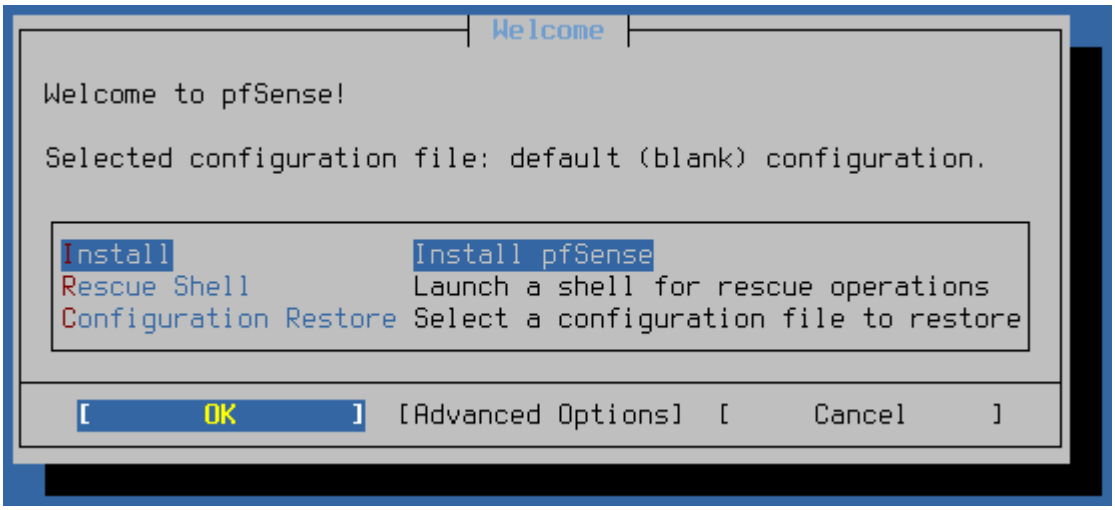


Fig. 2: Welcome Menu

Use the arrow keys to select an option, then press `Enter`. The options on this screen are:

Install
: Continue installing pfSense software

Rescue Shell
: Starts a basic shell prompt where advanced users can perform tasks to prepare the hardware in ways not possible directly in the installer, or to perform diagnostic tests or repairs on the firewall.

Configuration Restore
: Attempts to restore a configuration file recovered from a prior installation or copied from other media and then use that configuration in the target installation.

  The installer will hide this menu option if it cannot locate any configurations to restore.

  See *Configuration Restore* for details.

Advanced Options
: This option is in the bottom row of buttons. It loads another menu which contains extra options to control the behavior of the installer.

  See *Advanced Options* for details.

Configuration Restore

The installer searches for available configurations to recover and use for the target installation. This can be an existing prior installation of pfSense software or a configuration file on a FAT/FAT32 partition on a USB drive. The installer lists every configuration file it can locate and offers the user a choice of which to use, or to proceed without recovering a configuration.

Tip: When restoring a configuration from a prior installation, this option also searches for and copies SSH host keys and DHCP lease data to the new installation.

To recover a configuration and copy it to the target installation, use the arrow keys to select it from the list and press the `Enter` key.

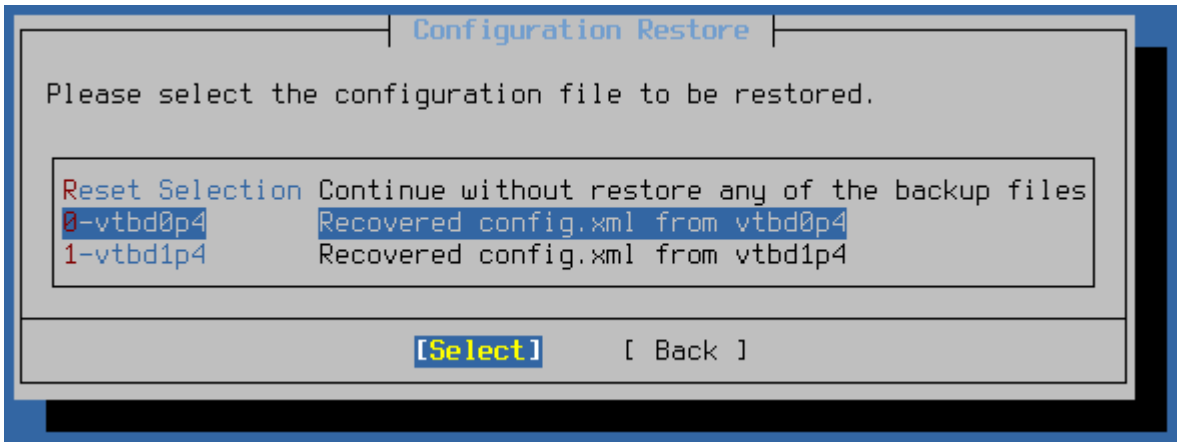


Fig. 3: Configuration Restore - List of Configuration Files

After selecting a configuration to restore, the installer displays this choice on the welcome screen:

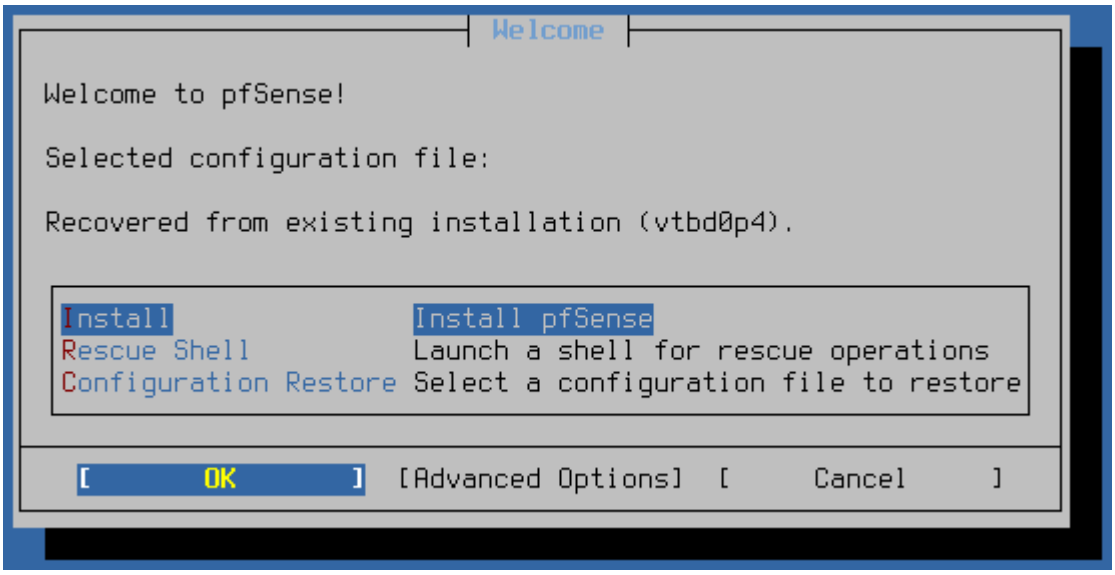


Fig. 4: Configuration Restore - Selected Configuration

To deselect the configuration file and proceed with a default configuration instead, enter the **Configuration Restore** menu again and choose **Reset Selection**.

Advanced Options

The options on the **Advanced Options** menu fine-tune the target installation.

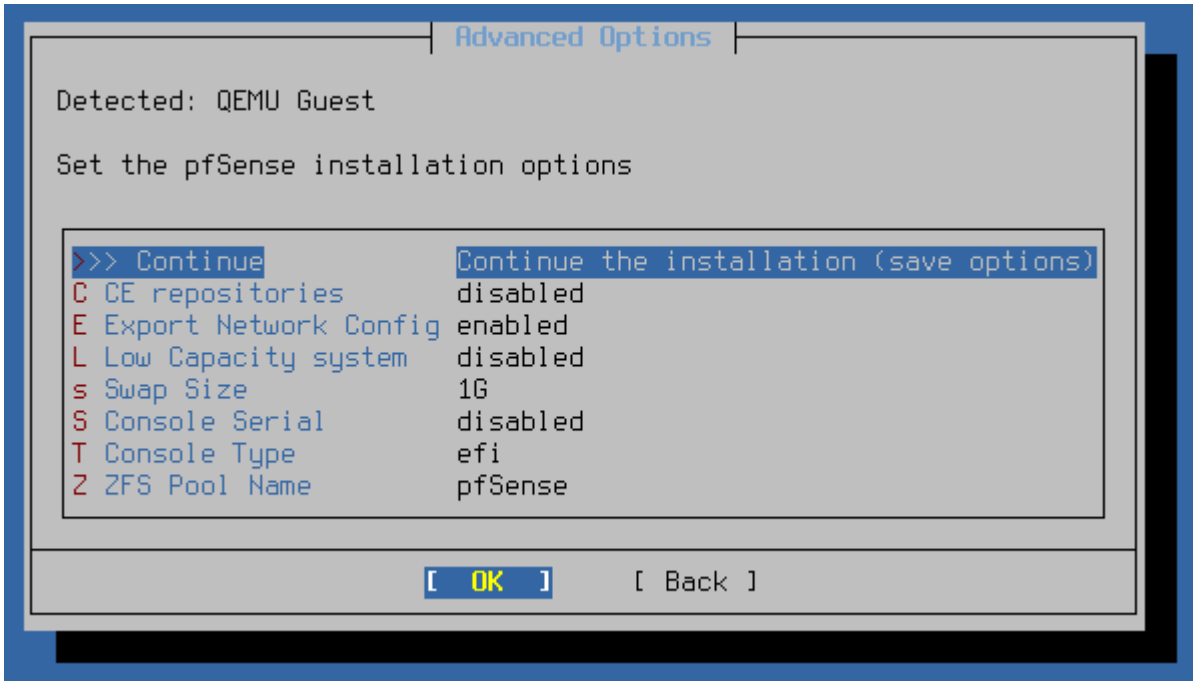


Fig. 5: Advanced Options

Use the arrow keys to select an option, then press `Enter` to set or toggle the value. The options on this screen are:

CE Repositories
: For devices eligible to install pfSense Plus software, this option toggles the availability of CE repositories in the list of versions the installer will offer.

  This allows someone with a device capable of running pfSense Plus software to install pfSense CE software instead.

Export Network Config
: This option instructs the installer to export its network configuration settings into the default configuration file for the target system during the installation process.

  For example, interface assignments along with the WAN and LAN configuration set inside the installer will carry over to the host for a smoother post-installation experience.

  The default value is **Enabled**. The installer automatically disables this option when recovering an existing configuration file.

Low Capacity System
: This option controls whether the installer automatically adjusts certain default options and behaviors when a device has low storage capacity (e.g. only contains disks with less than 4GB capacity).

  When enabled or set to automatic (default) and the installer detects a low capacity device, the installer will make the following changes:

  - Disables Swap
  - Sets the default file system to UFS
  - Will not configure EFI unless necessary

- Sets the default partition method to MBR

Swap Size
Sets the size of the swap partition the installer creates on the target disk.

Swap space is used for holding crash dump data as well as for virtual memory to supplement available RAM.

Enter a value with a size suffix, such as `1G` for 1 GiB of swap space. Use a value of `0` to disable swap.

Note: Swap usage can cause a higher volume of disk writes, but the best practice is to at least keep a small swap partition for crash dump data.

Console Serial
Controls whether the serial console will be enabled on the target installation. Toggles between `enabled` and `disabled`.

Console Type
Sets a specific type of console for the target installation.

EFI
EFI console, best suited for systems booting EFI with video and/or serial.

Video
Traditional VGA style console.

None
Do not set a specific console type.

ZFS Pool Name
The name the installer uses when it creates the ZFS pool on selected target disks. The default value is *pfSense*.

After setting options on this menu, choose **Continue** and **OK** and the installer will return to the *Welcome Menu*.

Network Setup

As this is an online installer it requires network connectivity to download installation packages from Netgate servers. To configure the network, the installer has to know at a minimum which port is a WAN with external connectivity, and configuration details to reach the Internet.

Note: The installer detects known models of Netgate hardware an automatically assigns the WAN and LAN to their default ports, skipping this manual assignment process and going right to *Confirm Network Configuration*.

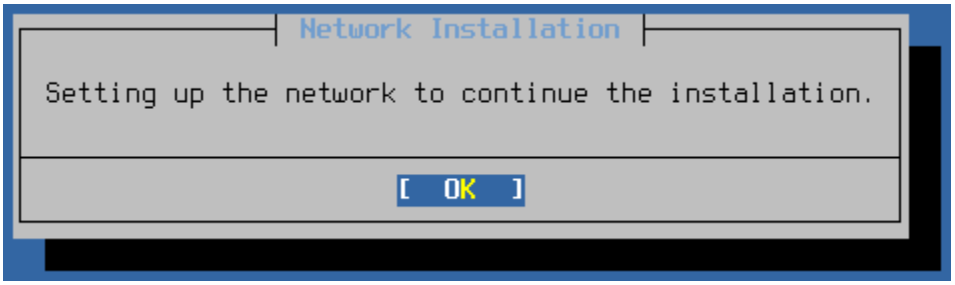


Fig. 6: Network Setup Prompt

Select WAN Interface

The first interface to assign is the WAN interface. This is the interface connected to the upstream network (e.g. Internet, modem, CPE, etc.). The installer presents a list of all detected interfaces and their MAC addresses, along with their current link state.

Use the up/down arrow keys to select the WAN interface and press `Enter` to continue.

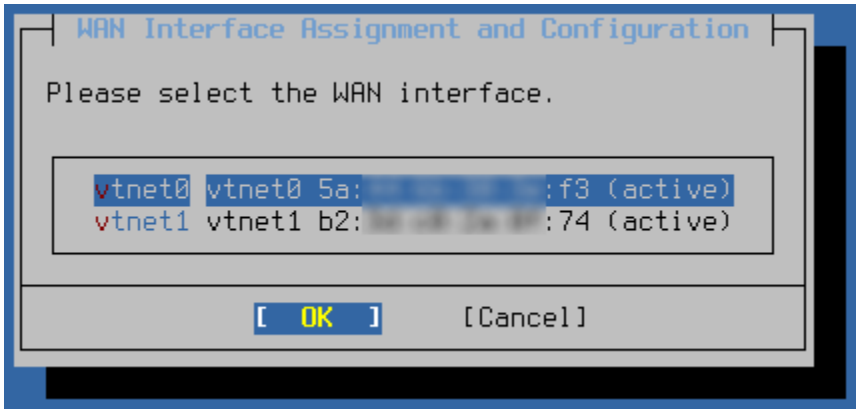


Fig. 7: Select WAN Interface

Note: When re-visiting this assignment screen later, for example to change the interface assignment or configuration, the list also includes the current assignment (e.g. WAN or LAN) at the end of each row.

Configure WAN Interface

The next step is to configure the WAN interface. The installer can utilize DHCP, static IP address, and PPPoE configurations for WAN interfaces. Additionally, interfaces may be VLAN tagged if necessary.

To change the type of interface configuration, select **Interface Mode** and press the `Enter` key. To configure a VLAN tag, select **VLAN Settings** and press the `Enter` key. To toggle use of the local resolver, select **Use local resolver**.

These options are explained in further detail in the following sections.

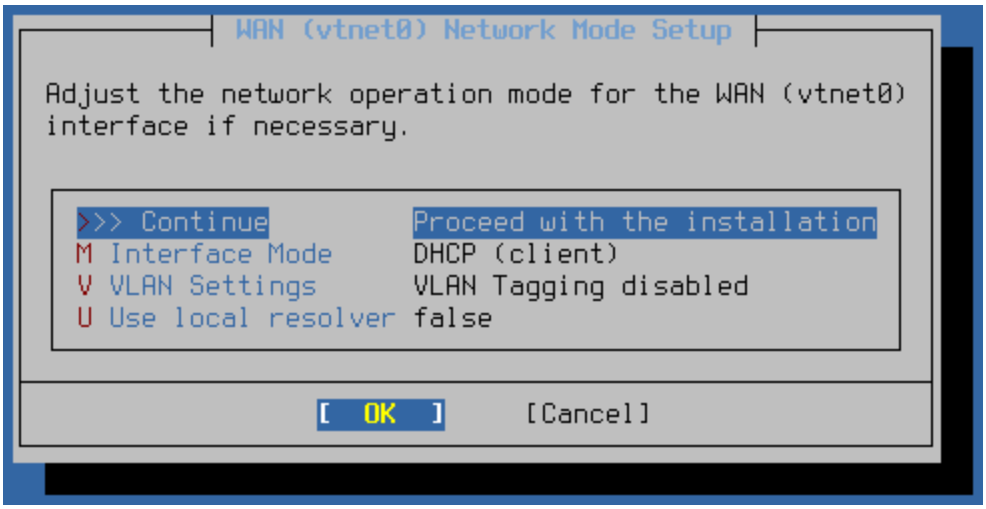


Fig. 8: WAN Interface Configuration

Interface Mode

The options on this screen change depending on the selected **Interface Mode** as certain types require additional configuration.

DHCP Client WAN

When the WAN interface is set to **DHCP (Client)** there are no additional options to configure, the behavior is automatic.

Static IP Address WAN

Changing the **Interface Mode** to **STATIC** presents several additional fields to configure static IP address WAN connectivity.

The available settings are:

IP Address
The IPv4 address and CIDR mask to use for external connectivity.

Note: The installer assumes a CIDR mask of `/24` if the address is given without a CIDR mask.

Default Gateway
The IPv4 address of the default gateway through which the installer can reach the Internet.

DNS Server
The IPv4 address of a DNS server, usually at the ISP or a public DNS server such as Google, CloudFlare, etc.

The figure above depicts a fully configured static IP address WAN.

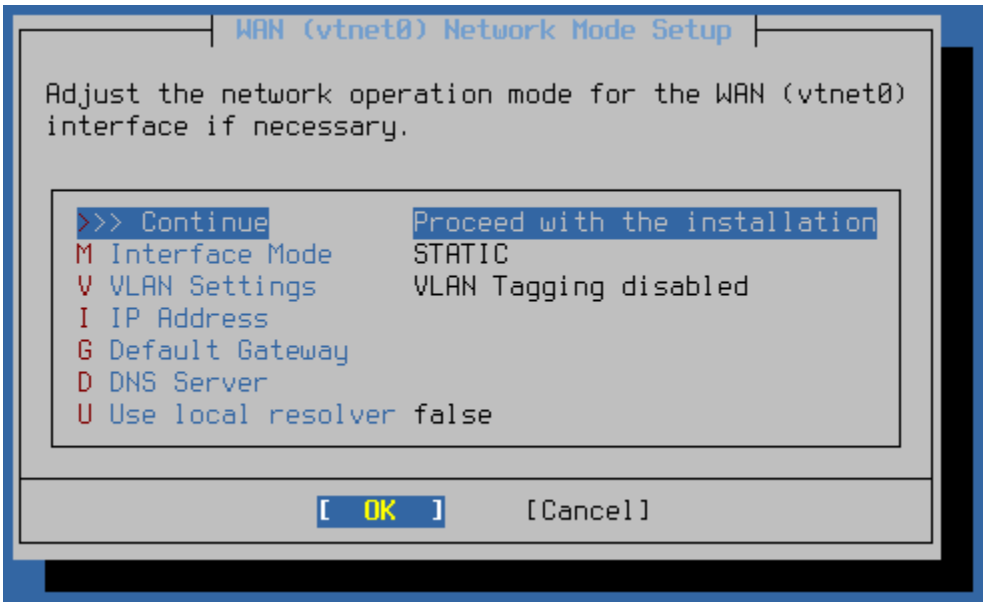


Fig. 9: Static IP Address

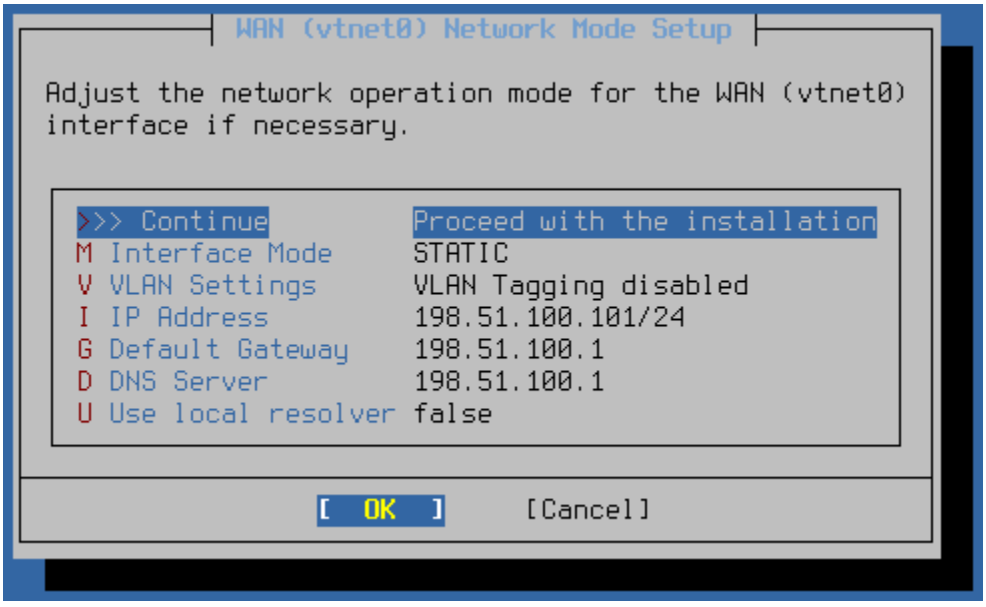


Fig. 10: Static IP Address (Configured)

PPPoE WAN

Changing the **Interface Mode** to **PPPoE** presents several additional fields to configure PPPoE WAN connectivity.

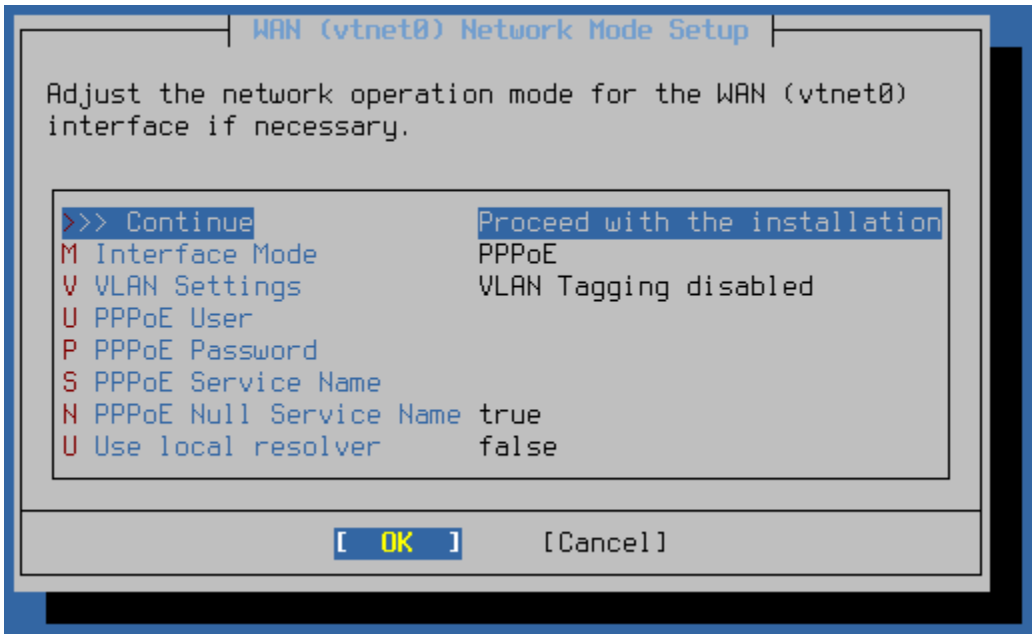


Fig. 11: PPPoE WAN

The available settings are:

PPPoE User
: The username to use when logging into the upstream PPPoE service.

PPPoE Password
: The password to use when logging into the upstream PPPoE service.

PPPoE Service Name
: Some PPPoE providers require a specific service name to be set in authentication requests. If the provider requires such a value, set it here. Otherwise, leave it empty.

PPPoE Null Service Name
: Configures the PPPoE client to send a null service name instead of an empty name when the **PPPoE Service Name** is empty. Certain providers may prefer one method or the other when they do not require a service name.

  Entering a **PPPoE Service Name** automatically sets this to **false**.

  Toggling this setting to **true** erases the **PPPoE Service Name**.

The figure above depicts a configured PPPoE WAN.

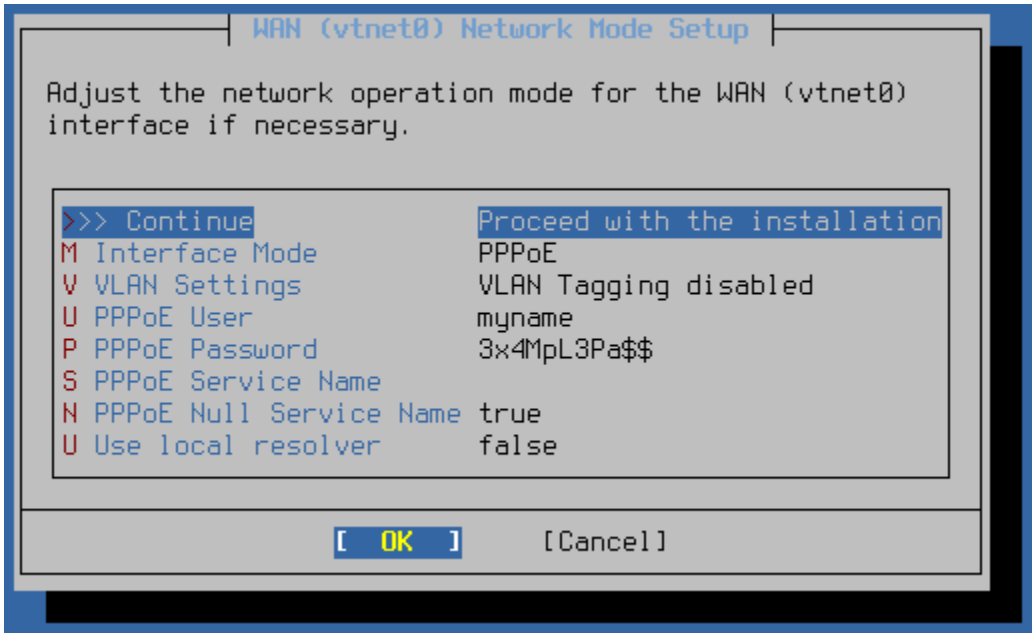


Fig. 12: PPPoE WAN (Configured)

VLAN Configuration

Each interface can be optionally configured to use a VLAN tag when communicating with the rest of the network connected to that interface.

To use a VLAN tag, first select **VLAN Settings** from the interface configuration screen to reach the VLAN settings screen.

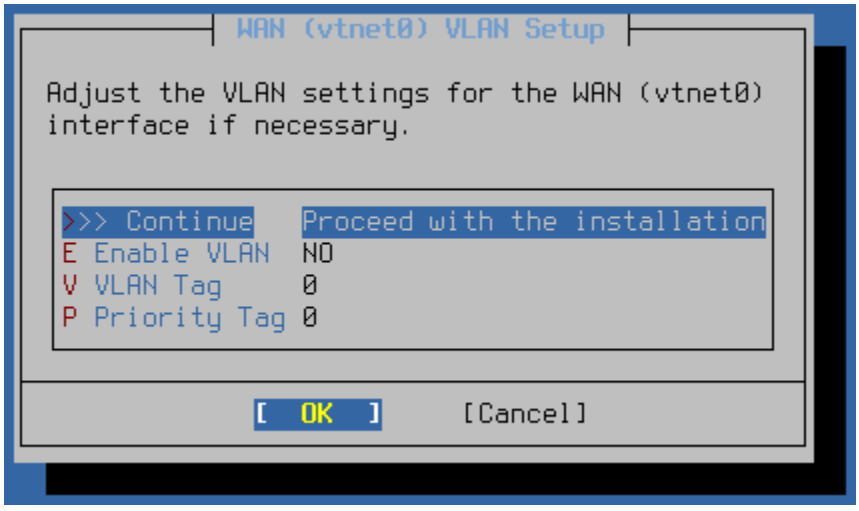


Fig. 13: VLAN Configuration

The **VLAN configuration** screen controls how installer uses VLANs on an interface. The following options are available:

Enable VLAN
Enables or disables VLAN tagging for the interface.

VLAN Tag
Sets the VLAN tag for traffic on the interface.

Priority Tag
Sets a VLAN priority value.

Select **OK** to return to the interface configuration.

Use Local Resolver

The **Use Local Resolver** option is present for every WAN type. It toggles the use of a local DNS resolver (Unbound) to handle DNS resolution rather than querying upstream DNS servers directly.

Select LAN Interface

The next step is to select the LAN interface. This is used for connecting to the installer from a local network if needed. While not used in this particular walkthrough, future installer features will rely on having a working LAN configuration, and it can also make obtaining information about installation problems easier to gather when obtaining assistance.

Selecting **None** will proceed without configuring a LAN, which is acceptable for installing from the console.

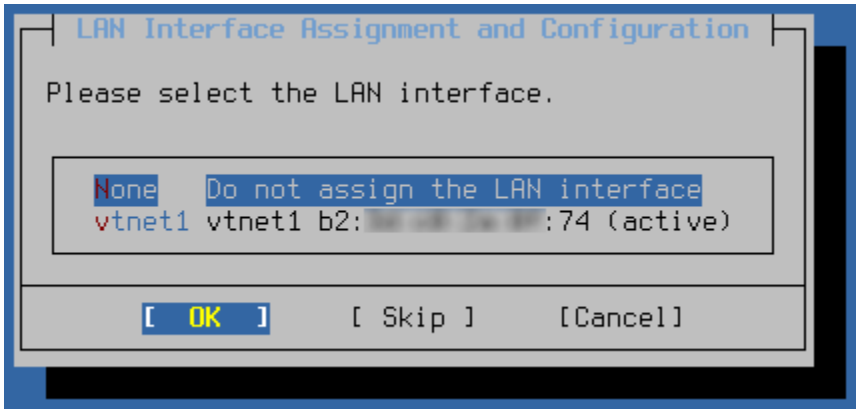


Fig. 14: Select LAN Interface if necessary

Note: When re-visiting this assignment screen later, for example to change the interface assignment or configuration, the list also includes the current assignment (e.g. WAN or LAN) at the end of each row.

Configure LAN Interface

The options to configure the LAN are similar to a WAN but not identical.

The following options are available when configuring the LAN interface:

Interface Mode
Select between DHCP Client and Static IP Address configuration types.

VLAN Settings
Enter *VLAN Configuration* mode for this interface.

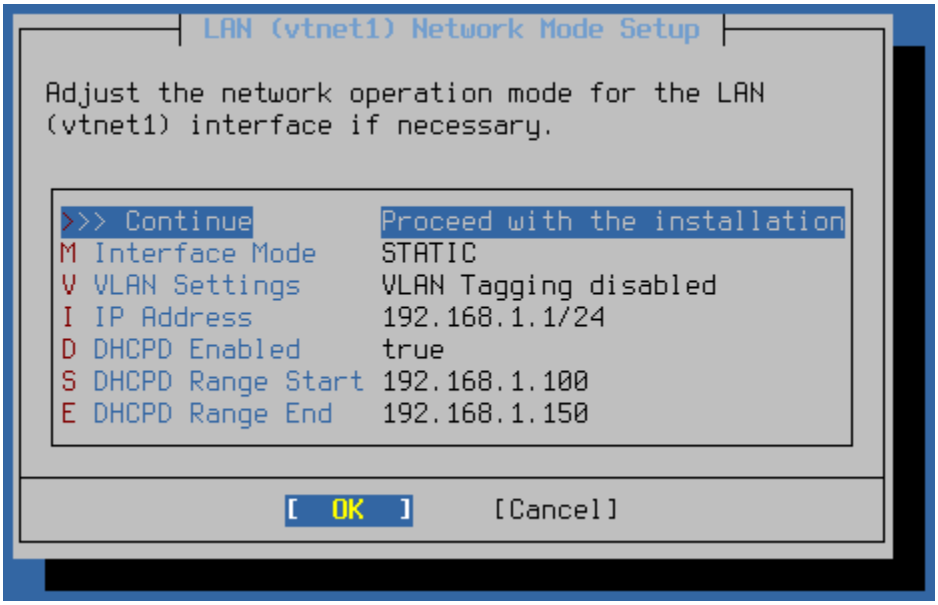


Fig. 15: LAN Interface Configuration

IP Address
Configure a static IP address and CIDR mask for the LAN. Default is `192.168.1.1/24`.

DHCPD Enabled
Toggles DHCP server behavior off/on (default: on)

Note: This option, along with the range start/end, are only available when LAN is set to a static IP address configuration.

DHCPD Range Start
Sets the starting address of the LAN DHCP range. Default is `192.168.1.100`.

DHCPD Range End
Sets the ending address of the LAN DHCP range. Default is `192.168.1.150`.

Confirm Network Configuration

This screen lists the current interface assignments, either after manual assignment or from being assigned automatically for known models of Netgate hardware.

If the default settings are OK, then choose to **Continue** from here by selecting it with the left/right arrows and pressing the `Enter` key.

The default settings are a DHCP client WAN, static IP Address LAN on `192.168.1.1/24` with DHCP server enabled on LAN from `192.168.1.100` to `192.168.1.150`.

To change the interface assignments or configuration, select the interface with the up and down arrows and then use the left/right arrows to highlight **Assign/Configure** then press the `Enter` key. Refer to the previous sections for information on how to assign and configure each interface.

At this point the installer should have Internet connectivity.

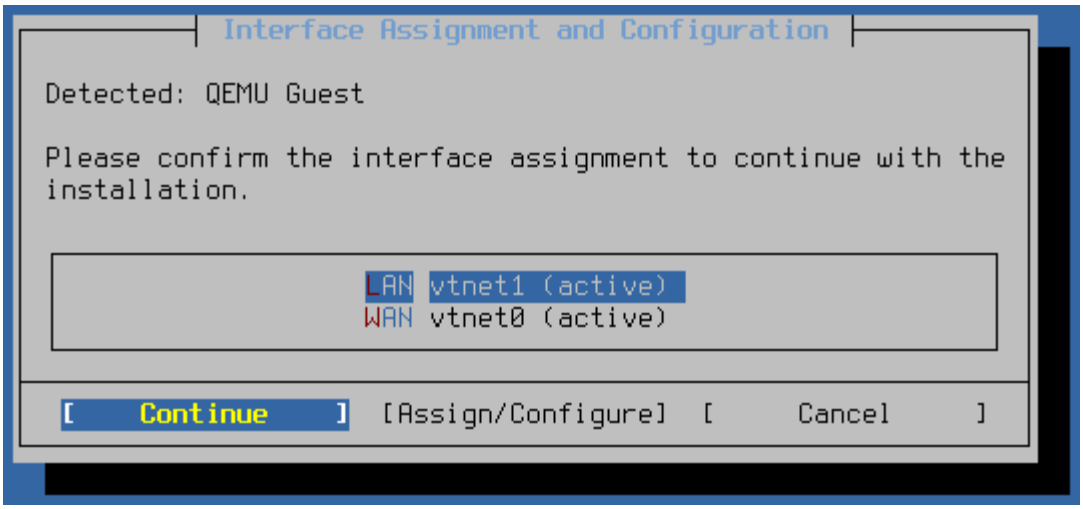


Fig. 16: Confirm Network Configuration

Ineligible Device Prompt

The installer gathers information about the device and communicates with Netgate servers to determine if the device is eligible to run pfSense Plus software. If the device is eligible, it moves forward to the *filesystem selection* screen. If the device is *not* eligible, the installer displays a prompt informing the user of this fact.

Warning: If the installer is unable to contact Netgate servers it will display an error saying "Cannot verify the eligibility of this system, please try again." For suggestions on how to correct that, see *Installer Network Connectivity Problems*.

If the device does not have an active subscription for pfSense Plus software, one can be purchased at this time by visiting https://www.netgate.com/purchase-plus and entering the Netgate Device ID (NDI), which is listed on this screen of the installer as well.

After subscribing, choose the **Retry Validation** option to allow the installer to check the subscription status again.

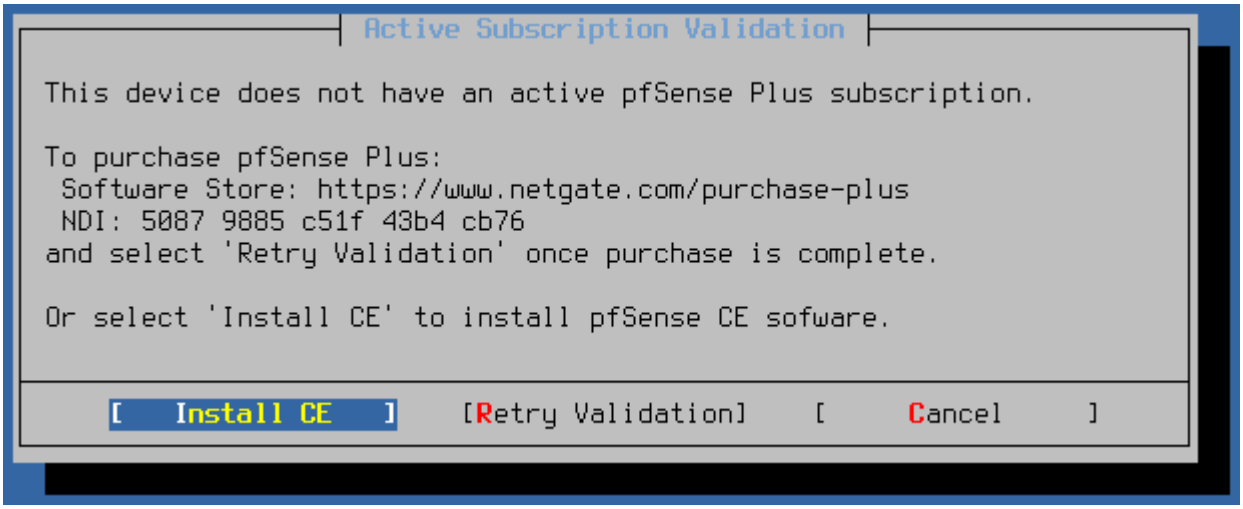


Fig. 17: Prompt displayed on systems not yet eligible to run pfSense Plus software

Alternately, users can choose the **Install CE** option to install pfSense CE Software, and that installation can upgrade to pfSense Plus software later after completing the subscription process.

File System and Partition Settings

After verifying the subscription, the next step is to choose the file system and partition type.

The available options are:

File System
The type of file system to use on the target disk.

ZFS
A robust modern file system with many advanced features, such as boot environments, but it uses a lot more resources. Even so, this is the default and best practice choice for nearly all cases.

UFS
An older file system that works well but can be fragile when it comes to sudden interruptions such as power loss. It uses less resources, but also doesn't have modern features such as boot environments.

Partition Scheme
The partition scheme to use on the target disk.

GPT
A modern partitioning method compatible with most modern AMD64 devices, but in rare cases it can have issues with older BIOS implementations. This is the default choice as there are very few incompatible modern devices.

MBR
A more basic partition scheme but one which is more widely compatible. This is also used on ARM-based systems.

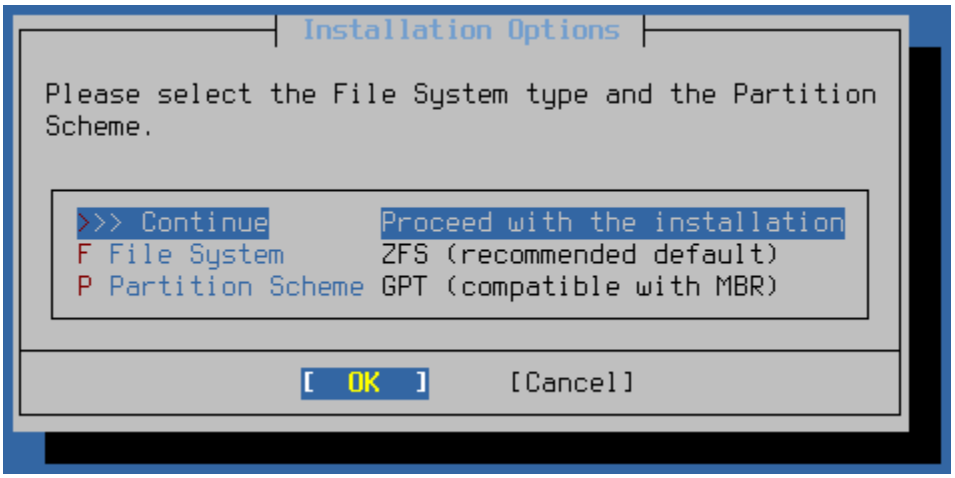


Fig. 18: File System and Partition Options

The process varies slightly depending on the selected filesystem type, so follow the section below that matches the filesystem type to be used by this firewall and then return to this document to complete the steps after.

ZFS

This section is a part of *Installation Walkthrough* and describes items specific to ZFS partitioning.

See also:

For information on tuning ZFS memory usage after installation, see *ZFS Tuning*.

Pool Type / Disks

When installing to ZFS the installer prompt to choose the **ZFS Configuration**. ZFS can utilize multiple disks in various ways for redundancy and/or extra capacity. Though using multiple disks with ZFS is software RAID, it is quite reliable and better than using a single disk.

The available types are:

stripe
A single disk, or multiple disks added together to make one larger disk (RAID 0).

Note: For devices with a single target disk, this is the correct choice.

---

mirror
Two or more disks that all contain the same content for redundancy. Can keep operating even if one disk dies. (RAID 1)

raid10
RAID 1+0, n x 2-way mirrors. A combination of stripes and mirrors, which gives redundancy and extra capacity. Can lose one disk from any pair at any time.

raidzX
Single, Double, or Triple redundant RAID. Uses 1, 2, or 3 parity disks with a pool to give extra capacity and redundancy, so either one, two, or three disks can fail before a pool is compromised. Though similar to RAID 5 and 6, the RAIDZ design has significant differences.

Select a type and press `Enter`

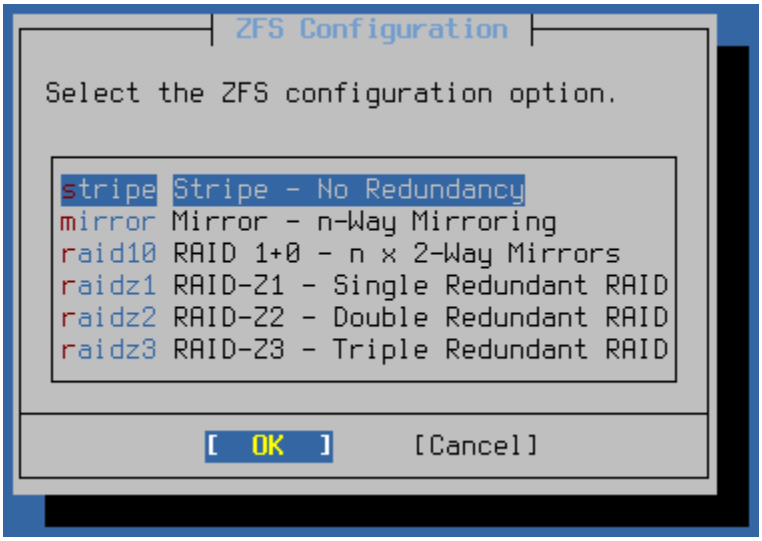


Fig. 19: ZFS Configuration Type

Select Disks

Next, the installer prompts for which disks it will include in the selected **ZFS Configuration**.

Use the `up` and `down` arrow keys to highlight a disk and `Space` to select disks. For mirrors or RAID types, select enough disks to fulfill the requirements for the chosen type.

Warning: Select a disk even if there is only one in the list!

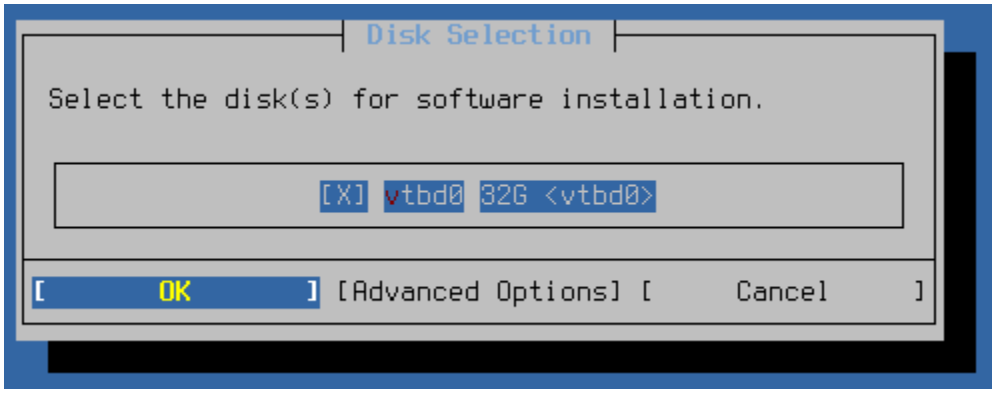


Fig. 20: ZFS Disk Selection

Note: If installer cannot find any drives, or if it shows incorrect drives, it is possible that the desired drive is attached to an incompatible controller or a controller set for an incompatible mode in the BIOS. See *Troubleshooting Installation Issues* for help.

See also:

The **Advanced Options** entry on this screen loads another menu containing additional disk operations. For information on these options, see *Advanced Disk Options*.

Continue Install

Proceed to *Final Confirmation*.

UFS

This section is a part of *Installation Walkthrough* and describes items specific to the UFS choices for partitioning.

When installing to UFS, the installer will prompt to select the target disk where the installer will write out the pfSense® software, e.g. `ada0`. The installer will show all compatible drives.

Note: Unlike ZFS, UFS can only utilize a single disk. Some setups, such as those using a RAID controller, may still use multiple disks, so long as they present a single virtual volume to the installer.

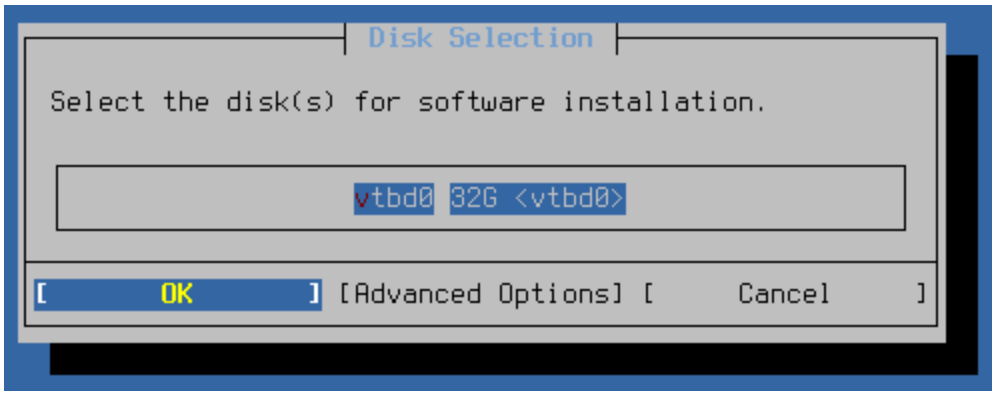


Fig. 21: UFS Disk Selection

Note: If installer cannot find any drives, or if it shows incorrect drives, it is possible that the desired drive is attached to an incompatible controller or a controller set for an incompatible mode in the BIOS. See *Troubleshooting Installation Issues* for help.

See also:

The **Advanced Options** entry on this screen loads another menu containing additional disk operations. For information on these options, see *Advanced Disk Options*.

Continue Install

Proceed to *Final Confirmation*.

Advanced Disk Options

When selecting disks for either UFS or ZFS the installer offers an **Advanced Options** selection with additional disk operations.

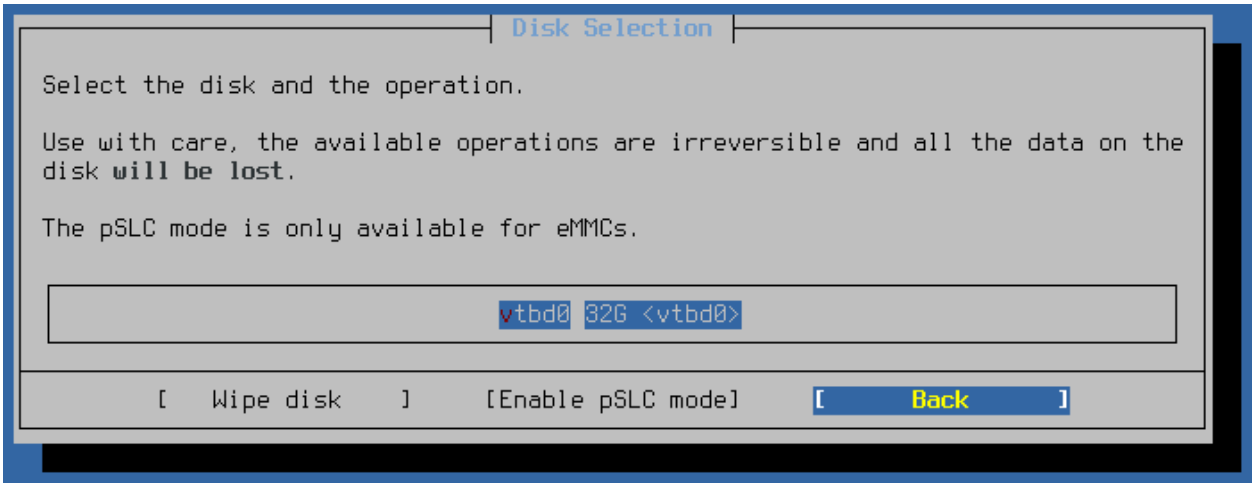


Fig. 22: Advanced Disk Options

Wipe Disks

This option clears partition and filesystem metadata from the selected disk, effectively making the operating system see the disk as empty.

On devices with both built-in and add-on storage, this is useful for making sure that a previous installation on an unused disk does not conflict with the new installation. For example, if a device has both eMMC storage and an NVMe SSD, the eMMC disk should be wiped, so the device only uses the SSD.

See also:

Troubleshooting Multiple Disks

Enable pSLC Mode

Some devices include eMMC disks with a "pseudo single-level cell" (pSLC) mode, which greatly increases write endurance at the cost of reduced usable capacity. Devices in pSLC mode can also be read twice as fast as triple-level cell (TLC) devices and written at six times the speed.

> **Danger:** Changing a disk to pSLC mode is **PERMANENT** and **CANNOT** be reversed. Proceed with caution.

> **Warning:** Disks can only change to pSLC mode if they are capable of using pSLC mode and have a low write count. The exact write count threshold varies by device.
>
> The installer attempts to set this mode on the disk, but there is no mechanism for the installer to detect pSLC mode compatibility or the write threshold.
>
> If the disk still shows the original capacity after enabling pSLC mode and power cycling the device, it either is not compatible with pSLC mode or its write cycle count was too high.

Switching a disk with multi-level cell (MLC) storage to pSLC mode increases its endurance from approximately 3,000 erasures per sector to 30,000, but the disk loses **half** of its usable capacity. For example, a 32GB MLC eMMC disk becomes a 16GB pSLC eMMC disk.

Switching a disk with triple-level cell (TLC) storage to pSLC mode increases its endurance from approximately 3,000 erasures per sector to 100,000, but the disk loses **two thirds** of its usable capacity. For example, a 32GB TLC eMMC disk becomes a 9GB pSLC eMMC disk.

Final Confirmation

After selecting the target disk the installer prompts for confirmation one final time before it makes destructive changes to the disk.

> **Danger:** Choosing to continue from this point will destroy anything left on the target disk!

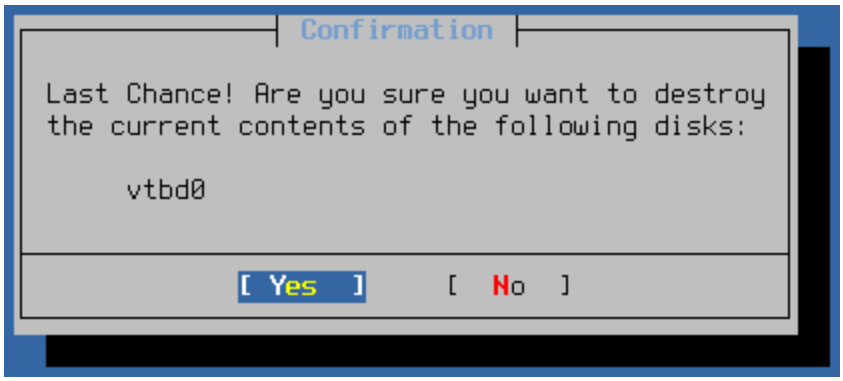


Fig. 23: Final Confirmation before Installing

Version Selection

At this point the installer presents a list of pfSense software that this device is eligible to run. This list will typically include the current version of pfSense software and one prior release. Depending on the current status of an upcoming release cycle, the installer may also offer development snapshots.

Select the version to install from the list with the up/down arrow keys, select OK with the left/right arrow keys, then press `Enter`

Tip: In most cases the correct selection will be the one labeled "Current Stable Version".

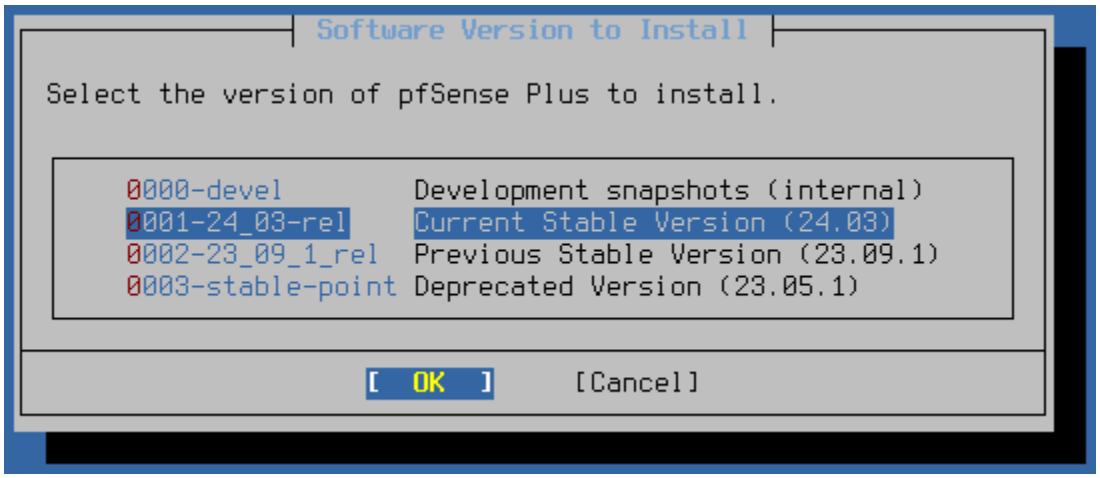


Fig. 24: Select Software Version to Install

Installation

After picking the version, the installer proceeds to download the installation data for that version and installs it on the target disk.

Sit back, wait, and have a few sips of a refreshing beverage while the installation process formats the drive(s) and copies pfSense software files to the target disk(s).

The installer displays the output from this process as it works. When finished, the installed presents an OK button which will continue to post-installation tasks.

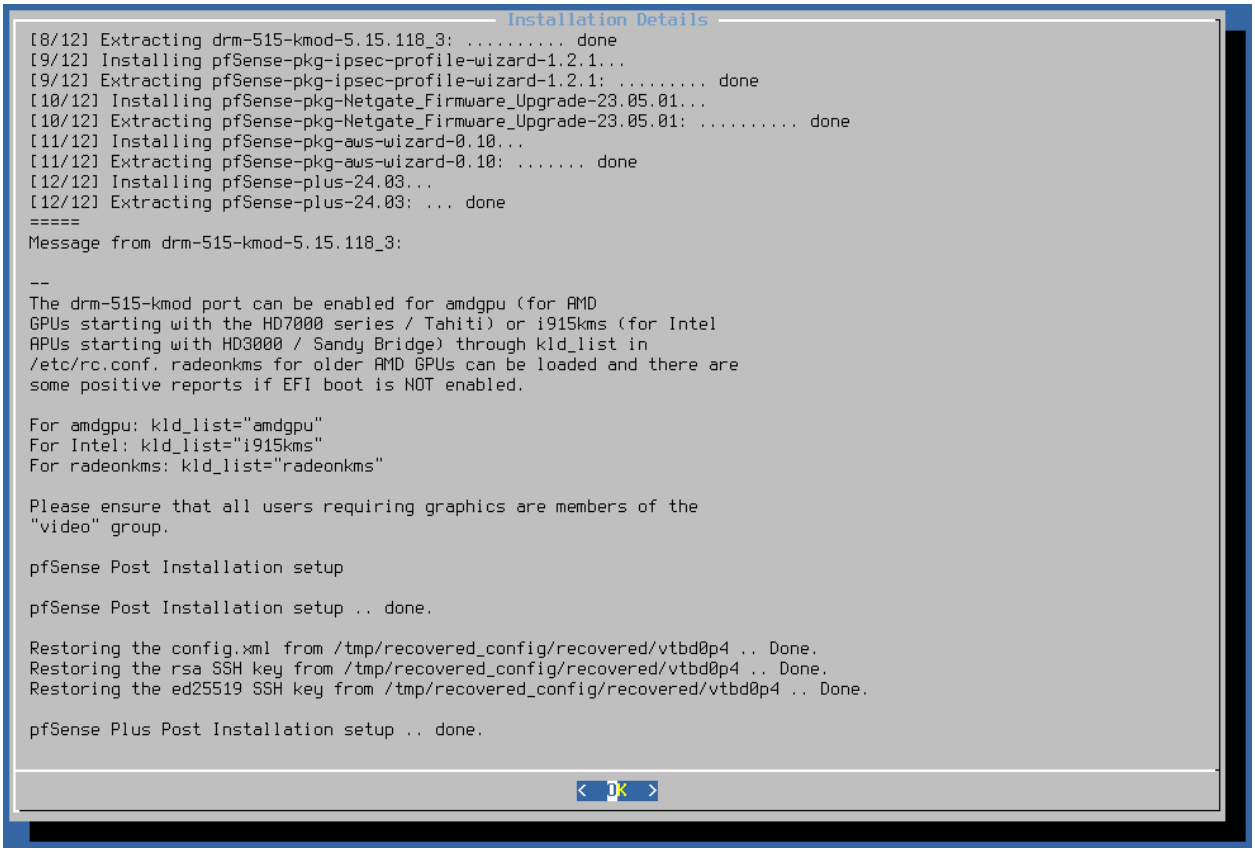


Fig. 25: Output After Installation Completes

Finish Up

At this point the installation is complete. The installer will prompt one final time to either reboot into the new installation or to start a shell prompt for any manual adjustments advanced users may wish to make.

Remove the installation media from the firewall during the reboot. The ideal time to remove it is while the hardware is initializing, but before it boots from the disk.

Once the device has booted from its own internal disk the device is ready for use.

Congratulations, the installation is complete!

The next step is to connect to the GUI and configure the device as described in *Configuration*.

See also:



Fig. 26: Reboot Prompt

- *Troubleshooting Installation Issues*
- *Copy Files to a USB Drive*
- *Alternate Remote Backup Techniques* (for an example of using SCP)

pfSense Software Default Configuration

After installation and interface assignment, pfSense software has the following default configuration:

- WAN is configured as an IPv4 DHCP client.
- WAN is configured as an IPv6 DHCP client and will request a prefix delegation.
- LAN is configured with a static IPv4 address of **192.168.1.1/24**.
- LAN is configured to use a delegated IPv6 address/prefix obtained by WAN (Track IPv6) if one is available.
- All incoming connections to WAN are *blocked* by the firewall.
- All outgoing connections from LAN are *allowed* by the firewall.
- The firewall performs NAT on IPv4 traffic leaving WAN from the LAN subnet
- The firewall will act as an IPv4 *DHCP Server*
- The firewall will act as an IPv6 *DHCPv6 Server* if a prefix delegation was obtained on WAN, and also enables SLAAC
- The *DNS Resolver* is enabled, so the firewall can accept and respond to DNS queries.
- SSH is disabled.
- GUI is running on port 443 using HTTPS.
- Default credentials are set as described in *Default Username and Password*.

7.5 Assign Interfaces

After the installer completes and the firewall reboots, the firewall software looks for network interfaces and attempts to assign interface mappings automatically.

The automatic interface assignment profiles used by the firewall are:

Netgate Hardware sold with pfSense® Plus Software
pfSense Plus software for devices from the Netgate Store includes default mappings appropriate to the hardware, which varies depending upon the hardware ordered with the device. Consult the Netgate Product Manuals for specific details on each model.

RCC-VE 4860/8860
WAN: igb1, LAN: igb0

RCC-VE 2220/2440
WAN: igb0, LAN: igb1

APU
WAN: re1, LAN: re2

Other Devices
For other devices the firewall looks for common interfaces and attempts to assign them appropriately, for example:

- WAN: igb0, LAN: igb1
- WAN: em0, LAN: em1
- WAN: re1, LAN: re2

If the firewall cannot automatically determine the network interface layout, it will present a prompt for interface assignment as in Figure *Interface Assignment Screen*. This is where the network cards installed in the firewall are given their roles as WAN, LAN, and Optional interfaces (OPT1, OPT2 … OPTn).

```
Valid interfaces are:

igb0   00:08:a2:09:95:b5   (up) Intel(R) PRO/1000 Network Connection, Version
igb1   00:08:a2:09:95:b6   (up) Intel(R) PRO/1000 Network Connection, Version
igb2   00:08:a2:09:95:b1 (down) Intel(R) PRO/1000 Network Connection, Version
igb3   00:08:a2:09:95:b2 (down) Intel(R) PRO/1000 Network Connection, Version
igb4   00:08:a2:09:95:b3 (down) Intel(R) PRO/1000 Network Connection, Version
igb5   00:08:a2:09:95:b4 (down) Intel(R) PRO/1000 Network Connection, Version

Do VLANs need to be set up first?
If VLANs will not be used, or only for optional interfaces, it is typical to
say no here and use the webConfigurator to configure VLANs later, if required.

Should VLANs be set up now [y|n]? n

If the names of the interfaces are not known, auto-detection can
be used instead. To use auto-detection, please disconnect all
interfaces before pressing 'a' to begin the process.

Enter the WAN interface name or 'a' for auto-detection
(igb0 igb1 igb2 igb3 igb4 igb5 or a):
```

Fig. 27: Interface Assignment Screen

The firewall displays a list of detected network interfaces and their MAC (Media Access Control) addresses, along with an indication of their link state if that is reported by the network card. The link state is denoted by `(up)` appearing after the MAC address if a link is detected on that interface.

Note: The Media Access Control (MAC) address of a network card is a unique identifier assigned to each card, and no two network cards *should* have the same MAC address. If a duplicate MAC address is present on a network, either by chance or by intentional spoofing, all conflicting nodes will experience connectivity problems.

After printing the network interface list, the firewall prompts for VLAN configuration. If VLANs are desired, answer `y`, otherwise, type `n`, then press `Enter`.

See also:

For information about configuring VLANs, see *Virtual LANs (VLANs)*.

The firewall prompts to set the WAN interface first. As the firewall typically contains more than one network card, a dilemma may present itself: How to tell which network card is which? If the identity of each card is already known, enter the proper device names for each interface. If the difference between network cards is unknown, the easiest way to figure it out is to use the auto-detection feature.

For automatic interface assignment, follow this procedure:

- Unplug all network cables from the firewall
- Type `a` and press `Enter`
- Plug a network cable into the WAN interface of the firewall
- Wait a few moments for the firewall to detect the link-up event
- Press `Enter`

If all went well, the firewall can determine which interface to use for the WAN.

Repeat the same process for the LAN and optional interfaces, if any are necessary. If the firewall prints a message stating "No link-up detected", see *Manually Assigning Interfaces* for more information on sorting out network card identities.

Once the list of interfaces for the firewall is correct, press `Enter` at the prompt for additional interfaces. The firewall will ask **Do you want to proceed (y|n)?** If the network interface assignment list is correct, type `y` then press `Enter`. If the assignment is incorrect, type `n` and press `Enter` to repeat the assignment process.

Note: In addition to the normal routing/firewall mode with multiple interfaces, a firewall may also run in **Appliance Mode** where it has only a single interface (**WAN**). The firewall places the GUI anti-lockout rule on the WAN interface, which allows clients to access the firewall web interface from that network. The usual routing and NAT functions are not active in this mode since there is no internal interface or network. This type of configuration is useful for VPN appliances, DHCP servers, and other stand-alone roles.

7.5.1 Manually Assigning Interfaces

If the auto-detection feature did not work, there is still hope of telling the difference between network cards prior to installation. One way is by MAC address, which the firewall prints next to the interface names on the assignment screen:

```
vmx0    00:0c:29:50:a4:04
vmx1    00:0c:29:50:ec:2f
```

The MAC address is sometimes printed on a sticker somewhere physically on the network card. For virtualized systems, the virtual machine configuration usually contains the MAC address for each network card. MAC addresses are assigned by manufacturer, and there are several online databases which offer reverse lookup functionality for MAC addresses in order to find the company which made the card: http://www.8086.net/tools/mac/, http://www.coffer.com/mac_find/, and http://aruljohn.com/mac.pl, among many others.

Network cards of different makes, models, or sometimes chipsets may be detected with different drivers. It may be possible to tell an Intel-based card using the `igb` driver apart from a Broadcom card using the `bge` driver by looking at the cards themselves and comparing the names printed upon the circuitry.

The probe order of network cards can be unpredictable, depending on how the hardware is designed. In a few cases, devices with numerous ports may use different chipsets that probe in different ways, resulting in an unexpected order. Add-on and Multi-port NICs are generally probed in bus order, but that can vary from board to board. If the hardware has onboard NICs that are the same brand as an add-in NIC, be aware that some systems will list the onboard NIC first, and others will not. In cases when the probe order makes multiple NICs of the same type ambiguous, it may take trial and error to determine the port placements and driver name/number combinations.

After the network cards have been identified, type the name of each card at the interface assignment screen when prompted. In the above example, `vmx0` will be WAN and `vmx1` will be LAN. To assign them these roles, follow this procedure:

- Type `vmx0` and press `Enter` when prompted for the WAN address
- Type `vmx1` and press `Enter` when prompted for the LAN address
- Press `Enter` again to stop the assignment process, since this example does not contain any optional interfaces.
- Type `y` and press `Enter` to confirm the interface assignments

7.6 Upgrade Guide

Users can reliably upgrade pfSense® software from an older release to a current release.

Netgate periodically releases new versions of pfSense software which contain new features, updates, bug fixes, and various other changes. In most cases, updating an installation is easy. If the firewall is updating to a new release that is only a point release (e.g 2.x.3 to 2.x.4), the update is typically minor and unlikely to cause problems.

Note: Netgate TAC can only officially assist with the most recent stable releases of pfSense software, so updating is also important to ensure that TAC can help resolve problems encountered by customers.

Upgrades use the same software edition that the firewall is currently running. For example, pfSense CE software installations will upgrade to the latest version of pfSense CE software. pfSense Plus software will upgrade to the latest version of pfSense Plus software. The only exception to this is when following the special procedure to *Migrate from pfSense® CE software to Netgate pfSense Plus software*.

The most common problems users encounter during upgrades are hardware-specific regressions from one FreeBSD version to another, though those are uncommon. New software releases fix more hardware than they break, but regressions are always possible. Larger jumps, such as from 2.3.x to 2.8.0-RELEASE must be handled with care, and ideally tested on identical hardware in a test environment prior to use in production.

Warning: Firewalls must be connected to the Internet to update.

- *Update Settings*
  - *Branch / Tracking Snapshots*
  - *Boot Environments*
  - *Dashboard Check*
  - *GitSync*

7.6.1 Upgrade Process Overview

This is a brief overview of the process that takes place behind the scenes when an administrator initiates an upgrade. The process varies depending on the version of pfSense software and whether the firewall is capable of utilizing ZFS Boot Environments. Unless otherwise stated, all steps are performed by the firewall automatically.

See also:

- *ZFS Boot Environments (Plus Only)*
- *Boot Environments Settings*

Traditional Upgrade Process (CE, Plus with UFS)

This is the traditional upgrade process using `pkg`, which is used by pfSense CE software as well as installations of pfSense Plus software installed using UFS:

- Download update files
- Upgrade kernel
- Reboot (Offline downtime starts)
- Upgrade base files (Operating System, PHP code, etc.)
- Upgrade base and add-on packages
- Finish boot (Offline downtime ends, back online and ready)

ZFS Boot Environments (First Generation, Plus 22.05-23.09.1)

This is the first generation of ZFS Boot Environments, available from pfSense Plus software version 22.05 until 23.09.1:

- Create and activate new ZFS Boot Environment from current point
- Download update files
- Upgrade kernel
- Reboot (Offline downtime starts)
- Upgrade base files (Operating System, PHP code, etc.)
- Upgrade base and add-on packages
- Finish boot (Offline downtime ends, back online and ready)

This process essentially makes an automatic local backup before starting so that administrators can manually roll back if there are problems with the upgrade.

ZFS Boot Environments (Second Generation, Plus 24.03 and later)

This is the second generation of ZFS Boot Environment upgrade behavior, which started on pfSense Plus software version 24.03:

- Create new ZFS Boot Environment from current point
- Download update files
- Upgrade kernel inside the new Boot Environment
- Upgrade base files (Operating System, PHP code, etc.) inside the new Boot Environment
- Upgrade base and add-on packages inside the new Boot Environment
- Check for errors inside the new Boot Environment and stop if any problems have occurred
- Activate the new Boot Environment
- Reboot (Offline downtime starts)
- Monitor for problems during reboot
  - If problems are detected, automatically activate the last known good Boot Environment and reboot
- Finish boot (Offline downtime ends, back online and ready)

This method involves less downtime as the only downtime is during the reboot and the bulk of the upgrade process is already complete by that point. It is more robust since it can automatically detect problems and roll back without manual intervention if need be, which makes remote upgrades safer. Earlier ZFS Boot Environments are still available to roll back manually if users encounter problems not detected automatically.

7.6.2 Pre-Upgrade Tasks

Make a Backup ... and a Backup Plan

Before making any modifications to a firewall, the best practice is to make a backup using the GUI:

- Navigate to **Diagnostics > Backup/Restore**
- Set the **Backup Area** to *All* in the **Backup Configuration** section of the page
- Click  **Download**
- Save this file somewhere safe

Keep multiple copies of the backup file in different secure locations. Consider using the free **Auto Config Backup** service (*Automatic Configuration Backup Service*). Auto Config Backup can create a manual backup with a note identifying the change, which is encrypted and stored on Netgate servers.

Another good practice is to have a copy of the Netgate installer on-hand, in case something goes awry and reinstallation is necessary. Should that happen, have the backup file on hand and refer to *Backup and Recovery*.

Check and Clean Up ZFS Boot Environments

Systems running pfSense Plus software installed using ZFS will automatically create *ZFS Boot Environment* entries during the upgrade process as a safety measure. These entries allow users to easily boot back into the previous version if necessary. These entries consume space which presents itself as shrinking disk capacity and less free space on the disk. Eventually if left unchecked there may not be sufficient room to contain an upgrade without removing older boot environments.

See also:

See *Boot Environment Disk Space Usage* for more information about ZFS Boot Environment space usage.

Tip: The amount of space consumed by a boot environment varies depending on how much the disk contents change during and after an upgrade. It is safer to estimate on the high side and consider it may use at least 1.5-2GB for a major upgrade. Upgrades that involve fewer changes on the file system will use less disk space.

Before any upgrade, navigate to **System > Boot Environments** and review the list. Remove older entries which are no longer necessary (*Managing Boot Environments in the GUI*). Typically, users keep the latest 2-3 entries, but preferences vary.

Low Memory Hardware and AWS/Azure Instances

Hardware with **1 GiB or less** available memory may have issues upgrading depending on which features, services, or packages are running. This includes some Netgate hardware such as the Netgate 1100 when running with ZFS and/or certain services/packages. For the best chance of success in these cases, temporarily disable any non-critical services before starting the upgrade.

Tip: A *Pre-Upgrade Reboot* can also temporarily reduce memory used for ZFS caching, which can help in this situation as well.

---

pfSense Plus software can no longer run on AWS ".nano" size instances as they lack sufficient RAM to upgrade properly. Attempting to upgrade a ".nano" instance to pfSense Plus software version 24.03 will fail before the upgrade is performed. Migrate the instance to a ".micro" or larger size **before** attempting to upgrade, or redeploy instead.

Similar to the above, pfSense Plus software can no longer run on Azure A0 instances. Migrate to instances with more memory.

Remove / Dismount any Installation Media

Some users leave an installation disk plugged in or mounted for various reasons, but this can interfere with the upgrade process. Be sure to dismount and remove any installation media such as a USB thumb drive, optical disk, or ISO in a virtual drive. This includes items mounted via Hypervisors/Virtual Machine emulated or passed-through hardware, IPMI virtual media, and other similar mechanisms.

VM Snapshots

An easy fall-back plan for virtualized firewalls is to take a snapshot of the VM before performing an upgrade. This way, it can easily roll back to a known-good state if the VM encounters a problem.

Note: Before rolling back a VM due to problems, ensure the hardware compatibility of the VM is current and update the VM Guest operating system to match the upgraded OS if there is a matching choice in the Hypervisor.

Pre-Upgrade Reboot

Reboot the firewall before applying an update. This step is optional, but a best practice.

If the hardware has a problem, such as a disk issue, then performing a reboot before the upgrade will allow that to be identified early. Otherwise, a hardware issue could be confused with an issue that occurred as a result of the upgrade process.

If the installation is using ZFS, a reboot will reset the amount of memory the OS is using for ZFS caching as well, which can help ensure the upgrade runs smoothly.

There is still a chance that the upgrade could draw out a hardware issue, such as a disk failing from the writes that happen in the upgrade process, but that is much less common to see in practice.

Packages

Warning: When the firewall is configured to pull packages from a release newer than the one current running, **Do not upgrade packages before upgrading pfSense® software**. Either remove all packages or leave the packages alone before running the update.

The safest practice is to remove **all** packages before upgrading to a new release. The upgrade process will handle packages automatically, but packages are frequently a source of problems. To ensure a smooth upgrade, note the installed packages, remove them, perform the upgrade, and then reinstall when the upgrade is complete.

Danger: Installing packages from repositories other than those provided by Netgate is **not supported** and carries a **significant** risk of breaking the installation.

Packages installed from manually configured non-Netgate repositories may result in failed upgrades or unbootable systems. Even if the third party repository configurations and packages are removed, packages from those repositories may have already damaged the base system in unknown ways. If an administrator chooses to ignore these warnings and use such packages, they must take extra caution and have installation media ready in case the device is unable to upgrade or reboot successfully.

7.6.3 Version-Specific Notes

This document covers specific concerns which must be taken into account by administrators when upgrading pfSense® software from an older version.

Review sections for each intermediate version between the version running on the firewall and the current release.

- *Upgrading from versions older than pfSense Plus 21.02.2 or pfSense CE 2.5.1*
- *Upgrading from versions older than pfSense 2.5.0*
- *Upgrading from versions older than pfSense 2.4.5-p1*
- *Upgrading from versions older than pfSense 2.4.4*
- *Upgrading from versions older than pfSense 2.4.0*
- *Older Version Upgrade Notes*

Upgrading from versions older than pfSense Plus 21.02.2 or pfSense CE 2.5.1

Warning: WireGuard was removed from the base system in releases after pfSense Plus 21.02-p1 and pfSense CE 2.5.0, when it was removed from FreeBSD.

If upgrading from a version that has WireGuard active, the upgrade will abort until all WireGuard tunnels are removed. For more details, see the *Release Notes*

WireGuard is available as an add-on package on pfSense Plus 21.05, pfSense CE 2.5.2, and later versions. The settings for the WireGuard add-on package are not compatible with the older base system configuration.

Note: The WireGuard package is still under active development. Follow the development progress on the developer's YouTube channel

Upgrading from versions older than pfSense 2.5.0

- The built-in `relayd` load balancer has been deprecated and removed as it does not compile or run on pfSense 2.5.0. A copy of the load balancer configuration will be left in `/conf/deprecated_load_balancer.xml` for reference when converting to an alternate solution, such as HAProxy (*HAProxy package*).
- PHP was migrated from PHP 7.2 to PHP 7.4. A number of PHP errors were fixed along the way, but certain combinations of configuration parameters may result in further errors. Note any problems on the Netgate Forum, and if possible, try to include relevant portions of `config.xml` with personal data removed.
- Due to the significant nature of the changes in this version of pfSense software, warnings and error messages, particularly from PHP and package updates, are likely to occur during the upgrade process. These errors are primarily seen on the console as the upgrade is applied, but may appear in a crash report once the upgrade completes. In nearly all cases these errors are a harmless side effect of the changes between FreeBSD 11.2 and 12.x and between PHP 7.2 and PHP 7.4.
- See the FreeBSD 12 Release Notes for information on deprecated hardware drivers that may impact firewalls upgrading to pfSense version 2.5.0. Some of these were renamed or folded into other drivers, others have been removed, and more are slated for removal in FreeBSD 13 in the future.
- OpenSSL was upgraded to 1.1.1a as a part of upgrading to FreeBSD 12.0, this will impact all packages which depend on OpenSSL, especially those not obtained from Netgate. Be aware that this will require obtaining new versions of such packages after the upgrade.

Upgrading from versions older than pfSense 2.4.5-p1

- Upgrading to pfSense software version 2.4.5-p1 requires `pfSense-upgrade` version 0.70 or later. Most installations will automatically pick up the new version and upgrade normally. If this does not happen automatically and the upgrade to version 2.4.5-p1 is not offered, use the following procedure:
  - Navigate to **System > Updates**
  - Set **Branch** to *Previous stable version*
  - Wait a few moments for the upgrade check to complete
  - Optional: Confirm that the latest version of *pfSense-upgrade* is present (version >= 0.70) using `pkg-static info -x pfSense-upgrade`.

    If the correct version is not present, wait a bit longer and check again as that package may be updating in the background.
  - Set **Branch** to *Latest stable version*
  - Wait a few moments for the upgrade check to complete

  At this point, the upgrade check should see 2.4.5-p1 and the upgrade can proceed.
- pfSense software version 2.4.5-p1 includes `pkg` version 1.13.x which introduces a new metadata version. Most installations will automatically pick up the new version and upgrade normally. In certain cases, especially coming from much older versions, the `pkg` utility may require a manual update before it can correctly process the new metadata.

  The `pkg` utility can be upgraded manually with the following command run from an ssh or console shell:

```
# pkg-static bootstrap -f
```

  See *Repository Metadata Version Errors* for more details.

Upgrading from versions older than pfSense 2.4.4

- Third party packages from **alternate repositories** are causing problems for users with the upgrade process and also with post-upgrade behavior. These packages have never been supported, and had to be manually added by users outside the GUI.

  Due to the major changes required for FreeBSD 11.2 and PHP 7.2, third party packages from alternate repositories cannot be present during the upgrade. There is no way to predict if a third party package is compatible with the new version or will cause the upgrade itself to fail.

  The upgrade process will automatically remove `pfSense-pkg-*` packages installed from alternate repositories. After the upgrade completes, the user can reinstall these packages. Packages from **alternate repositories** will not appear in the **Installed Packages** list in the GUI, and must be entirely managed in the command line.

  This change does not affect packages installed from the official pfSense package repository.

- *Automatic Configuration Backup Service* is integrated into pfSense version 2.4.4 and free for all to use. It is no longer an add-on package. It is now located under **Services > Auto Config Backup**.
- PHP was migrated from PHP 5.6 to PHP 7.2. A number of PHP errors were fixed along the way, but certain combinations of configuration parameters may result in further errors. Note any problems on the Netgate Forum or the pfSense subreddit, and if possible, try to include relevant portions of `config.xml` with personal data removed.
- Due to the significant nature of the changes in this version of pfSense software, warnings and error messages, particularly from PHP and package updates, are likely to occur during the upgrade process. These errors are primarily seen on the console as the upgrade is applied, but may appear in a crash report once the upgrade completes. In nearly all cases these errors are a harmless side effect of the changes between FreeBSD 11.1 and 11.2 and between PHP 5.6 and PHP 7.2.
- Gateway handling changes in 2.4.4 may result in different default gateway behavior than previous releases. Nearly all cases should behave properly, but be aware that it may be necessary to re-select the default gateway after upgrade.
- The FEC LAGG Protocol is deprecated, and its options have been removed #8734
- The login protection daemon was changed from `sshlockout_pf` to `sshguard` and the behavior may be more sensitive in some cases to SSH and GUI login failures. For example, be aware of possible issues where probes from monitoring systems may end up triggering a block.
- Major changes to RADIUS for the base system and specifically Captive Portal could lead to behavior changes in certain cases. Read the release notes and associated bug reports for more details. Note any problems on the Netgate Forum or the pfSense subreddit.
- A crash report containing no data (empty) may appear after the upgrade completes. See #8915
- Intel Atom systems containing HD Graphics chipsets may experience console problems after the update. Affected systems will boot successfully, but fail to display console output after the boot menu. To fix the problem, add the following line as a *Loader Tunable* to use the `syscons` console type:

```
kern.vty=sc
```

  - Alternately, try using i915 driver with the standard VT console using these lines as *Loader Tunables*:

```
i915kms_load="YES"
drm.i915.enable_unsupported=1
```

> **Warning:** This driver will consume extra bus resources and may cause resource hungry add-on hardware to fail, such as multi-port network adapters.

  - Systems with similar console problems not containing a graphics chip compatible with the i915 driver may need to reinstall 2.4.4 to use a UEFI console.

- An ISP that supplies a bogus interface MTU via DHCP may cause interface problems with certain network interface types when **Advanced Configuration** options are present on DHCP interfaces, such as a DHCP WAN. The typical default case is handled automatically, but advanced options override the corrected default behavior. To fix the problem, apply the patch from #8507 or add `supersede interface-mtu 0` to the **Option modifiers** box in the WAN interface advanced DHCP options. If a custom `dhclient.conf` is in use, add `supersede interface-mtu 0` on a line inside the `interface` block. See #8507. The **Advanced Configuration** case has been corrected for the next release.

Upgrading from versions older than pfSense 2.4.0

- To use ZFS, a reinstallation of the operating system is required. It is not possible to upgrade in-place from UFS to ZFS.
- **Wireless** interfaces **must** be created on the **Wireless** tab under **Interfaces > Assignments** before they are available for assignment
- Some hardware devices may not boot 2.4.0 installation images, for example, due to UEFI compatibility changes. These are primarily BIOS issues and not issues with the installer images. Upgrading in place from 2.3.x typically allows affected hardware to run version 2.4.
- To upgrade Firewalls in place which are running pfSense software version 2.2.x or earlier, first upgrade the firewall to pfSense 2.3.4 and then perform an update to pfSense 2.4.x afterward. Alternately, reinstall 2.4.x directly and restore the configuration.

> **Warning:** When upgrading to 2.4.x from 2.2.x or earlier, remove all packages before attempting the update. Even when upgrading from 2.3.x this is the best practice to ensure a smooth upgrade process. Package settings are retained.

Older Version Upgrade Notes

Versions of pfSense software prior to 2.3 used a different upgrade method. For "full" installations, a `tgz` file was used by the firewall to copy in the new files. This method was problematic and is no longer used.

The best practice in these cases is to take a backup and reinstall with a current, supported version of pfSense software.

The following information is for upgrading from outdated and unsupported versions of pfSense software. They may still be of use to users attempting to upgrade from an older release to a current, supported, release.

When upgrading from a very old release, read every document below that covers versions between the older one being upgraded and the new version.

Upgrading from versions older than pfSense 2.3

See also:

For information about upgrading to current versions, see *Upgrade Guide*.

Warning: Uninstalling *all* packages is **required** when upgrading from old releases. Packages **must be removed** before the upgrade is performed. After the upgrade is complete, packages can be reinstalled. Package configuration is automatically retained.

See *2.3 New Features and Changes* for a larger list of changes.

- Due to the GUI overhaul, older themes have been removed. All previously chosen themes are reset on upgrade to the default "pfSense" 2.3 theme.
- **Status > RRD Graphs** moved to **Status > Monitoring** and has been revamped. The same data, and more, is still accessible but with a modern interface.
- **System > Firmware** is now **System > Update**
- **System > Packages** is now **System > Package Manager**

Limiters

- On pfSense® software versions 2.2 and 2.3, limiters cannot be used on firewall rules residing on interfaces where NAT applies. This limits their use to LAN-type interfaces only, and not WANs, in most circumstances. This has been fixed on pfSense 2.4. Bug #4326
- On pfSense software versions 2.2 and 2.3, limiters cannot be used where pfsync is enabled. This has been fixed on pfSense 2.4.3. Bug #4310

NanoBSD

Warning: NanoBSD has been deprecated as of pfSense 2.4.0-RELEASE. This section remains only for users on i386 hardware with NanoBSD who must upgrade to pfSense 2.3.5-p2.

In most cases, a normal installation may be used in place of NanoBSD. Activating the option to keep `/var` and `/tmp` in RAM can typically yield the same net benefits for older/slower CF and SD media. Firewalls with modern SSDs should have no concerns with writes.

1GB NanoBSD images have been removed as they were too small to properly function and upgrade. If a 1GB NanoBSD image is in use, it cannot be upgraded. It must be re-imaged on a larger card using the 4GB or 2GB image or converted to a full installation.

Package System

- Due to the package system overhaul, any custom package repository settings are removed. The firewall will pull package information directly from Netgate servers.
- We highly recommend uninstalling all packages before upgrading.

Removed features that are disabled on upgrade

- Groups with spaces are no longer permitted. They are not allowed at the OS level and were not functioning properly. On upgrade, such groups are renamed with an underscore ('_') in place of a space.
- The "Enable" checkbox for IPsec has been removed. If IPsec was disabled, all Phase 1 entries are disabled automatically on upgrade.
- The **Unity** plugin for IPsec has been disabled by default, where it was previously enabled by default. This is preferable for the vast majority of users, however those using mobile IPsec with IKEv1 may need to enable it under **VPN > IPsec**, **Advanced** tab.
- The `apinger` daemon for gateway monitoring has been replaced by `dpinger`. Due to the differences in settings between the two, many advanced gateway parameters are reset on upgrade.
- The PPTP *Server* has been removed, if the PPTP server was in use, seek alternate solutions such as IPsec or OpenVPN. **Do not continue to use PPTP**.
  - The PPTP server settings, firewall rules, and so on have all been removed
  - If the "Redirect" PPTP server type was in use, add manual NAT rules for TCP/1723 and GRE to point to the actual server.
- Layer 7 classification has been removed and any configuration using L7 is automatically removed on upgrade.
- WEP has been removed from Wireless interfaces, and if a wireless interface was using WEP, the interface is deactivated on upgrade.
- Single DES has been removed from IPsec, if a Phase 1 or Phase 2 entry was using DES, it is deactivated on upgrade.
  - Note: 3DES is still present. Only the older and insecure, **single** DES option was removed.
- The Live CD platform has been removed. The ISO is a bootable installer, as always, but it cannot run a live system.
  - For the very few people who were still using Live CD: If the hardware can boot from USB, install to a USB thumb drive and run from it instead. Use the options to keep `/var` and `/tmp` in RAM, and do not install packages, then net result should be similar but ultimately more functional.
- Some obsolete password hashes, such as `nt-hash`, are removed from users on upgrade. There was no remaining code on pfSense that utilized these hashes, so there should be no loss of functionality.
- `fifolog` was removed, and will revert to clog format on upgrade.
- The `net.inet.ip.fastforwarding` tunable is no longer present, and is unset on upgrade.
- Some PHP modules, such as MySQL, were included by default on previous versions but are no longer a part of the base system on 2.3. They are available as packages that may be installed manually from the shell (e.g. `pkg install php56-mysql`)

New features that may require action

- The default system password hash has been changed to `bcrypt`. Current passwords will continue to work. Existing users need to reset their password to convert to the new, more secure, hash. #4120
- A new option was added to Captive Portal for stop/start RADIUS accounting updates compatible with FreeRADIUS. This option solves problems with user session time limits. If a Captive Portal was using stop/start RADIUS accounting with FreeRADIUS, change it to use the new option manually.

Upgrading from a 2.3 Snapshot

- If a firewall was upgraded to 2.3 before Jan 21, 2016, some files from 2.2.x or earlier packages may still be left behind that can prevent new packages from installing properly. Run the following command to clean outdated symlinks that are not relevant for 2.3:

```
find / -type l -lname '/usr/pbi/*' -delete
```

Multi-WAN Weighted Load Balancing

There is a quirk in PF handling of weighted load balancing where Load balancing fails when one gateway has a weight of 1 and another gateway has a weight >1. Coming from 2.2.x, if this scenario applies, simply double the assigned weights. For example: WAN1 = `1`, WAN2 = `5` on 2.2.x should be WAN1 = `2`, WAN2 = `10` on 2.3.

Captive Portal

Due to the change in the web server from `lighttpd` to `nginx`, in some cases the portal HTML must be updated to include the zone parameter. On 2.3.1 and later the web server process attempts to handle this automatically, but it is best to include the HTML in the portal page directly, inside the form tag:

```
<input name="zone" type="hidden" value="$PORTAL_ZONE$">
```

Upgrading from versions older than pfSense 2.2

See also:

For information about upgrading to current versions, see *Upgrade Guide*.

Warning: Uninstalling **all** packages is **required** when upgrading from old releases. Packages **must be removed** before the upgrade is performed. After the upgrade is complete, packages can be reinstalled. Package configuration is automatically retained.

Limiters

- On pfSense® software versions 2.2 and 2.3, limiters cannot be used on firewall rules residing on interfaces where NAT applies. This limits their use to LAN-type interfaces only, and not WANs, in most circumstances. This has been fixed on pfSense 2.4. Bug #4326
- On pfSense software versions 2.2 and 2.3, limiters cannot be used where pfsync is enabled. This has been fixed on pfSense 2.4.3. Bug #4310

IPsec Changes

The IPsec daemon was changed from **racoon** to **strongSwan**. Existing configurations work the same as always, but if any unusual configurations are present, take care in testing after the upgrade. Changes in behavior because of this change may trigger bugs in remote endpoints that weren't previously an issue. Configurations that were always technically incorrect may exhibit problems now where they didn't previously. We have listed the circumstances we are aware of here, and will expand upon this list if anything new is found.

Problem in racoon with aggressive mode and NAT-D

Those using racoon (pfSense 2.1.x and earlier, among a variety of other similar products) on remote endpoints with aggressive mode may encounter a bug in racoon related to NAT-D and aggressive mode. Any site to site IPsec VPNs using aggressive mode with racoon as a remote endpoint should change to main mode to prevent this from being an issue. Main mode is always preferable for its stronger security.

`glxsb` Crypto Accelerator Warning

For those using the `glxsb` crypto accelerator in the ALIX and other devices with Geode CPUs, they are only compatible with 128-bit AES. Any key length greater than 128 bits has never worked, and must not be configured. There appear to be circumstances where AES on "auto" with racoon preferred 128-bit AES where strongSwan prefers the strongest-available and is choosing 256-bit AES, which `glxsb` breaks. Input validation in 2.2.1 prevents such invalid configurations when adding configurations or making changes, however existing configurations are not changed. If using `glxsb` and AES, ensure both phase 1 and phase 2 configurations all use AES 128 only and never auto.

Mobile client users, verify Local Network

For mobile IPsec clients, clients could pass traffic in certain circumstances without having specified the necessary matching local network in the mobile phase 2 configuration. The "Local Network" specified in mobile IPsec phase 2 must include all networks mobile clients need to reach. If mobile IPsec clients need to access the Internet via IPsec, the mobile phase 2 must specify `0.0.0.0/0` as the local network.

Stricter Phase 1 Identifier Validation

In 2.1.x and earlier versions, racoon could accept mismatched phase 1 identifiers where using *IP Address* as the identifier. This is most commonly a problem where one of the endpoints is behind NAT and phase 1 is using *My IP Address* and *Peer IP Address* for identifiers. On the side with the private IP WAN, *My IP Address* will be its private WAN IP address. On the opposite end, *Peer IP Address* will be the public IP address of the opposite side. Hence, these two values do not match, and should have resulted in a connection failure. racoon would fall back to checking the source IP address of the initiating host as an identifier, where it found the match. To resolve this issue, change the phase 1 identifiers so they actually match.

Phase 2 behavior change with incorrect network addresses

In 2.1.x and earlier versions a phase 2 configuration with an incorrect network address would still be presented by racoon with the corrected network address. e.g. if `192.168.1.1/24` is set in a phase 2, which should be `192.168.1.0/24`, racoon used it as `192.168.1.0/24`. In 2.2.x and newer versions, strongSwan sends it exactly as configured. This may result in a phase 2 mismatch where configured with an incorrect network address.

Disk Driver Changes

The disk drivers in FreeBSD changed between the underlying OS versions and now the CAM-based ATA drivers and AHCI are used by default. As such, ATA disks are labeled as `/dev/adaX` rather than `/dev/adX`. The `ada` driver for ATA disks and GEOM keeps legacy aliases in place so that old disk references will still work post-upgrade. This does not always extend to virtualized disk drivers, however (see the Xen note below.). The upgrade process on pfSense 2.3 and 2.4 also attempts to automatically correct for this change.

A manual workaround is also possible. Running `/usr/local/sbin/ufslabels.sh` **before** the upgrade will convert `/etc/fstab` to UFS labels rather than disk device names bypassing any device name issues that could arise due to the switch.

There is a chance that the new driver stack will have issues with certain controller/disk combinations that were not present in prior releases. There may be BIOS changes or other workarounds to help. See *Boot Troubleshooting*.

The methods used to disable DMA and write caching have both changed on FreeBSD 10.x. For most, disabling these manually is no longer necessary.

If disabling DMA is necessary, the following may be used in a *Loader Tunable*:

```
hint.ata.X.mode=PIO4
```

Change `X` to be the ATA controller ID, typically `0` or `1`.

If write caching must be disabled, the following may be used as a *Loader Tunable*:

```
kern.cam.ada.write_cache=0
```

Xen Users

The FreeBSD base used by pfSense 2.2 and later includes PVHVM drivers for Xen in the kernel. This can cause Xen to automatically change the disk and network device names during an upgrade to pfSense 2.2 or later, which the Hypervisor should not do but does anyway.

The disk change can be worked around by running `/usr/local/sbin/ufslabels.sh` **before** the upgrade to convert `/etc/fstab` to UFS labels rather than disk device names.

The NIC device change issue has no workaround. Manual reassignment is required.

vmxnet3 (VMware/ESX) users

Users who manually installed *VMware Tools* to use `vmxnet3` network adapters may encounter an issue with interface name changes when upgrading to pfSense 2.2 or later, similar to those with Xen mentioned above. In pfSense 2.1.x the `vmxnet3` interfaces were named starting with `vmx3f` and on pfSense 2.2.x they are `vmx` using the built-in drivers. Manually reassigning the interfaces or correcting them in `config.xml` followed by a restore is required.

Old/Broken GEOM Mirrors

If a manual `gmirror` configuration was performed post-install and not using the pfSense installer `gmirror` option before install, there is a chance that the mirror will not function on pfSense 2.2 or later because the manual post-install method did not create a proper mirror setup. If an upgraded mirror does not boot or function on pfSense 2.2 or later, use the following entry to work around the integrity check that would otherwise fail.

Add the following line as a *Loader Tunable*:

```
kern.geom.part.check_integrity=0
```

If the disks are configured in this way, we **strongly** recommend backing up the configuration and reinstalling, using one of the mirrored disk options in the pfSense installer.

CARP Changes

Due to the new CARP subsystem, the old method of having a virtual interface for CARP VIPs is no longer available. CARP VIPs work more like IP Alias style VIPs, existing directly on the main interface. For most, the changes made to accommodate this new system will be transparent, but there are some potential issues, such as:

- With no separate interface available, monitoring a CARP VIP status via SNMP is no longer possible.

FTP Proxy

The FTP proxy is not included in pfSense 2.2-RELEASE or later, due to changes in the kernel and state table handling that made it more difficult to implement. Use of FTP is strongly discouraged as credentials are transmitted insecurely in plain text. #4210

See *FTP without a Proxy* for additional information and workarounds.

Another option is the recently added FTP Client Proxy package which leverages in FreeBSD to allow clients on local interfaces to reach remote FTP servers with active FTP.

LAGG LACP Behavior Change

LAGG using LACP in FreeBSD 10.0 and newer defaults to "strict mode", which means the lagg does not come up unless the attached switch is speaking LACP. This will cause a LAGG to not function after upgrade if the switch is not using active mode LACP.

To retain the lagg behavior in pfSense 2.1.5 and earlier versions, add a new system tunable under **System > Advanced**, **System Tunables** tab for the following:

```
net.link.lagg.0.lacp.lacp_strict_mode
```

With value set to `0`.

This can be added before upgrading to 2.2 to ensure the same behavior on first boot after the upgrade. It will result in a harmless cosmetic error in the logs on 2.1.5 since the value does not exist in that version.

If a firewall has more than one LAGG interface configured, enter a tunable for each instance since that is a per-interface option. For `lagg1`, add the following:

```
net.link.lagg.1.lacp.lacp_strict_mode
```

Also with the value set to `0`.

Intel 10Gbit/s ixgbe/ix users with Unsupported SFP modules

The `sysctl` OID to allow unsupported SFP modules changed in FreeBSD between the versions used for pfSense 2.1.x and 2.2.

The old *Loader Tunable* was:

```
hw.ixgbe.unsupported_sfp=1
```

This must be changed to:

```
hw.ix.unsupported_sfp=1
```

Edit the *Loader Tunable* before applying the update and the behavior will be retained.

Layer 7

Layer 7 is deprecated and has been removed. For layer 7 application identification and filtering we recommend using the *Snort IDS/IPS* package with OpenAppID detectors and rules.

Microsoft Load Balancing / Open Mesh Traffic

Windows Network Load Balancing and Open Mesh access points can use multicast MAC address destinations which rely on broken behavior that was incorrectly allowed by default in earlier versions of FreeBSD and pfSense. The fact it worked before was technically a bug, acting in violation of RFC 1812.

> A router MUST not believe any ARP reply that claims that the Link Layer address of another host or router is a broadcast or multicast address.

The default behavior on pfSense 2.2 is correct, but it may be changed.

If this behavior be required, manually add a tunable as follows:

- Navigate to **System > Advanced**, **System Tunables** tab
- Click 
- Enter the following values:
  - **Tunable**: `net.link.ether.inet.allow_multicast`
  - **Description**: Optional. It would be wise to enter the URL to this note or a similar note.
  - **Value**: 1
- Click **Save**

Upgrading from versions older than pfSense 2.1

See also:

For information about upgrading to current versions, see *Upgrade Guide*.

> **Warning:** Uninstalling **all** packages is **required** when upgrading from old releases. Packages **must be removed** before the upgrade is performed. After the upgrade is complete, packages can be reinstalled. Package configuration is automatically retained.

See the HA section at the end of this document for a High Availability-specific pfsync note about pfSense® software version 2.1 upgrades.

The **State Killing on Gateway Failure** feature (**System > Advanced**, **Miscellaneous tab**) now kills ALL states when a gateway has been detected as down, not only states on the failing WAN. This is done because otherwise the LAN-side states were not killed appropriately, and thus some connections would be in limbo, especially SIP. Due to the change in its behavior, **State Killing on Gateway Failure** is disabled by default in new configurations and is disabled during upgrade to pfSense 2.1.x from 2.0.x or before regardless of the user's previously chosen setting. If the feature is desired even with its new behavior, it must be manually re-enabled post-upgrade.

The **Allow IPv6** checkbox is **NOT** changed on upgrade unlike it was in early pfSense 2.1 BETA snapshots. This was changed so that the user's chosen existing behavior is preserved. To allow IPv6 traffic after an upgrade, the setting must be changed manually. This setting is located on **System > Advanced** on the **Networking** tab. It defaults to allowed for new configurations.

Changes to policy route negation between pfSense 2.0.x and 2.1 may result in local/private traffic hitting policy routing rules that would not have happened on pfSense 2.0.x. This most commonly presents as an inability to reach local networks after upgrading. The automatic policy route negation rules on pfSense 2.0.x were too lenient, and that behavior was corrected. To ensure proper routing to other local interfaces, VPNs, or static route networks rules must be added to the local interfaces to pass traffic to these destinations without a gateway set. And that rule must be above any others that would match and have a gateway set.

Upgrading High Availability Deployments

If upgrading from any previous version of pfSense (1.2.x, 2.0.x, etc.) to pfSense 2.1 or later in an HA environment, ensure that the pfsync interface has a rule to pass the correct traffic for state synchronization to work properly. pfSense 2.1 removed the automatic pfsync rule, since the documentation always recommended adding it manually and to add it behind the scenes with no way to block it could be counter-productive and potentially insecure. If the documentation was not followed, and a pfsync or allow all rule was not created on the sync interface, state synchronization may break after this upgrade. Add an appropriate rule to the sync interface and it will work again.

At a minimum, pass traffic of the *pfsync* protocol from a source of the synchronization subnet to all cluster nodes.

Upgrading from versions older than pfSense® 2.0

See also:

For information about upgrading to current versions, see *Upgrade Guide*.

Warning: Uninstalling **all** packages is **required** when upgrading from old releases. Packages **must be removed** before the upgrade is performed. After the upgrade is complete, packages can be reinstalled. Package configuration is automatically retained.

Note for users of the OpenVPN Status Package

If a manual `management` directive was entered into the **Advanced Configuration** of an OpenVPN client or server, it must be removed. The OpenVPN status code is built into pfSense® software version 2.x and later, and it is handled internally. The management directive must be removed or the status of the VPN instance will not be properly reported.

Note for Captive Portal RADIUS WISPr Bandwidth Users

The WISPr RADIUS attributes were incorrectly applied to all versions prior to pfSense 2.0-RELEASE. They were applied as *Kbps* where WISPr is supposed to be *bps*, hence those using WISPr attributes will have one one-thousandth of the previous bandwidth unless the RADIUS server is corrected. The RADIUS server will need to have these values updated to bps for proper functionality once the firewall has been upgraded to pfSense 2.0-RELEASE or later.

International/Special Characters in 1.2.x Configurations

International characters, such as `åäö` and so on, were not compatible with pfSense software version 1.2.x, but were allowed in some places due to overly lenient input validation and less strict XML parsing. These characters cause invalid XML when they are stored directly, and as such if pfSense 1.2.x did not crash and toss out the configuration with such characters, it will not upgrade to any current version of pfSense software.

pfSense software version 2.0 and later will reset and toss out the config.xml on every reboot if it contains these characters bare, leaving the firewall at an "assign interfaces" prompt since it does not have a valid configuration.

The `config.xml` file can be run through an XML parser such as `xmllint` and the parser will show where problems exist, if any. Fix the errors, and then the corrected configuration can be used for an upgrade. The good news is that these characters are handled properly in most areas of the current pfSense software GUI, and they are `CDATA` escaped so they are safe from such errors.

Upgrading High Availability Deployments

When upgrading from pfSense 1.2.3 to 2.0 or later, Check the CARP VIPs to make sure they are actually on the proper interface. That is, that the interface chosen for the VIP properly matches the subnet in which the CARP VIP resides, and that the subnet mask is proper. pfSense 2.0 validates this more strictly than previous releases, and as a consequence if a CARP VIP was misconfigured on pfSense 1.2.3, it may not upgrade cleanly.

7.6.4 Perform the Upgrade

There are several methods available to update an installation of pfSense® software. Either the GUI or the console can be used.

Note: Before performing an upgrade, read through the entire *Upgrade Guide*.

If problems occur during the upgrade process, consult *Troubleshooting Upgrades* for assistance.

Upgrading using the GUI

The **Automatic Update** check feature contacts a Netgate server and determines if there is a release version newer than the version on the firewall. This check is performed when an administrator visits the dashboard or **System > Update**.

To perform the upgrade in the GUI:

- Navigate to **System > Update** or click  in the System Information dashboard widget next to the new version notification.
- Click  **Confirm** to start the update
- Wait for the upgrade to complete

The update takes a few minutes to download and apply, depending on the speed of the Internet connection being used and the speed of the firewall hardware. The firewall will reboot automatically when finished.

Tip: Monitor the firewall console during the upgrade if possible to watch for potential problems.

Upgrading using the Console

An update may also be run from the console. The console option is available from any means available for console access: Video/Keyboard, Serial Console, or SSH.

- Connect to the firewall console or login via ssh
- Enter menu option `13`
- Wait for the upgrade to complete

Alternately, from a shell prompt running as root, manually execute the following command:

```
# pfSense-upgrade
```

Tip: When upgrading from SSH, the GNU screen utility can be a useful tool to monitor the upgrade process in environments where the connection to the firewall is unstable:

```
# pkg install screen
# rehash
# screen pfSense-upgrade
```

Reinstalling / Upgrading Configuration

If an upgrade will not function properly on an existing installation, the configuration file can be restored to a freshly installed copy of pfSense software. An older configuration can always be imported into a new version. The upgrade code will make necessary changes to the configuration so it will work with the current version of the software. See *Backup and Recovery* for details.

7.6.5 Upgrading High Availability Clusters

This document provides guidance on upgrading redundant firewalls in a high availability configuration across major versions of pfSense® software.

Upgrading from one version to another generally follows this procedure, exceptions are noted later in the page:

- Review release notes, blog, and upgrade guide.
- Take a backup from both nodes.

  Danger: Do not skip this step, backups are critical!

- Upgrade the secondary node as described in the *Upgrade Guide*.
- Test the secondary node to be sure it is operating as expected – correct packages present, services running, no obvious errors in logs, etc.
- Enter Persistent CARP maintenance mode on the primary node from **Status > CARP**.
- Ensure traffic is still flowing properly and that the network is functional.

  If it is not, then exit maintenance mode on the primary node, fix the secondary node, then try again.
- Upgrade the primary node as described in the *Upgrade Guide*.
- Check the primary node to ensure it upgraded and is operating as expected – correct packages present, services running, no obvious errors in logs, etc.
- Leave Persistent CARP maintenance mode on the primary node.
- Test everything again.

Tip: If this cluster is using the ISC DHCP backend, consider converting to the Kea DHCP backend as the ISC DHCP backend is deprecated and will be removed in the near future. See *Converting High Availability DHCP from ISC to Kea* for details.

XMLRPC Config Sync Considerations

Upgrade either the primary or the secondary first, leaving the other on the older version until testing is complete.

Current versions of pfSense software properly check for and prevent unintentionally synchronizing data between incompatible versions.

pfsync considerations

The underlying pfsync protocol often changes between FreeBSD versions. Versions of pfSense software with a different base OS version of FreeBSD may not be capable of synchronizing their states between each other. Failover will still function, but not stateful failover so all existing connections will be dropped.

pfsync and State Policy

The State Policy (*Firewall State Policy*) of the firewall can introduce a conflict in state synchronization if nodes using pfsync do not have identical hardware.

See *pfsync and Physical Interfaces* for details and a potential workaround.

CARP considerations

CARP is generally the same between versions and will fail over and back as expected.

See also:

- *Troubleshooting Upgrades*

7.6.6 Update Settings

Branch / Tracking Snapshots

By default, the update check looks for officially released versions of pfSense software, but this method can also be used to track development snapshots.

To change the branch used for updates:

- Navigate to **System > Update**
- Set the **Branch** to the desired type of updates
- Wait for the page to refresh and perform a new update check

The branch list will vary depending on the current development cycle. Examples of options that may be found in the list include:

Latest Stable Version
: Stable versions are the best option, as they see the most testing and are reasonably safe and trouble-free. However, as with any upgrade, read the *changelog* and *update notes* for that release.

pfSense Plus Upgrade
: Upgrade a system from pfSense CE software to pfSense Plus software. Present on registered systems with access to pfSense Plus software repositories.

See also:

See *Migrate from pfSense® CE software to Netgate pfSense Plus software* for details on migrating to pfSense Plus software.

Previous Stable Version (Deprecated)
A pointer to the previous release so that firewalls may pull packages and update files from the previous release while waiting for a maintenance window or similar upgrade opportunity. May also be labeled "Legacy".

Latest Development Snapshots
Tracks development snapshot builds. These may either be snapshots for the next minor or major version depending on the status of the development cycle.

Next Major Version
Tracks snapshots for the next major update version. This is riskier, but in some cases may be required for newer hardware or new features that are not yet released. Consult the forum and test in a lab to see if these snapshots are stable in a particular environment.

> **Warning:** Do not run development versions of pfSense software in production environments.

Boot Environments

There are a handful of options related to *ZFS Boot Environments* which only appear on systems running pfSense Plus software installed using ZFS.

See also:

- *How Boot Environments Work*
- *Upgrade Process Overview*

The available options are:

Defer Automatic Reboot
When checked, the firewall will not automatically reboot after finishing the upgrade. Instead, it will wait for an administrator to reboot it manually.

Boot Verification
These options control the boot time verification for a Boot Environment. If verification fails, the firewall will automatically roll back to a previous known-good Boot Environment and reboot.

The firewall displays a prompt on the Dashboard to verify or fail the Boot Environment.

Manual Boot Verification
When set, the Boot Environment is not automatically verified and must be verified manually before the verification interval expires.

Boot Verification Interval
This option defines the amount of time the administrator has to verify that the Boot Environment is in working order. If the timer expires, the firewall will automatically roll back to a previous known-good Boot Environment and reboot.

The default value is `300` seconds (5 minutes).

See also:

See *ZFS Boot Environments (Plus Only)* for more information.

Dashboard Check

The **Dashboard Check** checkbox on **System > Update**, **Update Settings** tab controls whether the **System Information** widget on the dashboard performs an update check. On firewalls with low resources or slow disks, disabling this check will reduce the load caused by running the check each time an administrator views the dashboard.

GitSync

This section is for developers and should not be used by end users. Leave settings in this area empty or disabled.

7.7 Migrate from pfSense® CE software to Netgate pfSense Plus software

Netgate now offers the ability to migrate from the Community Edition (CE) of pfSense® software to pfSense Plus software.

This enables users with virtual machines or hardware not sold by Netgate to utilize the advantages of pfSense Plus software.

pfSense Plus Software Migration Procedure

- *Requirements*
- *Obtain an Activation Token*
- *Register and Migrate*

7.7.1 Requirements

To perform this migration:

- The firewall must be running pfSense CE software version 2.6.0 or later.

  Before starting, take one of the following steps:

  - Perform a fresh installation of at least pfSense CE software version 2.6.0 by following the *installation guide*.
  - Upgrade an existing installation of pfSense CE software to version 2.6.0 or later by following the *upgrade guide*.
- The firewall must be connected to the Internet to perform the migration.

Warning: The migration process preserves the existing filesystem type, so ensure that a firewall is in the intended state before upgrading. For example, install pfSense CE software using ZFS so that it can use pfSense Plus software with ZFS.

7.7.2 Obtain an Activation Token

Activation tokens are generated by the Netgate Store. To obtain a token, follow these steps:

- Visit the Netgate Store
- Create a new account or log into an existing account
- Visit the pfSense Plus Software Subscription product page
- Select the desired **Software Type**
- Add the product to the cart
- Complete the checkout process

After completing the checkout process the store will send an activation token by e-mail to the address on the Netgate Store account.

Tip: If the activation e-mail does not arrive in a timely manner, check spam or junk mail folders in the e-mail client.

Warning: Activation tokens are single use. Ensure the pfSense CE software installation is functional and is in the intended configuration before performing the migration.

7.7.3 Register and Migrate

- Navigate to **System > Register** in the pfSense CE software GUI
- Paste the **Activation Token** into the text area on the page
- Click **Register**

  The page will display a message indicating the registration results. If the registration was successful, continue. If registration failed, contact Netgate TAC.
- Navigate to **System > Update**

  The page will contain a message announcing the pfSense Plus software migration branch.
- Set **Branch** to *pfSense Plus Upgrade* as seen in figure *pfSense Plus Software Branch Selection*.
- Wait for the firewall to complete the update check
- Click **Confirm** to confirm and start the migration process

The migration process will proceed from there and reboot when it is complete. This may take several minutes to complete, especially in locations with slow download speeds. Monitor the console for progress.

Warning: Do not manually reboot or remove power from the device until the migration completes as this may interrupt the process and cause it to fail.

Congratulations, the firewall is now running pfSense Plus Software!

See also:

- *Virtualization*
- *Connect to the Console*

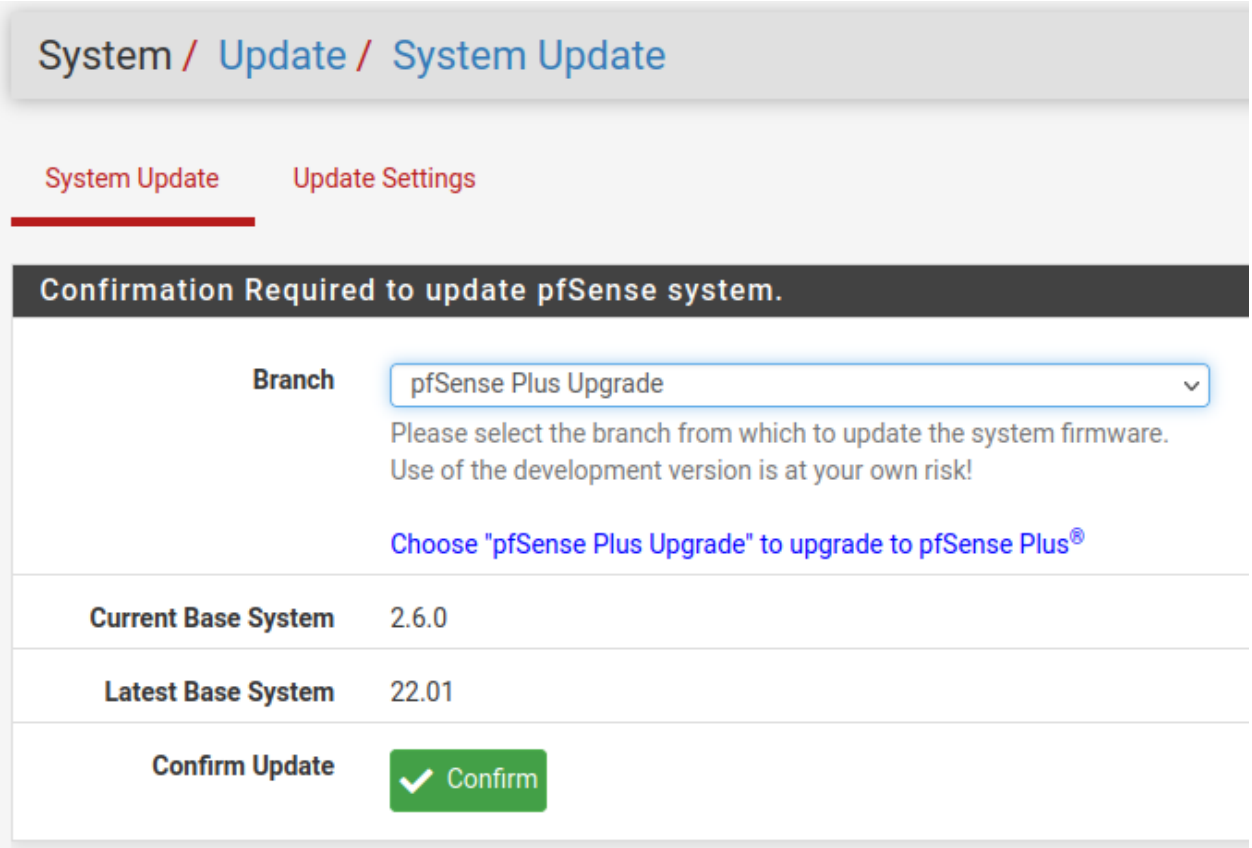


Fig. 28: pfSense Plus Software Branch Selection

- *Troubleshooting Installation Issues*
- *Troubleshooting Upgrades*
- *Troubleshooting Disk and File System Issues*

CHAPTER

EIGHT

CONFIGURATION

8.1 Setup Wizard

The first time a user logs into the pfSense® software GUI, the firewall presents the Setup Wizard automatically. The first page of the wizard is shown in Figure *Setup Wizard Starting Screen*.

Click » **Next** to proceed.

Tip: Using the setup wizard is optional. Click the logo at the top left of the page to exit the wizard at any time.

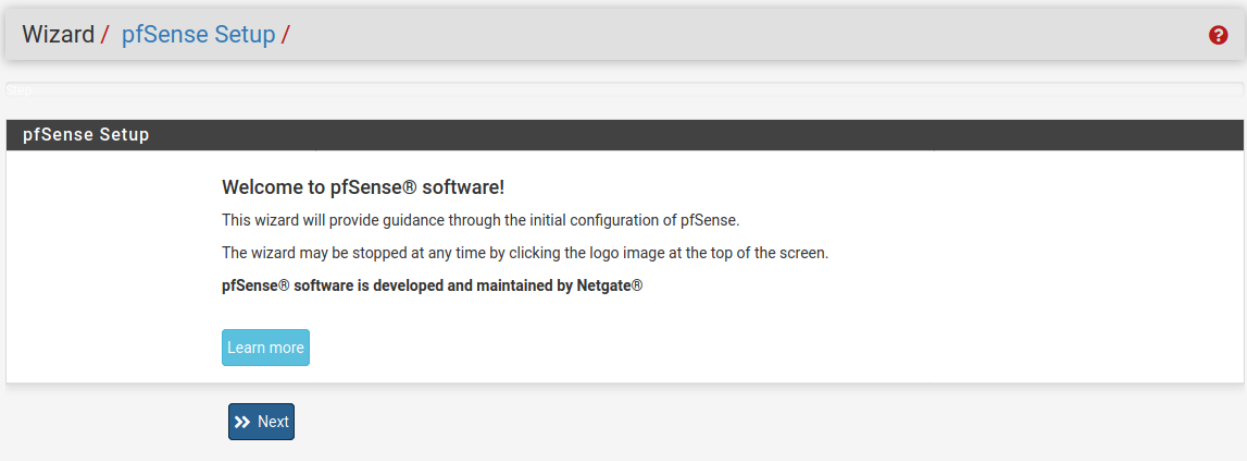


Fig. 1: Setup Wizard Starting Screen

The next screen of the wizard explains the availability of assistance from Netgate TAC. Click » **Next** again to start the configuration process using the wizard.

8.1.1 General Information Screen

The next screen (Figure *General Information Screen*) configures the name of this firewall, the domain in which it resides, and the DNS servers for the firewall.

Hostname
The **Hostname** is a name that should uniquely identify this firewall. It can be nearly anything, but must start with a letter, and it may contain only letters, numbers, or a hyphen.

Domain
Enter a Domain, e.g. `example.com`. If this network does not have a domain, use `<something>.home.arpa`, where `<something>` is another identifier: a company name, last name, nickname, etc. For example, `company.home.arpa` The hostname and domain name are combined to make up the fully qualified domain name of this firewall.

Primary/Secondary DNS Server
The IP address of the Primary DNS Server and Secondary DNS Server, if known.

These DNS servers may be left blank if the DNS Resolver will remain active using its default settings. The default configuration has the DNS Resolver active in resolver mode (not forwarding mode), when set this way, the DNS Resolver does not need forwarding DNS servers as it will communicate directly with Root DNS servers and other authoritative DNS servers. To force the firewall to use these configured DNS servers, enable forwarding mode in the DNS Resolver or use the DNS Forwarder.

If this firewall has a dynamic WAN type such as DHCP, PPTP or PPPoE these may be automatically assigned by the ISP and can be left blank.

Note: The firewall can have more than two DNS servers, add more under **System > General Setup** after completing the wizard.

Override DNS
When checked, a dynamic WAN ISP can supply DNS servers which override those set manually. To force the use of only the DNS servers configured manually, uncheck this option.

See also:

For more information on configuring the DNS Resolver, see *DNS Resolver*

Click  **Next** to continue.

8.1.2 NTP and Time Zone Configuration

The next screen (Figure *NTP and Time Zone Setup Screen*) has time-related options.

Time server hostname
A Network Time Protocol (NTP) server hostname or IP address. Unless a specific NTP server is required, such as one on LAN, the best practice is to leave the **Time server hostname** at the default `2.pfsense.pool.ntp.org`. This value will pick a set of random servers from a pool of known-good NTP hosts.

To utilize multiple time server pools or individual servers, add them in the same box, separating each server by a space. For example, to use three NTP servers from the pool, enter: `0.pfsense.pool.ntp.org 1.pfsense.pool.ntp.org 2.pfsense.pool.ntp.org`

This numbering is specific to how `.pool.ntp.org` operates and ensures each address is drawn from a unique pool of NTP servers, so the same server does not get used twice.

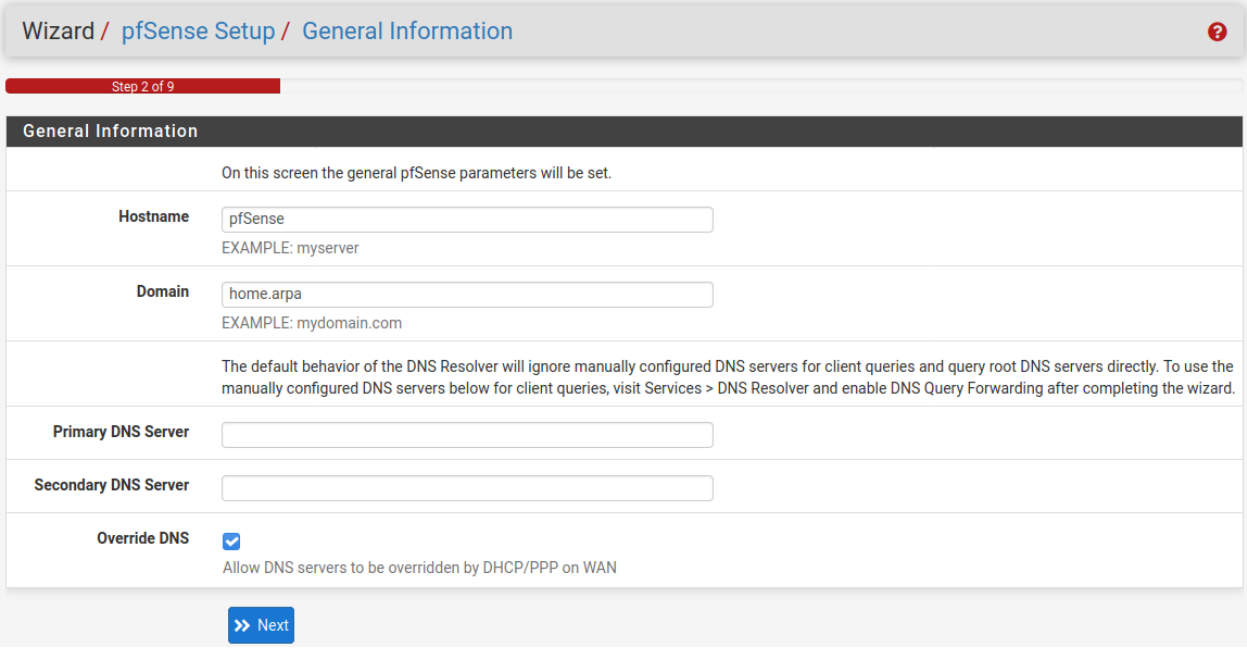


Fig. 2: General Information Screen

Timezone
: Choose a geographically named zone which best matches location of this firewall, or any other desired zone.

Click  **Next** to continue.

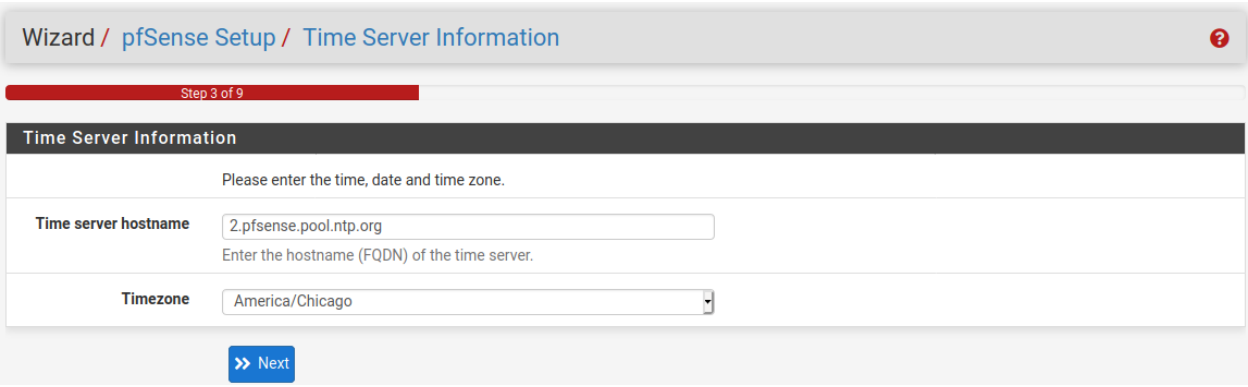


Fig. 3: NTP and Time Zone Setup Screen

8.1.3 WAN Configuration

The next page of the wizard configures the WAN interface of the firewall, which is the external network facing the ISP or upstream router. The wizard offers configuration choices for several common types of ISP connections.

WAN Type

The **Selected Type** (Figure *WAN Configuration*) must match the type of WAN required by the ISP, or whatever the previous firewall or router was configured to use. Possible choices are *Static*, *DHCP*, *PPPoE*, and *PPTP*. The default choice is *DHCP* due to the fact that it is the most common, and for the majority of cases this setting allows a firewall to "Just Work" without additional configuration. If the WAN type is not known, or specific settings for the WAN are not known, this information must be obtained from the ISP. If the required WAN type is not available in the wizard, or to read more information about the different WAN types, see *Interface Types and Configuration*.

Note: If the WAN interface is wireless, additional options will be presented by the wizard which are not covered during this walkthrough of the standard Setup Wizard. Refer to *Wireless*, which has a section on Wireless WAN for additional information. If any of the options are unclear, skip the WAN setup for now, and then perform the wireless configuration afterward.

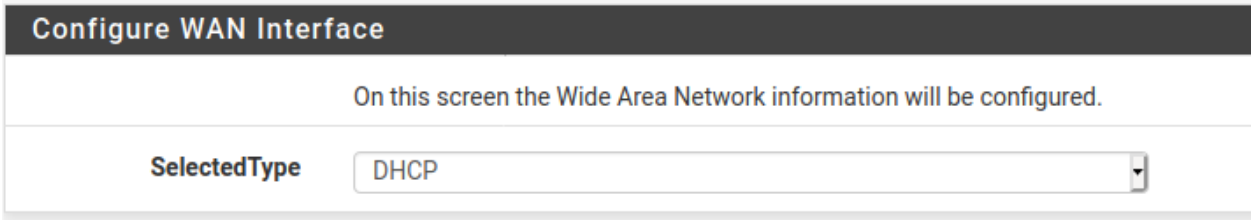


Fig. 4: WAN Configuration

MAC Address

This field, shown in Figure *General WAN Configuration*, changes the MAC address used on the WAN network interface. This is also known as "spoofing" the MAC address.

Note: The problems alleviated by spoofing a MAC address are typically temporary and easily worked around. The best course of action is to maintain the original hardware MAC address, resorting to spoofing only when absolutely necessary.

Changing the MAC address can be useful when replacing an existing piece of network equipment. Certain ISPs, primarily Cable providers, will not work properly if a new MAC address is encountered. Some Internet providers require power cycling the modem, others require registering the new address over the phone. Additionally, if this WAN connection is on a network segment with other systems that locate it via ARP, changing the MAC to match and older piece of equipment may also help ease the transition, rather than having to clear ARP caches or update static ARP entries.

Warning: If this firewall will ever be used as part of a *High Availability Cluster*, do not spoof the MAC address.

Maximum Transmission Unit (MTU)

The MTU field, shown in Figure *General WAN Configuration*, can typically be left blank, but can be changed when necessary. Some situations may call for a lower MTU to ensure packets are sized

appropriately for an Internet connection. In most cases, the default assumed values for the WAN connection type will work properly.

Maximum Segment Size (MSS)

MSS, shown in Figure *General WAN Configuration* can typically be left blank, but can be changed when necessary. This field enables MSS clamping, which ensures TCP packet sizes remain adequately small for a particular Internet connection.

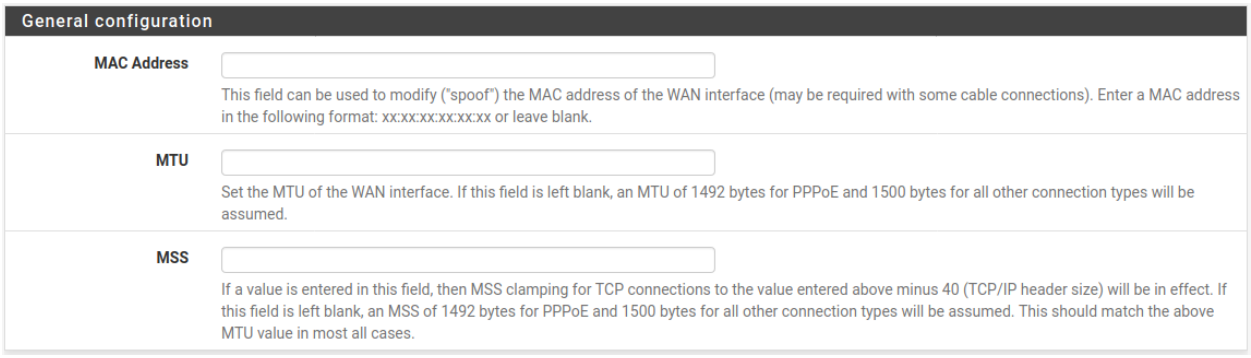


Fig. 5: General WAN Configuration

Static IP Configuration

If the "Static" choice for the WAN type is selected, the **IP address**, **Subnet Mask**, and **Upstream Gateway** must all be filled in (Figure *Static IP Settings*). This information must be obtained from the ISP or whoever controls the network on the WAN side of this firewall. The **IP Address** and **Upstream Gateway** must both reside in the same Subnet.

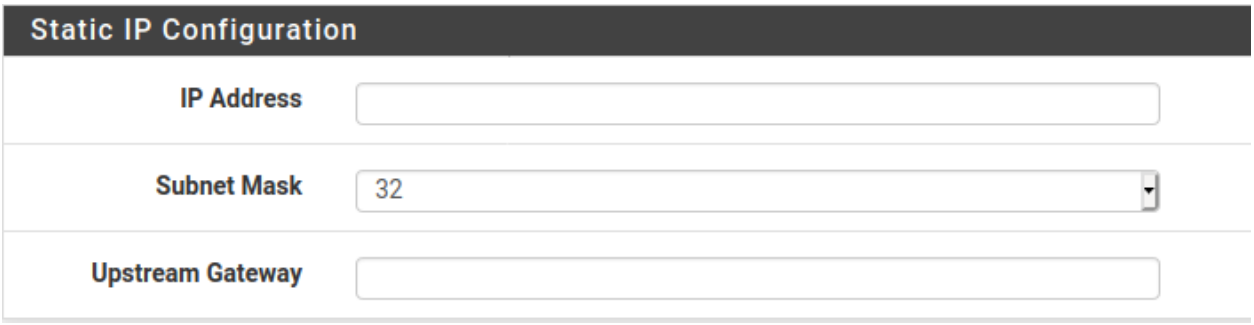


Fig. 6: Static IP Settings

DHCP Hostname

This field (Figure *DHCP Hostname Setting*) is only required by a few ISPs. This value is sent along with the DHCP request to obtain a WAN IP address. If the value for this field is unknown, try leaving it blank unless directed otherwise by the ISP.

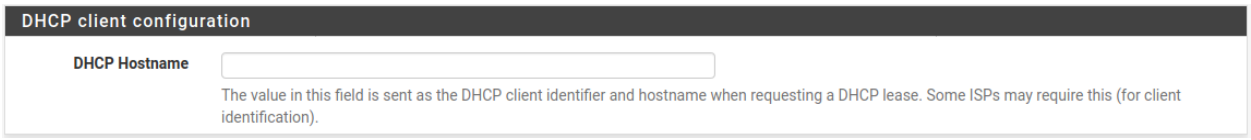


Fig. 7: DHCP Hostname Setting

PPPoE Configuration

When using the PPPoE (Point-to-Point Protocol over Ethernet) WAN type (Figure *PPPoE Configu-*

ration), The **PPPoE Username** and **PPPoE Password** fields are required, at a minimum. The values for these fields are determined by the ISP.

PPPoE Username
The login name for PPPoE authentication. The format is controlled by the ISP, but commonly uses an e-mail address style such as `myname@example.com`.

PPPoE Password
The password to log in to the account specified by the username above. The password is masked by default. To view the entered password, check **Reveal password characters**.

PPPoE Service Name
The PPPoE Service name may be required by an ISP, but is typically left blank. When in doubt, leave it blank or contact the ISP and ask if it is necessary.

PPPoE Dial on Demand
This option leaves the connection down/offline until data is requested that would need the connection to the Internet. PPPoE logins happen quite fast, so in most cases the delay while the connection is set up would be negligible. If public services are hosted behind this firewall, do not check this option as an online connection must be maintained as much as possible in that case. Also note that this choice will not drop an existing connection.

PPPoE Idle Timeout
Specifies how much time the PPPoE connection remain up without transmitting data before disconnecting. This is only useful when coupled with Dial on demand, and is typically left blank (disabled).

Note: This option also requires the deactivation of gateway monitoring, otherwise the connection will never be idle.

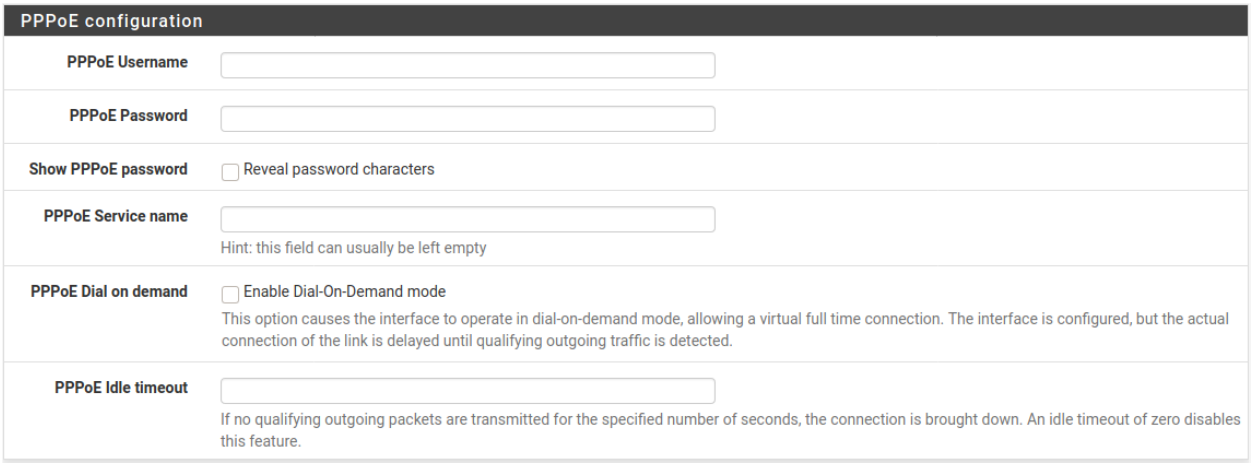


Fig. 8: PPPoE Configuration

PPTP Configuration
The PPTP (Point-to-Point Tunneling Protocol) WAN type (Figure *PPTP WAN Configuration*) is for ISPs that require a PPTP login, **not** for connecting to a remote PPTP VPN. These settings, much like the PPPoE settings, will be provided by the ISP. A few additional options are required:

Local IP Address

The local (usually private) address used by this firewall to establish the PPTP connection.

CIDR Subnet Mask
The subnet mask for the local address.

Remote IP Address
The PPTP server address, which is usually inside the same subnet as the **Local IP address**.

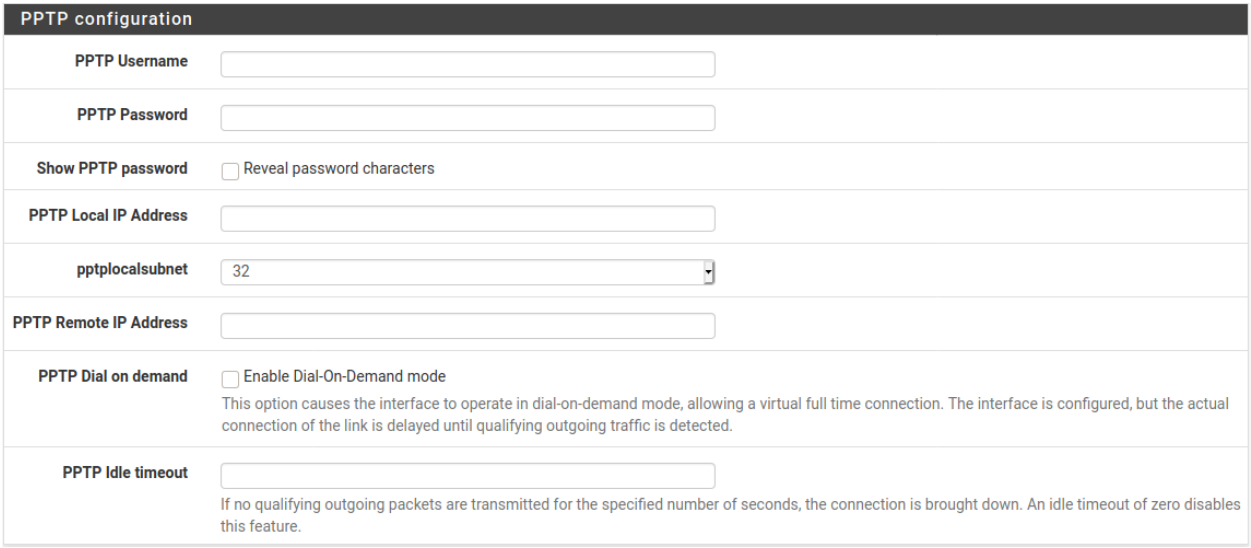


Fig. 9: PPTP WAN Configuration

These last two options, seen in Figure *Built-in Ingress Filtering Options*, are useful for preventing invalid traffic from entering the network protected by this firewall, also known as "Ingress Filtering".

Block RFC 1918 Private Networks
Blocks connections sourced from registered private networks such as `192.168.x.x` and `10.x.x.x` attempting to enter the WAN interface. A full list of these networks is in *Private IP Addresses*.

Block Bogon Networks
When active, the firewall blocks traffic from entering if it is sourced from reserved or unassigned IP space that should not be in use. The list of bogon networks is updated periodically in the background, and requires no manual maintenance. Bogon networks are further explained in *Block Bogon Networks*.

Click  **Next** to continue once the WAN settings have been filled in.

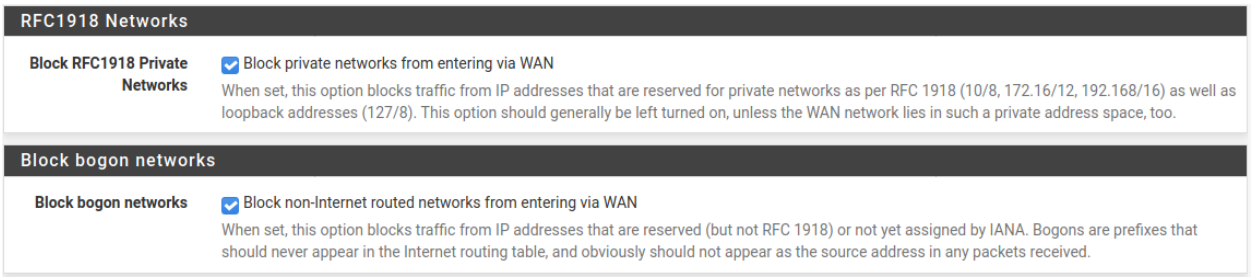


Fig. 10: Built-in Ingress Filtering Options

8.1.4 LAN Interface Configuration

This page of the wizard configures the **LAN IP Address** and **Subnet Mask** (Figure *LAN Configuration*).

If this firewall cannot connect to any other network via VPN, the default `192.168.1.0/24` network may be acceptable. If this network must be connected to another network, including via VPN from remote locations, choose a private IP address range much more obscure than the common default of `192.168.1.0/24`. IP space within the `172.16.0.0/12` RFC 1918 private address block is generally the least frequently used, so choose something between `172.16.x.x` and `172.31.x.x` to help avoid VPN connectivity difficulties.

If the LAN is `192.168.1.x` and a remote client is at a wireless hotspot using `192.168.1.x` (very common), the client will not be able to communicate across the VPN. In that case, `192.168.1.x` is the local network for the client at the hotspot, not the remote network over the VPN.

If the **LAN IP Address** must be changed, enter it here along with a new **Subnet Mask**. If these settings are changed, the IP address of the computer used to complete the wizard must also be changed if it is connected through the LAN. Release/renew its DHCP lease, or perform a "Repair" or "Diagnose" on the network interface when finished with the setup wizard.

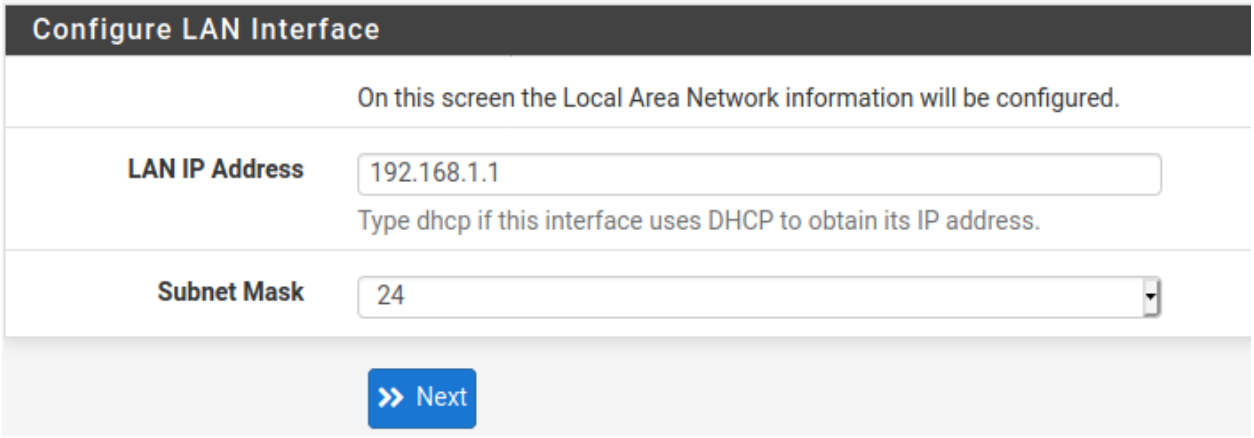


Fig. 11: LAN Configuration

Click » **Next** to continue.

8.1.5 Set admin password

Next, change the administrative password for the GUI as shown in Figure *Change Administrative Password*. The best practice is to use a strong and secure password.

> **Warning:** This password cannot be set to the same value as the username.
>
> Additionally, on pfSense Plus software version 24.03 and later, the password cannot be set to the default value (*Default Username and Password*).

Enter the password in the **Admin Password** and confirmation box to be sure that has been entered correctly.

> **Warning:** On pfSense Plus software version 24.03 and later changing the password is **mandatory**. The wizard will not proceed until the password is changed.

Click » **Next** to continue.

> **Warning:** **Do not leave the password set to the default** `pfsense`. If access to the firewall administration via GUI or SSH is exposed to the Internet, intentionally or accidentally, the firewall could easily be compromised if it still uses the default password.

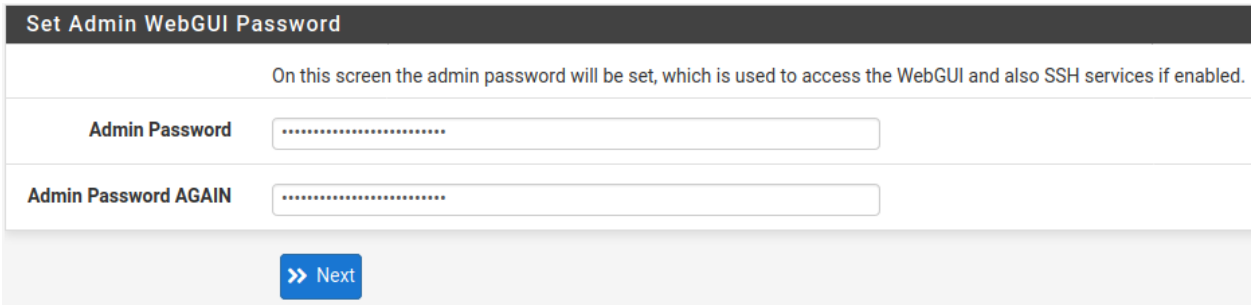


Fig. 12: Change Administrative Password

8.1.6 Completing the Setup Wizard

That completes the setup wizard configuration. Click **Reload** (Figure *Reload the GUI*) and the GUI will apply the settings from the wizard and reload services changed by the wizard.

Tip: If the LAN IP address was changed in the wizard and the wizard was run from the LAN, adjust the client computer's IP address accordingly after clicking **Reload**.

When prompted to log in again, enter the new password. The username remains `admin`.

After reloading, the final screen of the wizard includes convenient links to check for updates, get assistance, and other resources. Click **Finish** to complete and exit the wizard.

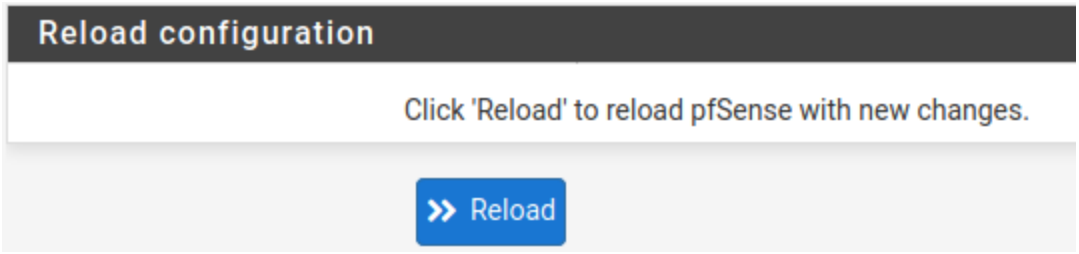


Fig. 13: Reload the GUI

At this point the firewall will have basic connectivity to the Internet via the WAN and clients on the LAN side will be able to reach Internet sites through this firewall.

If at any time this initial configuration must be repeated, revisit the wizard at **System > Setup Wizard** from within the GUI.

8.2 Interface Configuration

Basic aspects of interface configuration within pfSense® software can be performed at the console and in the setup wizard to start, but changes may also be made after the initial setup by visiting pages under the **Interfaces** menu. A few basics are covered here, the details can be found in *Interface Types and Configuration*.

8.2.1 Assign interfaces

Interfaces added after the initial setup may be assigned roles by visiting **Interfaces > Assignments**. There are numerous tabs on that page used for assigning and creating different types of interfaces. The two most commonly used tabs are **Interface assignments** and **VLANs**.

See also:

VLAN configuration is covered in *Virtual LANs (VLANs)*.

The **Interface assignments** tab shows a list of all currently assigned interfaces: WAN, LAN, and any OPTx entries configured on the firewall. Next to each interface is a drop-down list of all network interfaces/ports found on the system. This list includes hardware interfaces as well as VLAN interfaces and other virtual interface types. The MAC address, VLAN tag, or other identifying information is printed alongside the interface name to aid in identification.

The other tabs, much like the **VLAN** tab, are there to create additional interfaces which can then be assigned. All of these interface types are covered in *Interface Types and Configuration*.

To change an existing interface assignment to another network port:

- Navigate to **Interfaces > Assignments**
- Locate the interface to change in the list
- Select the new network port from the drop-down list on the row for that interface
- Click **Save**

To add a new interface from the list of unused network ports:

- Navigate to **Interfaces > Assignments**
- Select the port to use from the drop-down list labeled **Available Network Ports**

- Click  **Add**

This action will add another line with a new OPT interface numbered higher than any existing OPT interface, or if this is the first additional interface, *OPT1*.

8.2.2 Interface Configuration Basics

Interfaces are configured by choosing their entry from under the **Interfaces** menu. For example, to configure the WAN interface, choose **Interfaces > WAN**.

Every interface is configured in the same manner and any interface can be configured as any interface type (Static, DHCP, PPPoE, etc.). Additionally, the blocking of private networks and bogon networks may be performed on any interface. Every interface can be renamed, including WAN and LAN, to a custom name. Furthermore, every interface can be enabled and disabled as desired, so long as a minimum of one interface remains enabled.

See also:

For detailed interface configuration information, see *Interface Types and Configuration*

The **IPv4 Configuration Type** can be changed between *Static IPv4*, *DHCP*, *PPPoE*, *PPP*, *PPTP*, *L2TP*, or *None* to leave the interface without an IPv4 address. When *Static IPv4* is used, an **IPv4 Address**, subnet mask, and **IPv4 Upstream Gateway** may be set. If one of the other options is chosen, then type-specific fields appear to configure each type.

The **IPv6 Configuration Type** can be set to *Static IPv6*, *DHCP6*, *SLAAC*, *6rd Tunnel*, *6to4 Tunnel*, *Track Interface*, or *None* to leave the interface without any IPv6 configuration. When Static IPv6 is selected, set an **IPv6 address**, prefix length, and **IPv6 Upstream Gateway**.

If this interface is wireless, the page will contain many additional options to configure the wireless portion of the interface. Consult *Wireless* for details.

Note: Selecting a **Gateway** from the drop-down list, or adding a new gateway and selecting it, will direct the firewall to treat this interface as a WAN type interface for NAT and related functions. This is not desirable for internal-facing interfaces such as LAN or a DMZ. Gateways may still be utilized on those interfaces for static routes and other purposes *without* selecting a **Gateway** here on the interfaces page.

8.3 Managing Lists in the GUI

The pfSense® software GUI has a common set of icons which are used for managing lists and collections of objects throughout the firewall. Not every icon is used in every page, but their meanings are consistent based on the context in which they are seen. Examples of such lists include firewall rules, NAT rules, IPsec, OpenVPN, and certificates.



Add a new item to a list



Add an item to the beginning of a list



Add an item to the end of a list



Edit an existing item



Copy an item (create a new item based on the selected item)



Disable an active item



Enable a disabled item



Delete an item



Used for moving entries after selecting one or more items. Click to move the selected items above this row. Shift-click to move the selected items below this row.

Sections may have their own icons specific to each area. Consult the appropriate sections of this documentation for specifics about icons found in other parts of the firewall.

Tip: To determine which action an icon will perform, hover over the icon with the mouse pointer and a tooltip will display a short description of the icon's purpose.

8.4 Quickly Navigate the GUI with Shortcuts

Many areas of the GUI have shortcut icons present in the area known as the "Breadcrumb Bar", as seen in Figure *Shortcuts Example*. These shortcut icons reduce the amount of hunting required to locate related pages, allowing a firewall administrator to navigate quickly between the status page of a service, its logs, and configuration. The shortcuts for a given topic are present on each page related to that topic.

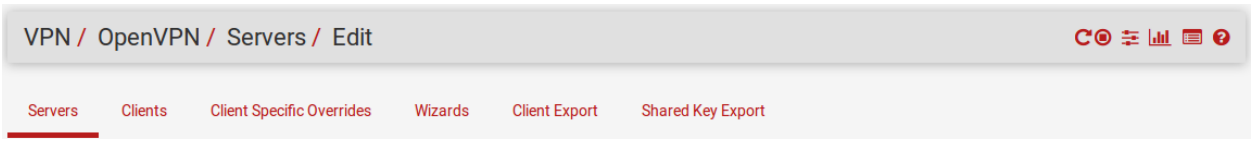


Fig. 14: Shortcuts Example

Note: Shortcut icons only appear when their respective actions are possible and the target pages exist. Not every section has every icon.

The shortcut icons have the following effects when they appear in the GUI:

Start Service
If the service is stopped, this icon starts the service.

Restart Service
If the service is running, this icon restarts the service.

Note: Some services will stop and start, others reload the configuration. Check the documentation of each service for details.

Stop Service
If the service is running, this icon stops the service.

Related Settings
This icon navigates to the **settings** page for this section.

Status Page Link
This icon navigates to the **status** page for this section.

Log Page Link
This icon navigates to the **logs** page for this section.

Help Link
This icon navigates to a related help topic for this page.

The *Service Status* page (**Status > Services**) also has shortcut controls for pages related to each service, as shown in Figure *Shortcuts on Service Status*. The icons have the same meaning as in the above section.

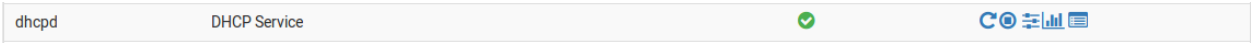


Fig. 15: Shortcuts on Service Status

8.5 General Configuration Options

System > General Setup contains basic configuration options for pfSense® software. A few of these options are also found in the *Setup Wizard*.

Hostname
The **Hostname** is the short name for this firewall, such as `firewall1`, `hq-fw`, or `site1`. The name must start with a letter, and it may contain only letters, numbers, or a hyphen.

Domain
The **Domain** name for this firewall, e.g. `example.com`. If this network does not have a domain, use `<something>.home.arpa`, where `<something>` is another identifier: a company name, last name, nickname, etc. For example, `company.home.arpa`

The **Hostname** and **Domain** name are combined to make up the Fully Qualified Domain Name (FQDN) of this firewall. For example, if the **Hostname** is `fw1` and the **Domain** is `example.com`, then the FQDN is `fw1.example.com`.

8.5.1 DNS Server Settings

Options in this section control how the firewall resolves hostnames using DNS.

Note: The DNS Resolver is active by default and uses **resolver** mode (*DNS Resolver Mode*). When set this way the DNS Resolver does not need forwarding DNS servers as it will communicate directly with root DNS servers and other authoritative DNS servers.

To use the servers in this list, switch the DNS resolver to forwarding mode. The DNS Forwarder (*DNS Forwarder*) is only capable of using forwarding mode and will always use the servers from this list.

DNS Servers

This page configures multiple DNS servers managed as a list. To add more DNS servers, click  **Add DNS Server**.

To remove an entry from the list, click  **Delete**.

The DNS server list may be left blank if the DNS Resolver is active in its default resolver mode. If this firewall has a dynamic WAN type such as DHCP or PPPoE these servers may be automatically assigned by the ISP and can also be left blank.

Each DNS server entry has the following properties:

Address
: The IP address of the DNS Server.

Hostname
: The FQDN of the DNS server, used to validate DNS server certificates when using DNS over TLS (*DNS Resolver Configuration*).

Gateway
: The gateway through which the firewall will reach this DNS server.

  This is useful in a Multi-WAN scenario where, ideally, the firewall will have at least one DNS server configured per WAN. More information on DNS for Multi-WAN can be found in *DNS Forwarding and Static Routes*.

DNS Resolution Behavior

These options fine tune the way the firewall utilizes DNS servers.

DNS Server Override
: When checked, a dynamic WAN ISP can supply DNS servers which override those set manually. To force the use of only the DNS servers on this page, uncheck this option. This does not apply to the DNS Resolver when acting in resolver mode.

DNS Resolution Behavior
: This option controls how the firewall itself resolves DNS queries.

Use Local DNS (127.0.0.1), fall back to remote DNS Servers (Default)
By default, the firewall will consult the DNS Resolver or DNS Forwarder running on this firewall to resolve hostnames for itself. It does this by listing localhost (`127.0.0.1`) as its first DNS server internally. If the local DNS server is unreachable, the firewall will send queries directly to the DNS servers configured on this page, or those received from dynamic WANs.

This method gives the firewall the best chance of having working DNS.

Use Local DNS (127.0.0.1), ignore remote DNS Servers
Like the option above, this option will make the firewall use its own DNS Resolver or DNS Forwarder to resolve hostnames. However, it will not attempt to use any other server.

This option is more secure as it forces DNS to be resolved using the configuration on the DNS Resolver or DNS Forwarder, which may have special requirements restricting or redirecting name resolution. For example, if the DNS Resolver is configured for *DNS over TLS*, using this option ensures that the firewall will not send queries to DNS servers without using TLS.

Use remote DNS Servers, ignore local DNS
This option forces the firewall to use the DNS servers configured on this page or from dynamic WANs, and it will not utilize the local DNS Resolver or DNS Forwarder.

This option is useful when the local DNS service is configured in a strict manner to control client behavior, but the firewall still needs unrestricted access to DNS for tasks such as updates and installing packages.

8.5.2 Localization

Options in this section control the firewall clock and language.

Time Zone
The time zone used by the firewall for its clock. Choose a geographically named zone which best matches location of this firewall, or a common zone such as UTC. The firewall clock, log entries, and other areas of the firewall base their time on this zone.

Note: Changing the zone requires a reboot to fully activate the new zone in all areas of the firewall.

Tip: Avoid using the GMT +/- zones as they do not operate in an intuitive manner. See *Troubleshooting Time Zone Configuration* for more information.

Time Servers
Network Time Protocol (NTP) server hostnames or IP addresses. Unless a specific NTP server is required, such as one on LAN, the best practice is to leave the **Time Servers** value at the default `2.pfsense.pool.ntp.org`. This value will pick random servers from a pool of known-good IPv4 and IPv6 NTP hosts.

To utilize multiple time servers or pools, add them in the same box, separating each entry by a space. For example, to use three NTP servers from the pool, enter:

```
0.pfsense.pool.ntp.org 1.pfsense.pool.ntp.org 2.pfsense.pool.ntp.org
```

This numbering is specific to how `.pool.ntp.org` operates and ensures each address is drawn from a unique pool of NTP servers, so the same server does not get used twice.

Language
The language used by the GUI. The GUI has been translated into multiple languages in addition to the default *English* language.

8.5.3 webConfigurator

Options in this section control various behaviors of the web-based GUI, which can be referred to as the GUI, WebGUI, or webConfigurator.

Theme
The **Theme** controls the look and feel of the GUI. Several themes are included in the base system, and they only make cosmetic not functional changes to the GUI.

Top Navigation
This option controls the behavior of the menu bar at the top of each page. There are two possible choices:

Scrolls with page
The default behavior. When the page scrolls, the navigation remains at the top of the *page*, so it is no longer visible as it scrolls off the top of the window.

This is the best option for most situations.

Fixed
When selected, the navigation remains fixed at the top of the *window*, always visible and available for use.

This behavior can be convenient, but can be problematic on smaller screens such as tablets and mobile devices. On low resolution browsers long menus can be cut off, leaving options at the bottom unreachable.

Hostname in Menu
Chooses if and how the GUI includes the firewall hostname in the menu. This can aid in quickly identifying a firewall when managing multiple firewalls in separate tabs or windows, but it consumes extra space in the menu.

Default (No hostname)
The GUI does not display the hostname or FQDN in the menu.

Hostname Only
When set, the GUI includes the firewall **Hostname** (no domain name) in the menu.

If all firewalls are in the same domain, or if they have unique hostnames, this may be sufficient.

Fully Qualified Domain Name
When set, the GUI includes the Fully Qualified Domain Name of the firewall in the menu.

This takes more space than displaying the hostname portion alone, but may be necessary to properly distinguish firewalls if they use similar hostnames in multiple domains.

Dashboard Columns
The dashboard is limited to 2 columns by default. On wider displays, additional columns can utilize extra horizontal screen space. The maximum number of columns is 4.

Interfaces Sort
When unset (default), the GUI presents interfaces in their natural order from the configuration. This

is critical for functions such as High Availability which require specific interface ordering. When this option is set, the GUI sorts the interface list alphabetically.

Associated Panels Show/Hide
A few GUI pages contain collapsible panels with settings or functions. These panels take up extra screen space so they are hidden by default. For firewall administrators who use the panels frequently, this can be slow and inefficient. The options in this group make the GUI show these panels by default instead of hiding them.

Available Widgets
Controls the **Available Widgets** panel on the Dashboard.

Log Filter



Controls the log filtering ( ) panel used for searching log entries under **Status > System Logs**.

Manage Log



Controls the per-log settings in the **Manage Log** ( ) panel available for each log under **Status > System Logs**.

Monitoring Settings
Controls the options panel used to change the graphs at **Status > Monitoring**.

Require State Filter
When set, the state table contents at **Diagnostics > States** are suppressed by the GUI unless a filter string is present. This helps the GUI handle large state tables which otherwise may fail to load.

Left Column Labels
When checked, the option labels in the left column are set to toggle options when clicked. This can be convenient if the firewall administrator is used to the behavior, but it can also be problematic on mobile or in cases when the behavior is unexpected.

Alias Popups
When set, the tooltip presented by the GUI when hovering over an alias in a rule list only shows the alias description. When unset, the contents of the alias are included in the tooltip. For firewalls with large aliases, this may cause performance or browser rendering issues.

Disable Dragging
When set, the GUI disables drag-and-drop on rule lists. Most users find drag-and-drop to be convenient and beneficial, thus the feature is enabled by default. Users who find the behavior undesirable can set this option.

Login Page Color
Controls the color of the login page, which is independent of the theme.

Login Hostname
When set, the GUI includes the hostname on the login form.

> **Warning:** This can be considered a security risk since it exposes information about the firewall to users who have not yet authenticated. If the firewall GUI is only reachable by authorized management clients, the convenience may outweigh the potential risk.

8.6 Advanced Configuration Options

System > Advanced contains numerous options to customize behavior for more complex environments. Most administrators will not need to adjust these options for basic deployments.

Some of these options are covered in more detail in other sections of the documentation where their discussion is more topical or relevant, but they are all mentioned here with a brief description.

8.6.1 Admin Access

The options on **System > Advanced**, **Admin Access** tab govern various methods for administering the firewall, including via the web interface, SSH, serial, and physical console.

webConfigurator (GUI)

Protocol

The protocol for connections between web browsers and the GUI. May be one of:

HTTP
: Plain unencrypted HTTP. Insecure and basic, but widely compatible. Should not be used in most cases, and should **never** be exposed to insecure networks.

HTTPS (SSL/TLS)
: Encrypted ("Secure") HTTP. Protects communication between the client browser and the firewall GUI. Requires an **SSL/TLS certificate** to function. Depending on the browser and certificate configuration, there may be compatibility issues, but typically these are easily overcome, such as by clicking through warnings about self-signed certificates.

Note: The best practice is to use HTTPS so **only** encrypted traffic is exchanged between the GUI and clients.

SSL/TLS Certificate

The **SSL/TLS Certificate** to be used by the GUI in **HTTPS (SSL/TLS)** mode.

The firewall automatically generates a default self-signed certificate on the first boot and again if a referenced certificate is missing. A self-signed certificate is not ideal, but it does still encrypt traffic and is better than communicating without encryption.

The primary disadvantage of a self-signed certificate is the lack of assurance of the identity of the host, since the certificate is not signed by a Certificate Authority trusted by the browser. Additionally, because for the bulk of Internet users such an invalid certificate should be considered a risk, modern browsers may restrict how such certificates are handled. Firefox, for example, displays a warning screen and forces the user to import the certificate and allow a permanent exception. Chrome shows a warning screen with a link to continue.

Tip: To use an externally signed SSL certificate and key, import them using the *Certificate Manager*, then select the certificate here.

Tip: The *ACME Package* can utilize the free Let's Encrypt service to automatically obtain and update a signed certificate for the GUI or for other purposes on the firewall.

Tip: Another alternate technique is to generate a self-signed CA and then generate a GUI certificate from that CA. Export the CA from the firewall and then import that CA into client browsers manually. Using this method, all certificates signed by that CA will be trusted by browsers. Specifics vary by client platform.

Tip: To generate a new self-signed certificate for the GUI, connect using the console or ssh and from a shell prompt, run the following command:

```
# pfSsh.php playback generateguicert
```

TCP Port

The port upon which the GUI listens for incoming connections from browsers. By default, the GUI uses HTTPS on port `443` with a redirect from port `80` for the best compatibility and ease of initial configuration. To change the port, enter a new port number into the **TCP Port** field.

Danger: **Do not** expose the GUI to untrusted networks such as the Internet even on non-default ports.

Some administrators move the GUI to an alternate port for security by obscurity reasons. Such practices **should not** be considered as offering any security benefit. It is trivial for port scanners to locate the service on any open port.

Tip: Moving the GUI to another port will free up the standard web ports for use with port forwards or other services such as HAProxy.

Max Processes

The number of web server worker processes used by the GUI when listening for client browser connections. The default value is `2`.

If multiple administrators view the GUI at the same time and pages take too long to load, or are failing to load, then increase the **Max Processes** value.

WebGUI Redirect

Controls whether the firewall runs a redirect on port `80` so that if a browser attempts to access the firewall with HTTP, the firewall will accept the request and then redirect the browser to the **TCP Port** used by the GUI (e.g. HTTPS on port `443`).

The redirect is enabled by default for ease of access and compatibility.

Disabling the redirect allows another daemon to bind to port `80`.

HSTS

Controls whether the GUI web server sends the `Strict-Transport-Security` HTTPS response header (HSTS) to the browser. Check this box to disable the behavior.

HSTS forces the browser to use only HTTPS for any future requests to the firewall FQDN to ensure it does not accidentally downgrade to an unencrypted connection.

> **Warning:** When disabling HSTS, clients which visited the GUI when HSTS was enabled must perform browser-specific steps for the change to take effect. Consult browser documentation for information on clearing cached HSTS behavior.

OCSP Must-Staple

Controls whether the GUI web server forcefully enables OCSP Stapling.

If the GUI **SSL/TLS Certificate** requires OCSP Stapling, this behavior is automatically enabled by the GUI web server. If the certificate property cannot be automatically determined by the firewall, this option can force the behavior.

Tip: Import the full CA and certificate chain or this option will be ignored by the GUI web server.

WebGUI Login Autocomplete

Controls whether the login form allows autocomplete so that browsers can save login credentials locally, for convenience.

In high-security environments, such as those that must adhere to specific security compliance standards, that behavior is not acceptable.

Note: This only controls autocomplete on the login form.

Warning: Few modern browsers respect this option. Many still offer to save passwords even when the form specifies that the browser must not allow the behavior. This behavior must be controlled or changed using browser options.

GUI login messages

Lowers the system log level for successful GUI login events.

By default, login events are logged at an emergency level and on hardware with a PC speaker, these emergency console messages generate a beep from the speaker.

When this is checked, successful logins to the GUI will be logged as a lower non-emergency level. This lower level will not trigger console bells and may help with processing log events on remote syslog servers.

Note: Changing this option is not necessary if the only goal is suppressing speaker beeps. The console bell behavior can be controlled independently on the *Notifications tab*.

Anti-lockout

Controls whether the firewall adds special rules to permit access to the GUI port and SSH port on the LAN interface by default.

These special rules override user-defined filter rules and prevent the user from accidentally locking themselves out of the firewall GUI or SSH. To control which LAN IP addresses may access the GUI and SSH using firewall rules, disable the anti-lockout rules.

When two or more interfaces are present, the firewall puts anti-lockout rules on the LAN interface; If only one interface is configured, the firewall places rules on that interface instead.

Warning: Filter rules must be in place to allow GUI access before enabling this option! If the LAN rules do not allow access to the GUI, removing the anti-lockout rule will block access to the GUI, potentially leaving the administrator without a means to reach the firewall.

Note: Resetting the LAN IP address from the console also resets the anti-lockout rule. If administrative access is unavailable after enabling this option, choose the console menu option `2`, then choose to set the LAN IP address, and enter in the exact same IP address and accompanying information.

DNS Rebind Check

Controls whether the DNS resolver or forwarder performs DNS rebinding checks. These checks prevent the firewall from receiving DNS responses containing private IP addresses from DNS servers to prevent DNS rebinding attacks.

Note: When accessing the firewall by IP address, these checks are not enforced because the attack is only relevant when using a hostname.

Check this box to disable DNS rebinding protection if it interferes with GUI access or name resolution.

See also:

More detail on DNS rebinding attacks may be found on Wikipedia.

The most common case for disabling DNS rebinding checks is when the firewall is set to use an internal DNS server which will return private (RFC1918) answers for hostnames.

Tip: Instead of disabling all DNS rebinding protections, the checks can be selectively disabled on a per-domain basis in the DNS Resolver or DNS Forwarder. See *DNS Resolver* and *DNS forwarder*.

Note: The DNS Resolver or Forwarder must be restarted after changing this option.

Browser HTTP_REFERER enforcement

Controls whether the GUI checks and enforces `HTTP_REFERER` request header contents.

The GUI checks the referring URL sent by a client browser to ensure that the form was submitted from this firewall. This check prevents a form on another site from submitting a request to the firewall, changing an option when the administrator did not intend for that to happen.

This also breaks some convenience behaviors, such as having a page that links to various firewall devices, though the benefits of the check typically outweigh the advantage of those behaviors.

Alternate Hostnames

A list of **Alternate Hostnames** for the firewall allowed by **DNS Rebind Checks** and **HTTP_REFERER Enforcement**. To keep these features active, but alter their behavior slightly, add **Alternate Hostnames**.

By default, the GUI allows access to the hostname configured on the firewall and all IP addresses configured on the firewall. Hostnames in this field are allowed by the firewall for GUI access and for referring URL purposes.

Browser Tab Text

By default, the GUI prints the firewall hostname first in the page/tab title, followed by the page name. To reverse this behavior and show the page name first and hostname second, check **Display page name first in browser tab**.

Administrators who access many firewalls at the same time in separate tabs tend to prefer having the hostname first (default). Administrators who access one firewall with many pages in separate tabs may prefer having the page name first.

Secure Shell (SSH)

The Secure Shell (SSH) server provides remote console access and file management. A user can connect with any standard SSH client, such as the OpenSSH command line ssh client, PuTTY, SecureCRT, or iTerm2.

When using SSH, both the `admin` username and `root` username are accessible using the `admin` account credentials.

Users in the User Manager that have the `User - System - Shell account access` privilege are also allowed to log in over SSH. These users do not have `root` access privileges, and do not print the menu when they log in because many of the options require `root` privileges.

Tip: To grant users additional shell privileges, use the *sudo package*.

File transfers to and from the firewall are also possible by using a Secure Copy (SCP) client such as the OpenSSH command line `scp`, FileZilla, WinSCP or Fugu. To use SCP, connect as the `root` or `admin` user. If a custom user has the `User - System - Copy files` permission, or all access, then they may also utilize SCP.

Tip: SSH clients must be kept up-to-date. As time goes on, security standards evolve and the SSH server settings utilized by SSH servers will change. Outdated clients may not be able to connect using the strong security keys and algorithms required by `sshd`. If a client cannot connect, check for an update from the vendor.

Enable Secure Shell

To enable the SSH daemon, check **Enable Secure Shell**. After saving with this option enabled, the firewall will generate SSH keys if they are not already present and then start the SSH daemon.

SSHd Key Only

This option controls which authentication methods the SSH daemon allows for clients. It can be set to one of the following values:

Password or Public Key
: Allows a user to authenticate with **either** a valid password or valid key. This is the default behavior.

Public Key Only
: Restricts authentication to only valid keys, passwords are not allowed.

Require Both Password and Public Key
: Requires a valid password **and** a valid key.

Key-based logins are a much more secure practice, though it does take more preparation to configure.

Add user keys for key-based login by editing users in the **User Manager** (*User Management and Authentication*). When editing a user, paste the allowed public keys into the **Authorized Keys** text field for the account.

Allow Agent Forwarding

Controls whether the SSH daemon allows agent forwarding for clients.

Agent forwarding allows a user to run an SSH agent on their client system and connect to the firewall, and then to other remote SSH servers using the key from their agent. In this case, the user does not need to have their private keys on the firewall but can still use key-based authentication to remote servers.

Use of an SSH agent can be considered a security issue in certain cases. Additionally, the firewall is not intended to be a general purpose SSH client or intermediate system, thus this feature is disabled by default.

SSH Port

Controls the port used by the SSH daemon to accept client connections. To change the port, type the new port into the **SSH Port** box.

Moving the SSH server to an alternate port provides a negligible security improvement, and frees up the port for other uses.

Tip: Brute force SSH scanners focus on hitting TCP port `22` but if the daemon is open to the Internet on another port, it will still be found and hit by scanners.

Best Practices for SSH

If this firewall is installed in an environment that requires leaving SSH access unrestricted by firewall rules, which is dangerous, the best practice is to take one of the following actions:

Change the SSH Port
: Moving to a random alternate port prevents log noise from many, but not all, brute-force SSH login attempts and casual scans. It can still be found with a port scan, however.

Force Key-Based Authentication
: Key-based authentication must always be used by publicly accessible SSH servers to eliminate the possibility of successful brute force attacks. Set **SSHd Key Only** to either *Public Key Only* or *Require Both Password and Public Key*.

Multiple unsuccessful logins from the same IP address will result in locking out the IP address trying to authenticate, but that alone is not considered sufficient protection.

Login Protection

pfSense software utilizes the `sshguard` daemon to protect against brute force logins for both the GUI and SSH connections. The options in this section fine-tune the behavior of this protection.

Threshold
: The total score value above which `sshguard` will block clients. Most attacks have a score of `10`, the default threshold value is `30`.

Blocktime
: The initial minimum number of seconds to block attackers who have exceeded the **Threshold** value. The default value is `120` seconds. Repeat offenders are blocked for increasingly longer amounts of time (1.5x for each repetition).

  Note: Attackers are unblocked at random intervals, so actual block time will be longer than stated. This prevents clients from predicting the timing to optimize targeted attacks.

Detection Time
: The amount of time, in seconds, attackers are remembered by `sshguard` since their last offense before it resets their score. Default is `1800` seconds.

Pass list
: A list of subnets which are excluded from login protection. This lowers security but is generally acceptable from specific secure management networks.

For example, it may be necessary to add entries for network monitoring systems which probe the SSH port but do not log in. Otherwise, such systems may be flagged as attackers.

Click  **Add address** to display additional form field lines for pass list entries.

Serial Communications

If the firewall hardware contains a hardware serial port, it can be used for a console connection. This is useful for running on hardware without video hardware or if it will be running "headless" (without keyboard and video attached). In these cases, the serial console can be enabled to maintain physical control, so long as the hardware has a viable serial port. The hardware visible to the operating system for such serial ports cannot be USB. Though some devices expose an external physical USB port dedicated to the console, typically such devices are wired to appear to the OS as a traditional non-USB serial port internally.

If hardware is detected which has no VGA port, the serial console is forced on and cannot be disabled, and the serial options are all hidden except for the speed.

Serial Terminal

When **Serial Terminal** is set, the operating system enables the console on the first serial port. This console will receive kernel boot messages and a menu after the firewall has finished booting. This will not disable the onboard keyboard and video console.

To connect to the serial console, use a null modem cable connected to a serial port or adapter on another PC or serial device.

See also:

For more information on connecting to a serial console, see *Connecting to a Serial Console* and *Start a Serial Client*.

When making any changes to the serial console, the firewall must be rebooted before they take effect.

Note: In some systems booting via EFI, the serial console can be activated or deactivated without a reboot.

Serial Console Speed

The default serial console speed is *115200* bps and almost all hardware works well at that speed. In rare cases, a slower speed may be required which can be set here by picking the desired speed from the **Serial Speed** drop-down.

When upgrading from an older version, this may remain at an older value such as *9600* or *38400* to match the BIOS on older hardware. Increasing the speed to *115200* is almost always safe and more useful than slower speeds.

Primary Console

On hardware with both the serial console enabled and a video port available, the **Primary Console** selector chooses which is the preferred console, so it will receive the boot log messages. Other OS kernel messages will show up on all console connections, and both consoles will have a usable menu.

In cases where the boot cannot complete, the preferred console must be used to resolve the problem, such as reassigning interfaces.

Console Menu

Normally the firewall always presents the menu on the console, and the menu will be available as long as someone has physical access to the console. In high-security environments this is not desirable.

This option adds password protection to the console. The console accepts the same usernames and passwords used to access the GUI. After setting this option, the firewall must be rebooted before it takes effect.

Note: While this will stop accidental key presses and keep out casual users, this is by no means a perfect security method. A knowledgeable person with physical access can still reset the passwords (see *Forgotten Password with a Locked Console*). Consider other physical security methods if console security is a requirement.

8.6.2 Firewall & NAT

The options on **System > Advanced**, **Firewall & NAT** tab control various aspects of how the firewall processes packets and connections.

Packet Processing

IP Do-Not-Fragment compatibility

This option is a workaround for operating systems which generate fragmented packets with the "don't fragment" (DF) bit set. Linux NFS (Network File System) is known to do this, as well as some VoIP implementations.

When this option is enabled, the firewall will not drop these malformed packets, but instead it will clear the DF bit. The firewall will also randomize the IP identification field of outgoing packets to compensate for operating systems that set the DF bit but set a zero IP identification header field.

IP Random ID generation

If **Insert a stronger ID into IP header of packets passing through the filter** is checked, the firewall replaces the IP identification field of packets with random values to compensate for operating systems that use predictable values. This option only applies to packets that are not fragmented after the optional packet reassembly.

Firewall Optimization Options

The optimization mode controls how the firewall expires state table entries:

Normal
: The standard optimization algorithm, which is optimal for most environments.

High Latency
: Used for high latency links, such as satellite links. Expires idle connections later than default.

Aggressive
: Expires idle connections quicker. More efficient use of CPU and memory but can drop legitimate connections earlier than expected. This option can also improve performance in high traffic deployments with lots of connections, such as web services.

Conservative
: Tries to avoid dropping any legitimate connections at the expense of increased memory usage and CPU utilization. Can aid in environments that require long-lived but mostly idle UDP connections, such as VoIP.

The table *Firewall Optimization Details* contains the values chosen by PF for each optimization algorithm. The values are taken from the PF source code. The first line is the raw value, second line is human-readable:

Table 1: Firewall Optimization Details

	Normal	High Latency	Conservative	Aggressive
tcp.first First TCP packet	60 1min	180 3min	3600 60min	30 30sec
tcp.opening No response yet	30 30sec	35 35sec	900 15min	5 5sec
tcp.established Established	86400 24h	86400 24h	432000 5days	18000 5h
tcp.closing Half closed	900 15min	905 15min + 5sec	3600 1h	60 60sec
tcp.finwait Got both FINs	45 45sec	50 50sec	600 10min	30 30sec
tcp.closed Got an RST	90 90sec	95 95sec	180 3min	30 30sec
tcp.tsdiff Allowed TS diff	30 30sec	60 60sec	60 60sec	10 10sec

Disable Firewall Scrub

When set, the scrubbing option in PF is disabled. The `scrub` action in PF can interfere with NFS, and in rare cases, with VoIP traffic as well. By default, the firewall uses the `fragment reassemble` option which reassembles fragmented packets before sending them on to their destination, when possible. More information on the `scrub` feature of PF can be found in the OpenBSD PF Scrub Documentation.

Note: Disabling `scrub` also disables other features that rely on `scrub` to function, such as DF bit clearing and ID randomization. Disabling `scrub` does not disable MSS clamping if it is active for VPNs, or when an MSS value is configured on an interface.

Firewall Adaptive Timeouts

Adaptive Timeouts control state handling in PF when the state table is nearly full. Using these timeouts, a firewall administrator can control how states are expired or purged when there is little or no space remaining to store new connection states.

Adaptive Timeouts are enabled by default and the default values are calculated automatically based on the configured **Firewall Maximum States** value.

Adaptive Start
: Adaptive scaling is started once the state table reaches this level, expressed as a number of states. **Adaptive Start** defaults to 60% of **Firewall Maximum States**.

Adaptive End
: When the size of the state table reaches this value, expressed as a number of state table entries, all timeout values are assumed to be **zero**, which causes PF to purge all state entries immediately. This setting defines the scale factor, it should be set greater than the total number of states allowed. **Adaptive End** defaults to 120% of **Firewall Maximum States**.

When the number of connection states exceeds the threshold set by **Adaptive Start**, timeout values are scaled linearly with factor based on the number of states used between the Start and End state counts. The timeout adjustment factor is calculated as follows: (Number of states until the **Adaptive End** value is reached) / (Difference between the **Adaptive End** and **Adaptive Start** values).

Note: As an example, consider a firewall with **Adaptive Start** set to `600000`, **Adaptive End** set to `1200000` and **Firewall Maximum States** set to `1000000`. In this situation, when the state table size reaches `900000` entries the state timeouts will be scaled to 50% of their normal values.

(1,200,000 - 900,000) / (1,200,000 - 600,000) = 300,000 / 600,000 = 0.50, 50%

Continuing the example, when the state table is full at 1,000,000 states the timeout values will be reduced to 1/3 of their original values.

Tip: The state table usage indicator on the dashboard will change color and text when the state table size crosses these thresholds.

Firewall Maximum States

This value is the maximum number of connections the firewall can hold in its state table. The default size is calculated based on 10% of total RAM. This default value is sufficient for most installations, but can be adjusted higher or lower depending on the load and available memory.

Each state consumes approximately 1 KB of RAM, or roughly 1 MB of RAM for every 1000 states. The firewall must have adequate free RAM to contain the entire state table before increasing this value. Firewall states are discussed further in *Stateful Filtering*.

Tip: On a firewall with 8GB of RAM the state table would have a default size of approximately 800,000 states. A custom **Firewall Maximum States** value of 4,000,000 would consume about 4GB of RAM, half the available 8GB total.

Firewall Maximum Table Entries

This value defines the maximum number of entries that can exist inside of address tables used by the firewall for collections of addresses such as aliases, ssh/GUI lockout records, hosts blocked by snort alerts, and so on. By default, this is 400,000 entries. If the firewall has features enabled which can load large blocks of address space into aliases such as URL Table aliases or the pfBlockerNG package, then increase this value to comfortably include at least double the total amount of entries contained in all aliases combined.

This maximum value of this setting depends on the amount of RAM available hold the entries. RAM usage for each alias entry is similar to, but less than, the state table. A safe assumption is approximately 1K of memory per entry to be conservative.

If the value of this setting is not large enough to contain the entire content all the aliases, the firewall may fail to load the ruleset. The aliases must fit in **twice** in the total area because of the way the firewall loads and reloads aliases; The firewall loads the new list alongside the old list and then removes the old list for a smoother transition.

Firewall Maximum Fragment Entries

When scrub is enabled the firewall maintains a table of packet fragments waiting to be reassembled. By default, this table can hold `5000` fragments. In rare cases a network may have an unusually high rate of fragmented packets which can require more space in this table. When this limit is reached, the following log message will appear in the main system log:

```
kernel: [zone: pf frag entries] PF frag entries limit reached
```

VPN Packet Processing

These settings affect traffic behavior with some types of VPNs, including IPsec and OpenVPN. It also affects the PPPoE Server.

IP Do-Not-Fragment compatibility

This option is the same as *IP Do-Not-Fragment compatibility*, but this option only applies that behavior to VPN networks.

IP Fragment Reassemble

Reassemble IP Fragments for normalization. In this case, fragments are buffered until they form a complete packet, and only the completed packet is passed on to the filter. The advantage is that filter rules have to deal only with complete packets, and can ignore fragments. The drawback of caching fragments is the additional memory cost and potential introduction of latency.

MSS Clamping

Enable maximum segment size clamping on TCP flows over IPsec tunnels. This helps overcome problems with path MTU discovery (PMTUD) on IPsec VPN links.

This is useful if large TCP packets have problems traversing the VPN, or if slow/choppy connections across the VPN are observed by users. Ideally it should be set to the same value on both sides of the VPN, but traffic will have MSS clamping applied in both directions.

Enable
When set, the **Maximum MSS** option is available and its value is used by the firewall configuration.

Note: For IPsec using VTI mode phase 2 entries, set the MSS value in the interface configuration for the assigned VTI interface under the **Interfaces** menu.

Maximum MSS
The maximum segment size set in TCP packets flowing across IPsec VPN tunnels. Defaults to `1400`. Must be low enough to account for the overhead of IPsec and the MTU of the link, but no so low that unnecessarily small segments are sent as that can be inefficient.

Advanced Options

Disable Firewall

When **Disable all packet filtering** is set, the firewall becomes a routing-only platform. This is accomplished by disabling PF entirely, and as a consequence, NAT is disabled since it is also handled by PF.

Tip: To disable *only* NAT, do not use this option. Consult *Disabling Outbound NAT* for more information on controlling outbound NAT behavior.

Firewall State Policy

Controls the default State Policy for states created by firewall rules. These policies fundamentally change how the firewall checks packets against existing state table entries to determine if a packet should be allowed. There are two options available, **Interface Bound States** and **Floating States**.

Note: There is no difference in the ability to view or kill states between either mode.

Tip: While this option controls the global default, there is a per-rule option to override this behavior as well. See *State Policy* for details.

Interface Bound States

Interface Bound States are more strict and secure. States are bound to specific interfaces by their OS/driver name (e.g. `igcX`). If a packet attempts to take a path through a different interface than the one to which it is bound, the packet is dropped.

As this is the most secure option, it is currently the default policy.

This policy is less likely to allow VPN or other traffic to "leak" or egress via unexpected paths (e.g. during interface events).

This policy has some drawbacks, however, as noted in the following subsections.

IPsec VTI Filtering

For firewalls utilizing IPsec VTI tunnels, due to the way the OS handles traffic on VTI interfaces in the default **IPsec Filter Mode** packets may appear to enter and exit on different interfaces (e.g. `ipsecX` vs `enc0`). This can cause issues with Interface Bound states but pass traffic OK with floating states.

This situation can be worked around a couple different ways:

Change IPsec Filter Mode
: The most secure option is only viable in certain cases. On installations which **only** contain VTI tunnels and **no** policy-based tunnels or mobile IPsec configurations, switch the **IPsec Filter Mode** option in the *Advanced IPsec Settings* to *Filter IPsec VTI and Transport on assigned interfaces, block all tunnel mode traffic*.

  When configured in this way, PF can track states and packets on VTI interfaces appropriately, but it is **not** compatible with tunnel mode or mobile IPsec.

Rules with Floating Policy Set
: The most compatible method is to add rules using the floating policy for IPsec VTI traffic:

  - Edit rules on the **IPsec** tab and change the per-rule *State Policy* in the **Advanced Options** to *Floating States* on each rule that matches inbound traffic from VTI peers.
  - Add new rule(s) on to the top of the **Floating** tab to match outbound traffic for VTI peers, for example:

    Action
    : *Pass*

    Quick
    : *Checked*

    Interface
    : *IPsec* (`enc0`)

    Direction
    : *Out*

    Advanced Options: State Policy
    : *Floating States*

    Set the other parameters to match the VTI peer traffic, or set the rule for *Any* protocol, source, and destination to match any outbound IPsec traffic.

Change Default State Policy
: Switching the default policy back to Floating states will allow it to work, but with the security implications mentioned previously.

High Availability State Synchronization

For High Availability configurations utilizing *state synchronization (pfsync)*, all nodes in the cluster must either be using identical hardware and physical interface assignments, or be using abstracted interfaces such as `laggX` that can mask hardware differences. Otherwise, synchronized states will never match traffic on the secondary node, making state synchronization ineffective. See *pfsync and Physical Interfaces* for more information on this limitation.

Floating States

Floating States are less secure, more lenient in their checks, and are not strictly associated with any interface. The interface is tracked in state properties, but it is informational and not enforced.

This policy is more forgiving when it comes to multi-WAN and asymmetric routing scenarios. It also allows HA nodes with different hardware to utilize state synchronization.

The most significant drawback to this relaxed policy is that it might allow packets to be misdirected or take unexpected paths. This is especially the case if routing can be manipulated in some way (e.g. by a dynamic WAN). This can be especially problematic for secure traffic, such as between local systems on separate interfaces, or for VPN traffic.

State Policy History

The state policy behavior in PF has changed over time as it changed in the base OS, but the option to explicitly choose the default behavior did not exist until recently.

Very old versions of pfSense software (2.1.5 and before) behaved in the "floating" style.

From pfSense software version 2.2 until pfSense Plus software version 21.05.2/CE 2.5.2 the behavior was closer to "interface bound" but not identical.

From pfSense Plus software version 22.01/CE 2.6.0 until pfSense Plus software version 23.09.1/CE 2.7.2 the behavior was closer to "floating".

Starting with pfSense Plus software version 24.03/CE 2.8.0 the default is explicitly set to "interface bound" for increased security.

The ability to change the policy, as well as the change in default policy, will also be available to some previous releases via the recommended patches mechanism in the *System Patches* package.

Ethernet Filtering (Plus Only)

Enable experimental rule-based pass/block filtering of packets based on Ethernet (Layer 2) header attributes (e.g. MAC addresses). These rules are processed before other (L3) rules on the inbound direction, and after those rules outbound. When enabled, Ethernet rules are managed on their own tab at **Firewall > Rules**, **Ethernet tab**.

Note: This feature is exclusive to pfSense Plus version 23.05 and later.

See also:

- *Ethernet (Layer 2) Rules*
- *Ordering of NAT and Firewall Processing*

Static Route Filtering

The **Bypass firewall rules for traffic on the same interface** option applies if the firewall has one or more static routes defined. If this option is enabled, traffic that enters and leaves through the same interface will not be checked by the firewall. This may be required in situations where multiple subnets are connected to the same interface, to avoid blocking traffic that is passed through the firewall in one direction only due to asymmetric routing. See *Bypass Firewall Rules for Traffic on Same Interface* for a more in-depth discussion on that topic.

Disable Auto-added VPN rules

By default, when IPsec is enabled firewall rules are automatically added to the appropriate interface which will allow the tunnel to establish. When **Disable Auto-added VPN rules** is checked, the firewall will not automatically add these rules. By disabling these automatic rules, the firewall administrator has control over which addresses are allowed to connect to a VPN. Further information on these rules can be found at *VPNs and Firewall Rules*.

Disable Reply-To

In a Multi-WAN configuration the firewall has a beneficial default behavior that ensures traffic leaves the same interface through which it arrived. This is accomplished using the PF keyword `reply-to` which is added automatically to interface tab firewall rules for WAN-type interfaces. When a connection matches a rule with `reply-to`, the firewall remembers the path through which the connection was made and routes the reply traffic back to the gateway for that interface.

Tip: WAN-type interfaces are interfaces which have a gateway set on their **Interfaces** menu entry configuration, or interfaces which have a dynamic gateway such as DHCP, PPPoE, or assigned OpenVPN, GIF, or GRE interfaces.

In situations such as bridging, this behavior is undesirable if the WAN gateway IP address is different from the gateway IP address of the hosts behind the bridged interface. Disabling `reply-to` will allow clients to communicate with the proper gateway.

Another case that has issues with `reply-to` involves static routing to other systems in a larger WAN subnet. Disabling `reply-to` in this case would help ensure that replies return to the proper router instead of being routed back to the gateway.

This behavior can also be disabled on individual firewall rules rather than globally using this option.

Disable Negate rules

In a Multi-WAN configuration traffic for directly connected networks and VPN networks typically must still flow properly when using policy routing. The firewall will insert rules to pass this local and VPN traffic without a gateway specified, to maintain connectivity. In some cases these negation rules can over-match traffic and allow more than intended.

Tip: The best practice is to create manual negation rules at the top of internal interfaces such as LAN. These rules should pass to local and VPN destinations without a gateway set on the rule, to honor the system routing table. These rules do not have to be at the top of the interface rules, but they must be above rules that have a gateway set.

Allow APIPA

Automatic Private IP Addressing (APIPA), or IPv4 Link-Local addressing, uses a special subnet of `169.254.0.0/16`. This traffic is for local links only (same L2), it must not be routed or traverse a firewall. As such, inbound traffic from these addresses is automatically blocked by internal firewall rules by default.

When **Allow APIPA traffic** is checked, the default block rules are removed, and user firewall rules can control the traffic.

There are some use cases which utilize these addresses for private communication on an interface, such as AWS VPC BGP, and in those cases, the option can be enabled along with carefully crafted manual firewall rules.

Warning: When this option is enabled, take care to never allow APIPA traffic to match policy routing rules. If APIPA traffic matches policy routing rules, behavior is unpredictable. There have been reports of such errors leading to packet loops and unexpectedly high resource usage. See Redmine Issue #2073 for more.

Aliases Hostnames Resolve Interval

This option controls how often hostnames in aliases are resolved and updated by the `filterdns` daemon. By default, this is `300` seconds (5 minutes). In configurations with few hostnames or a fast/low-load DNS server, decrease this value to pick up changes faster.

Check Certificate of Alias URLs

When **Verify HTTPS certificates when downloading alias URLs** is set, the firewall will require a valid HTTPS certificate for web servers used in URL table aliases. This behavior is more secure, but if the web server is private and uses a self-signed certificate, it can be more convenient to ignore the validity of the certificate and allow the data to be downloaded.

Warning: The best practice is to always use a server certificate with a valid chain of trust for this type of role, rather than weakening security by allowing a self-signed certificate.

Bogon Networks

The **Update Frequency** drop-down for **Bogon Networks** controls how often these lists are updated. Further information on bogon networks may be found in *Block Bogon Networks*.

Network Address Translation

NAT Reflection mode for Port Forwards

The **NAT Reflection mode for port forwards** option controls how NAT reflection is handled by the firewall. These NAT redirect rules allow clients to access port forwards using the public IP addresses on the firewall from within local internal networks.

See also:

Refer to *NAT Reflection* for a discussion on the merits of NAT Reflection when compared to other techniques such as Split DNS.

There are three possible modes for NAT Reflection:

Disabled
: The default value. When disabled, port forwards are only accessible from WAN and not from inside local networks.

Pure NAT
: This mode uses a set of NAT rules to direct packets to the target of the port forward. It has better scalability, but it must be possible to accurately determine the interface and gateway IP address used for communication with the target at the time the rules are loaded. There are no inherent limits to the number of ports other than the limits of the protocols. All protocols available for port forwards are compatible.

  When this option is enabled, **Automatic Outbound NAT for Reflection** must also be enabled if the clients and servers are in the same local network.

NAT + Proxy
: *NAT + proxy* mode uses a helper program to send packets to the target of the port forward. The connection is received by the reflection daemon, and it acts as a proxy, creating a new connection to the local server. This behavior puts a larger burden on the firewall, but is useful in setups where the interface and/or gateway IP address used for communication with the target cannot be accurately determined at the time the rules are loaded. *NAT + Proxy* reflection rules are not created for ranges larger than 500 ports and will not be used for more than 1000 ports total between all port forwards. This feature is only compatible with TCP port forwards.

Individual NAT rules have the option to override the global NAT reflection configuration, so they may have NAT reflection forced on or off on a case-by-case basis.

Reflection Timeout

The **Reflection Timeout** setting forces a timeout on connections made when performing NAT reflection for port forwards in *NAT + Proxy* mode. If connections are staying open and consuming resources, this option can mitigate that issue.

NAT Reflection for 1:1 NAT

When checked, this option adds additional reflection rules which enable access to 1:1 mappings of external IP addresses from internal networks. This gives the same functionality that already exists for port forwards, but for 1:1 NAT. There are complex routing scenarios that may render this option ineffective.

This option only affects the *inbound* path for 1:1 NAT, not outbound. The underlying rule style is similar to the *Pure NAT* mode for port forwards. As with port forwards, there are per-entry options to override this behavior.

Automatic Outbound NAT for Reflection

When checked, this option automatically creates outbound NAT rules which assist reflection rules that direct traffic back out to the same subnet from which it originated. These additional rules allow Pure NAT and 1:1 NAT Reflection to function fully when the clients and servers are in the same subnet. In most cases, this box must be checked for NAT Reflection to work.

Note: This behavior is necessary because when clients and servers are in the same subnet, the traffic source must be changed so that the connection appears to originate from the firewall. Otherwise, the return traffic will bypass the firewall and the connection will not succeed.

TFTP Proxy

The built-in TFTP proxy will proxy connections to TFTP servers outside the firewall, so that client connections may be made to remote TFTP servers. Ctrl-click or shift-click to select multiple entries from the list. If no interfaces are chosen, the TFTP proxy service is deactivated.

State Timeouts

The **State Timeout** section allows fine-tuning of the state timeouts for various protocols. These are typically handled automatically by the firewall and the values are dictated by the *Firewall Optimization Options* options. In rare cases, these timeouts may need adjusted up or down to account for irregularities in device behavior or site-specific needs.

All the values are expressed in *seconds*, and control how long a connection in that state will be retained in the state table.

See also:

Descriptions in the following options reference firewall state conditions as described in *Interpreting States*.

TCP First
: The first packet of a TCP connection.

TCP Opening
: The state before the destination host has replied (e.g. `SYN_SENT:CLOSED`).

TCP Established
: An established TCP connection where the three-way handshake has been completed.

TCP Closing
: One side has sent a TCP FIN packet.

TCP FIN Wait
: Both sides have exchanged FIN packets and the connection is shutting down. Some servers may continue to send packets during this time.

TCP Closed
: One side has sent a connection reset (TCP RST) packet.

TCP Tsdiff
: The allowed TCP timestamp difference.

UDP First
: The first UDP packet of a connection has been received.

UDP Single
: The source host has sent a single packet, but the destination has not replied (e.g. `SINGLE:NO_TRAFFIC`).

UDP Multiple
: Both sides have sent packets.

ICMP First
: An ICMP packet has been received.

ICMP Error
: An ICMP error was received in response to an ICMP packet.

Other First, Other Single, Other Multiple
: The same as **UDP**, but for other protocols.

8.6.3 Networking

DHCP Options

Server Backend

Selects the DHCP backend to use for allocating DHCP and DHCPv6 addresses to clients.

Kea DHCP
Kea is a modern DHCP server under active development, and the current best practice choice.

ISC DHCP (Deprecated)
This is the legacy DHCP server which has been the default in the past. ISC has stopped deprecated this server, and they will no longer release updates for bug fixes or security issues.

While it is deprecated, the pfSense software GUI still offers it as a choice for the time being. It will eventually be removed from pfSense software.

Ignore Deprecation Warning
As ISC DHCP software is deprecated and potentially insecure, the GUI displays a warning when this backend is active. Checking this box suppresses that warning message.

RADVD Debug

When checked, all log levels for RADVD will be logged to **System > Routing** logs. Otherwise, only log levels of `LOG_ERR` and higher will be logged.

DHCP6 Debug

When checked, the DHCPv6 client is started in debug mode to aid in troubleshooting.

Do not Allow PD/Address Release

By default, `dhcp6c` sends a `RELEASE` message to the ISP when it exits. Some ISPs then release the allocated address and/or prefix. This option prevents `dhcp6c` from sending that signal.

DHCP6 DUID

This option controls the DHCPv6 Unique Identifier (DUID) used by the firewall when requesting an IPv6 address. The firewall generates a DUID automatically, but in some cases, an administrator may want to use a different DUID. For example, if the operating system was reinstalled and the firewall should use the same DUID it had in the past, or if an upstream network administrator requires a specific DUID.

Note: Most users do not need to change this to any specific value, the default behavior is fine for nearly all environments. When in doubt, leave it alone unless directed to change it by an upstream network provider.

There are several possible DUID formats that this option can accept, chosen by the drop-down menu. When a format is chosen, the GUI displays a different set of input boxes specific to the selected format. The exact format depends upon the needs of the network administrator (e.g. ISP, datacenter, etc.) and they would provide the format and values.

The available DUID formats are:

Raw DUID
DUID represented exactly as observed in a DUID file or in logs. Entered as:

Raw DUID
A single text area in which the DUID can be entered.

This option also includes a  **Copy DUID** button which copies the DUID from the placeholder (automatically generated by the firewall) into the text box so that the existing DUID can easily be placed into the configuration.

DUID-LLT
DUID format with Link-Layer Address Plus Time. Entered as:

Time
Time (in seconds) since January 1st, 2000 UTC

Link-Layer Address
The link-layer address (MAC) of an interface on the firewall in the format `xx:xx:xx:xx:xx:xx`.

DUID-EN
DUID assigned by a vendor based on Enterprise Number. Entered as:

Enterprise Number
IANA Private Enterprise Number of the vendor.

Identifier
Variable length identifier in the format `xx:xx:xx:xx`. The length depends upon the vendor.

DUID-LL
DUID based on only Link-Layer Address. Entered as:

Link-Layer Address
The link-layer address (MAC) of an interface on the firewall in the format `xx:xx:xx:xx:xx:xx`.

DUID-UUID
DUID based on the host Universally Unique Identifier (UUID). Entered as:

DUID-UUID
The UUID for this host in nnnnnnnn-nnnn-nnnn-nnnn-nnnnnnnnnnnn format

IPv6 Options

Allow IPv6

The **Allow IPv6** option controls a set of block rules which prevent IPv6 traffic from being handled by the firewall. This option is enabled by default on new configurations.

Note: This option **does not disable IPv6** functions or prevent IPv6 from being configured, it only controls traffic flow using internal firewall rules.

When the option is checked, the firewall will allow IPv6 traffic when permitted by user rules and/or automatic rules, depending on the firewall configuration.

When the option is unchecked, the firewall will block all IPv6 traffic.

IPv6 over IPv4 Tunneling

The **Enable IPv6 over IPv4 Tunneling** option enables forwarding for IP protocol 41/RFC 2893 to an IPv4 address specified in the **IPv4 address of Tunnel Peer** field.

When configured, this forwards all incoming protocol 41/IPv6 traffic to a host behind this firewall instead of handling it locally.

Tip: Enabling this option does not add firewall rules to allow the protocol 41 traffic. A rule must exist on the WAN interface to allow the traffic to pass through to the local receiving host.

Prefer IPv4 over IPv6

When set, this option causes the firewall *itself* to prefer sending traffic to IPv4 hosts instead of IPv6 hosts when a DNS query returns results for both.

In rare cases when the firewall has partially configured, but not fully routed, IPv6 this can allow the firewall to continue reaching Internet hosts over IPv4.

Note: This option controls the behavior of the firewall itself, such as when polling for updates, package installations, downloading rules, and fetching other data. It cannot influence the behavior of clients behind the firewall.

IPv6 DNS Entry

This option controls whether the firewall creates local DNS entries for the firewall itself with IPv6 addresses, when available.

By default, (unchecked), the firewall automatically adds DNS entries for itself using its local IPv4 and IPv6 interface addresses. In some cases, such as with dynamic IPv6 addresses like tracked interfaces, the IPv6 address may disappear or change and clients may attempt to use an outdated address until their cached DNS response expires.

When the option is checked, the firewall only adds DNS entries for its IPv4 addresses.

Network Interfaces

Hardware Checksum Offloading

When checked, this option disables hardware checksum offloading on the network cards. Checksum offloading is usually beneficial as it allows the checksum to be calculated (outgoing) or verified (incoming) in hardware at a much faster rate than it could be handled in software.

Note: When checksum offloading is enabled, a packet capture will see empty (all zero) or flag incorrect packet checksums. These are normal when checksum handling is happening in hardware.

Checksum offloading is broken in some hardware, particularly Realtek cards and virtualized/emulated cards such as those on Xen/KVM. Typical symptoms of broken checksum offloading include corrupted packets and poor throughput performance.

Tip: In virtualization cases such as Xen/KVM it may be necessary to disable checksum offloading on the host as well as the VM. If performance is still poor or has errors on these types of VMs, switch the type of NIC if possible.

Hardware TCP Segmentation Offloading

Checking this option will disable hardware TCP segmentation offloading (TSO, TSO4, TSO6). TSO causes the NIC to handle splitting up packets into MTU-sized chunks rather than handling that at the OS level. This can be faster for servers and appliances as it allows the OS to offload that task to dedicated hardware, but when acting as a firewall or router this behavior is **highly undesirable** as it actually increases the load as this task has already been performed elsewhere on the network, thus breaking the end-to-end principle by modifying packets that did not originate on this host.

Warning: This option is not desirable for routers and firewalls, but can benefit workstations and appliances. It is disabled by default, and should remain disabled unless the firewall is acting primarily or solely in an appliance/endpoint role.

Do not uncheck this option unless directed to do so by a Netgate TAC representative. This offloading is broken in some hardware drivers, and can negatively impact performance on affected network cards and roles.

Hardware Large Receive Offloading

Checking this option will disable hardware large receive offloading (LRO). LRO is similar to TSO, but for the incoming path rather than outgoing. It allows the NIC to receive numerous smaller packets before passing them up to the operating system as a larger chunk. This can be faster for servers and appliances as it offloads what would normally be a processing-heavy task to the network card. When acting as a firewall or router this is **highly undesirable** as it delays the reception and forwarding of packets that are not destined for this host, and they will have to be split back up again on the outbound path, increasing the workload significantly and breaking the end-to-end principle.

Warning: This option is not desirable for routers and firewalls, but can benefit workstations and appliances. It is disabled by default, and should remain disabled unless the firewall is acting primarily or solely in an appliance/endpoint role.

Do not uncheck this option unless directed to do so by a Netgate TAC representative. This offloading is broken in some hardware drivers, and can negatively impact performance on affected network cards and roles.

hn ALTQ Support

Checking this option will configure the `hn(4)` driver for network interfaces in Hyper-V to work with ALTQ traffic shaping.

For ALTQ to work on `hn(4)` interfaces, the operating system must disable the multi-queue API which may reduce the system capability to handle traffic. The administrator must decide if this reduction in performance is worth the benefit of traffic shaping.

The firewall must be rebooted for this setting to take effect.

Suppress ARP messages

The firewall makes a log entry in the main system log when an IP address appears to switch to a different MAC address. This log entry notes that the device has moved addresses, and records the IP address and the old and new MAC addresses.

This event can be completely benign behavior (e.g. NIC teaming on a Microsoft server, a device being replaced) or a legitimate client problem (e.g. IP conflict), and it could show up constantly or rarely if ever. It all depends on the network environment.

The best practice is to allow these ARP messages to be printed to log since there is a chance it will report a problem worth the attention of a network administrator. However, if the network environment contains systems which generate these messages while operating normally, suppressing the errors can make the system log more useful as it will not be cluttered with unneeded log messages.

Reset All States

When set, if an interface IP address changes, the firewall will reset the entire state table instead of only clearing states for the old interface IP address.

This behavior is potentially disruptive, and is off by default. In single WAN environments, this is not typically any more disruptive than the WAN address changing, since clients already have to reestablish all connections.

In most cases, this behavior is not necessary, but it can help in certain situations where WAN addresses change rapidly, and the normal behavior misses states for former IP addresses.

Use `if_pppoe` Kernel Module

When set, pfSense software uses the kernel-based `if_pppoe` backend to handle PPPoE WAN interfaces and disables the MPD backend for PPPoE WANs. The `if_pppoe` backend is very fast and efficient, but it is not compatible with some advanced features found in the older deprecated MPD backend, such as MLPPP. The `if_pppoe` backend for PPPoE interfaces conflicts with the MPD backend, so the two cannot be used at the same time.

This option is disabled by default, but will be the default in future versions of pfSense Plus software.

8.6.4 Miscellaneous

Proxy Support

If this firewall resides in a network which requires a proxy for outbound Internet access, enter the proxy options in this section so that requests from the firewall for items such as packages and updates will be sent through the proxy.

Proxy URL
: This option specifies the location of the proxy for making outside connections. It must be an IP address or a fully qualified domain name.

Proxy Port
: The port to use when connecting to the proxy URL. By default, the port is `8080` for HTTP proxy URLs, and `443` for SSL proxy URLs. The port is determined by the proxy, and may be a different value entirely (e.g. `3128`). Check with the proxy administrator to find the proper port value.

Proxy Username
: If required, this is the username that is sent for proxy authentication.

Proxy Password
If required, this is the password associated with the username set in the previous option.

Load Balancing

When pfSense® software is directed to perform load balancing, successive connections will be redirected in a round-robin manner to a gateway, balancing the load across all available paths. The options in this section alter or fine-tune that behavior.

Sticky Connections
When active, connections from the same source are sent through the same gateway, rather than using multiple gateways in a round-robin manner.

This "sticky" association will exist as long as states are in the table for connections from a given source address (e.g. the IP address of a user). Once the states for that source expire, so will the sticky association. Further connections from that source host will be redirected to the next available gateway in the group.

This behavior can help with protocols such as HTTPS and FTP, where the server may be strict about all connections coming from the same source. The downside of this behavior is that balancing is not as efficient, a heavy user could dominate a single WAN rather than having their connections spread out.

Source Tracking Timeout
Controls how long the sticky association will be maintained for a host after the all the states from that host expire. The value is specified in seconds.

By default, this value is not set, so the association is removed as soon as the states expire. If sticky connections appear to work initially but seem to stop partway through sessions, increase this value to hold an association longer. Web browsers often hold open connections for a while as users are on a site, but if there is a lot of idle time, connections may be closed and states may expire.

Power Savings

There are currently two different types of power configuration available, Speed Shift and SpeedStep (PowerD). Using Speed Shift is the best practice, but it is only compatible with certain models of recent hardware.

Tip: The GUI displays options for Speed Shift only when the proper hardware is present. SpeedStep/PowerD settings are always available.

With either method of power management, if processes need the power, the CPU speed will be increased as needed and lowered when demand decreases. This lowers the amount of heat a CPU generates, and can also lower power consumption.

Warning: Speed Shift takes precedence over SpeedStep and only one of the two should be enabled at a time. If Speed Shift is enabled, SpeedStep is automatically disabled.

Speed Shift

Speed Shift configures hardware-controlled performance states which enable the CPU itself to control changes in clock speed. This allows the system to respond much faster to changes in load. This method also does not rely on the BIOS passing useful values for power management since it is handled entirely at the CPU level.

Note: Speed Shift responds so fast to load that even the act of checking the current CPU frequency status raises the clock speed temporarily, making the results appear higher!

Enable Speed Shift
This checkbox enables or disables Speed Shift.

The current status of Speed Shift is printed below the help text for this option.

> **Warning:** Changing this setting requires a reboot.

Control Level
Chooses between per-core or per-package frequency control. Core-level control is the best practice in most cases, especially for hardware with only a single physical CPU.

The active control level is printed below the help text for this option.

Core Level Control
Each CPU core can run at a different frequency.

Package Level Control
All cores on the same physical CPU package are locked to the same speed, but each package may run at a different speed.

> **Warning:** Changing this setting requires a reboot.

Power Preference
This option slider influences the bias of the hardware performance state toward performance (left, lower values) or energy efficiency (right, higher values).

A numerical representation of the preference level (from 0-100) is printed below the slider.

The default value is `80` to help keep heat and power usage down. For increased responsiveness to performance demands, move the slider farther to the left as needed.

PowerD

PowerD manages CPU frequency changes via SpeedStep. This daemon monitors the system, and it lowers or raises the CPU frequency based on activity levels.

Note: The behavior of this option depends greatly on the hardware in use. In some cases, the CPU frequency may lower but have no measurable effect on power consumption and/or heat, where others will cool down and use considerably less power. It is considered safe to run, but is left off by default unless pfSense detects compatible hardware.

Enable PowerD
Checking this option enables the `powerd` daemon.

Modes

The mode for `powerd` may also be selected for three system states:

AC Power
Normal operation connected to AC power.

Battery Power
Mode to use when the firewall is running on battery. Battery power detection capabilities vary by hardware.

Unknown Power
Mode used when `powerd` cannot determine the power source.

Four modes choices exist for each of these states:

Maximum
Keeps the performance as high as possible at all times.

Minimum
Keeps performance at its lowest, to reduce power consumption.

Adaptive
Tries to balance savings by decreasing performance when the system is idle and increasing when busy.

Hiadaptive
Similar to adaptive but tuned to keep performance high at the cost of increased power consumption. It raises the CPU frequency faster and drops it slower. This is the default mode.

Note: Some hardware requires `powerd` running to operate at its maximum attainable CPU frequency. If the firewall device does not have `powerd` enabled but always runs at what appears to be a low CPU frequency, enable `powerd` and set it to *Maximum* for at least the **AC Power** state.

Watchdog

Certain firewall hardware includes a **Watchdog** feature which can reset the hardware when the watchdog daemon can no longer interface with the hardware after a specified timeout. This can increase reliability by resetting a unit when a hard lock is encountered that might otherwise require manual intervention.

The downside to any hardware watchdog is that any sufficiently busy system may be indistinguishable from one that has suffered a hard lock.

Enable Watchdog
When checked, the `watchdogd` daemon is run which attempts to latch onto a compatible hardware watchdog device.

Watchdog Timeout
The time, in seconds, after which the device will be reset if it fails to respond to a watchdog request. If a firewall regularly has a high load and triggers the watchdog accidentally, increase the timeout.

Cryptographic & Thermal Hardware

On systems with configurable cryptographic and/or thermal hardware, this section provides options to control the hardware. On systems without configurable cryptographic modules, this section only displays options for thermal hardware.

IPsec-MB

The checkbox for **IPsec-MB** enables IPsec Multi-Buffer (IPsec-MB, IIMB) Cryptographic Acceleration.

IPsec-MB assists VPN performance by replacing the cryptographic functions provided by the kernel for AES-CBC, AES-GCM, and ChaCha20-Poly1305 with accelerated functions that utilize the optimal CPU SIMD instruction set.

This benefits any VPN utilizing the accelerated algorithms in the kernel which includes IPsec, OpenVPN DCO, and WireGuard.

IPsec-MB is faster than AES-NI and can even meet or exceed the performance of dedicated acceleration hardware such as QAT on current versions of pfSense software.

Note: If IPsec-MB and QAT are both enabled, IPsec-MB will take over handling of AES-GCM acceleration. Depending on the hardware, QAT may accelerate AES-GCM faster than IPsec-MB, but IPsec-MB can accelerate ChaCha20-Poly1305 which is not compatible with QAT. Depending on the required performance of each algorithm it may be better to only enable QAT, or to enable both, but it depends on the environment and workload.

See also:

- *Cryptographic Accelerators*
- *Tuning IPsec-MB*

Cryptographic Hardware

There are a few options available for accelerating cryptographic operations via hardware. Some are built into the kernel, and others are loadable modules.

Note: Some modules and hardware require pfSense® Plus software.

See also:

Cryptographic Accelerators

The following choices are available, depending on hardware:

Intel QuickAssist (QAT)
: Loads the Intel QuickAssist (QAT) driver for compatible hardware. QAT accelerates many types of AES and SHA operations and is ideal for use with IPsec and OpenVPN DCO.

BSD Crypto Device
: Loads the BSD Crypto device module (`cryptodev`) so it can be used by other available acceleration devices. Most accelerator drivers hook into the `crypto(9)` framework in FreeBSD, so many aspects of the system will automatically use acceleration for compatible ciphers when this module is loaded.

AES-NI CPU-based Acceleration
: Loads the **AES-NI** (Advanced Encryption Standard, New Instructions) kernel module. Notably, the `aesni` module will accelerate operations for AES-GCM, available in IPsec.

AES-NI instructions are built into many recent Intel and some AMD CPUs. Check with the OEM for specific CPU or SoC compatibility.

Speeds with AES-NI vary depending on the underlying software. IPsec speed will be greatly increased with AES-NI loaded provided that AES-GCM is used and properly configured.

AES-NI and BSD Crypto Device
Loads both the AES-NI and BSD Crypto Device modules together, which is the optimal configuration in most cases. Choose this unless a specific environment or configuration is found to work better without it.

SafeXcel and BSD Crypto Device [Plus Only]
Loads both the `safexcel` and the BSD Crypto Device modules. SafeXcel acceleration hardware is found on some ARM systems sold by Netgate, such as the SG-3100.

There are other compatible cryptographic devices with drivers built into the kernel. One example is the driver for the Marvell Cryptographic Engine and Security Accelerator (CESA) chipset, which is found on some ARM systems sold by Netgate, such as the SG-1100 and SG-2100.

In most cases, if the firewall detects a compatible accelerator chip, it will show it in the **System Information** widget on the dashboard or in the system log at boot time.

Note: Certain special cases also exist where software can detect and use acceleration hardware directly, even without drivers loaded. One example is OpenSSL, which can use AES-NI CPU instructions itself. Thus, even without the driver loaded, software which utilizes encryption through OpenSSL can still take advantage of AES-NI acceleration.

Thermal Sensors

The firewall can read temperature data from a few sources to display on the dashboard. If the firewall has a compatible CPU, selecting a thermal sensor will load the appropriate driver to read its temperature.

Note: Temperature data can be displayed by the **Thermal Sensors** *dashboard widget* or via `sysctl`.

The following sensor types are available:

None/ACPI
The firewall will attempt to read the temperature from an ACPI-compliant motherboard sensor if one is present, otherwise no sensor readings are available.

Intel Core
Loads the `coretemp` module which can read thermal data from Intel core-series CPUs and other modern Intel CPUs using their on-die sensors, including Atom-based processors.

AMD K8, K10, and K11
Loads the `amdtemp` module which can read thermal data from modern AMD CPUs using their on-die sensors.

If the firewall does not have a compatible thermal sensor chip, this option will have no effect. To unload the selected module, set this option to *None/ACPI* and then reboot.

Note: The `coretemp` and `amdtemp` modules report thermal data directly from the CPU core. This may or may not be indicative of the temperature elsewhere in the system. Case temperatures can vary greatly from temperatures on the CPU die.

Kernel Page Table Isolation (PTI)

Kernel PTI is a method for working around CPU vulnerabilities such as Meltdown. By exploiting that vulnerability without Kernel PTI, kernel memory could be accessed by unprivileged users on affected CPUs.

Note: While more secure, this protection can incur a performance penalty. If untrusted users do not have access to run arbitrary code on the firewall, it can be disabled without significant security risk.

Kernel PTI is active by default only on CPUs affected by the vulnerability.

This option forces the workaround off, and requires a reboot to change.

If a vulnerable CPU is not detected, PTI is disabled by default and this option will have no effect.

The current state of Kernel PTI is printed below the option.

Microarchitectural Data Sampling (MDS) Mitigation

Microarchitectural Data Sampling (MDS) mitigation is a method for working around weaknesses in Intel CPUs with Hyper-Threading Technology (HTT). By exploiting MDS without mitigation in place, kernel memory could be accessed by unprivileged users on affected CPUs.

Note: While more secure, this protection can incur a performance penalty. If untrusted users do not have access to run arbitrary code on the firewall, it can be disabled without significant security risk.

This option controls which method of MDS mitigation is used, if any. Changing the option requires a reboot to activate. The following modes are available:

Default
: The default operating system behavior. As of this writing, the default behavior is to disable MDS mitigation.

Mitigation Disabled
: Forcefully disable MDS mitigation.

VERW instruction (microcode) mitigation enabled
: Use VERW instruction mitigation, implemented in CPU microcode, to mitigate MDS. This is the fastest and most optimal way to mitigate MDS, but it requires support in the CPU microcode for this instruction.

Software sequence mitigation enabled
: Mitigates MDS by using software sequences, which is much slower, but safer.

Automatic VERW or Software selection
: When set to Automatic, the operating system will attempt to use VERW instructions if they are available and software in all other cases.

The current state of MDS mitigation is printed below the option.

Schedules

The **Do not kill connections when schedule expires** option controls whether states are cleared when a scheduled rule transitions into a state that would block traffic. If unchecked, connections are terminated when the schedule time has expired. If checked, connections are left alone and will not be automatically closed by the firewall.

Gateway Monitoring

State Killing on Gateway Recovery

States from the firewall itself

Connections from the firewall itself can fail over to other gateways by setting a failover gateway group as the system's default gateway. By killing states on lower-priority gateways after a higher-priority gateway recovers, these connections can re-establish on the preferred gateway. This option overrides the failover gateway group's state-killing behavior by affecting all states, not only those created by policy routing rules.

Don't kill states from the firewall itself (default)
: States from the firewall itself are unaffected. The configured failover gateway group determines the state-killing behavior for states created by policy routing rules.

Kill all states for lower-priority gateways
: All states on lower-priority gateways are killed when a higher-priority gateway returns to an online state.

Only kill states with the same Address Family as the gateway group
: States of the same Address Family as the gateway group are killed for lower-priority gateways.

States from policy routing rules

The state-killing behavior on gateway recovery handles policy routing states separately. This allows for applying different behaviors to different types of traffic. For example, one gateway group can handle general internet traffic by killing states on both gateway failure and recovery. Separately, another group handling VoIP traffic would only kill states on gateway failure to avoid interrupting active calls.

Don't kill policy routing states for lower-priority gateways
: Controls the default state-killing behavior for states created by policy routing rules using a failover gateway group. This behavior may also be controlled per gateway group. If unchecked (default), policy routing states on lower-priority gateways are killed when a higher-priority gateway recovers.

State Killing on Gateway Failure

When using Multi-WAN, clearing states on failed WANs can help redirect traffic for long-lived connections such as VoIP phone/trunk registrations to another WAN. However, clearing states can also disrupt ongoing connections if a lesser-used gateway is unstable or there is a gateway which is down long term but is not disabled, which would still states when it fails or is down during a filter reload.

There are several choices for this behavior, including:

Do not kill states on gateway failure (Default)
: The monitoring process will not flush states when a gateway is in a down state during a filter reload. This is the default behavior and is the least disruptive, though clients may have to wait for connections to timeout after a WAN failure.

Kill states for all gateways which are down
Selectively kill states using gateways that fail or are down during a filter reload, so long as those states were created by policy routing rules.

This function can only kill states which contain gateway information populated by policy routing rules (e.g. gateways or gateway groups on firewall rules, or even `reply-to`.). It cannot kill states created by default gateway switching as in that case the gateway in the state is `0.0.0.0/::` and not a specific gateway.

Flush all states on gateway failure
Clears **all** states for existing connections when **any** gateway fails or is in a down state during a filter reload.

> **Warning:** When this is triggered the firewall clears the **entire** state table if any gateway is down, which can be highly disruptive.

More information on how this impacts Multi-WAN can be found in *State Killing/Forced Switch*.

Skip Rules When Gateway is Down

By default, when a rule has a specific gateway set and this gateway is down, the gateway is omitted from the rule and traffic is sent via the default gateway.

The **Do not create rules when gateway is down** option overrides that behavior and the entire rule is omitted from the ruleset when the gateway is down. Instead of flowing via the default gateway, the traffic will match a different rule instead. This is useful if traffic must only ever use one specific WAN and never flow over any other WAN.

Tip: When utilizing this option, create a reject or block rule underneath the policy routing rule with the same matching criteria. This will prevent the traffic from potentially matching other rules below it in the ruleset and taking an unintended path.

Static Routes

By default, the firewall adds static routes for gateway monitor IP addresses to ensure traffic to the monitor IP address leaves via the correct interface. Enabling this checkbox overrides that behavior. When this option is set, the user will have to ensure the traffic exits the correct interface in some other way.

PHP Settings

Memory Limit
Configures the amount of memory an individual PHP process may consume. Default value varies based on available system memory. The GUI prints the default memory limit for the device in the option description.

Certain environments may require increasing this amount so that the GUI can better handle large volumes of items such as states, routes, etc.

> **Warning:** Do not adjust this value without careful consideration. A PHP memory error may not be related to load, but may be due to some other problem (e.g. a bug) or user error (user-created run process producing infinite output). In those cases, allowing PHP to use more memory will only further strain the device and not correct the problem.

RAM Disk Settings

The `/tmp` and `/var` directories are used for writing files and holding data that is temporary and/or volatile. Using a RAM disk can reduce the amount of writing that happens on disks in the firewall. Modern SSDs do not have disk write concerns as older drives once did, but it can still be a concern when running from lower quality flash storage such as USB thumb drives.

This behavior has the benefit of keeping most of the writes off of the disk in the base system, but packages may yet write frequently to the drive. It also requires additional handling to ensure data such as RRD graphs and DHCP leases are retained across reboots. Data for both is saved during a proper shutdown or reboot, and also periodically if configured.

Use RAM Disks
When checked, a memory disk is created at boot time for `/tmp` and `/var/` and the associated structure is initialized. When this setting is toggled, a reboot is required and forced on save.

> **Warning:** The size of RAM disks is limited by the amount of available kernel memory. The actual limit is calculated and printed in the GUI underneath the size options.

/tmp RAM Disk Size
The size of the `/tmp` RAM disk, in MiB. The default value is `40`, but should be set higher if there is available RAM and kernel memory.

/var RAM Disk Size
The size of the `/var` RAM disk, in MiB. The default value is `60`, but should be set much higher, especially if packages will be used. `512-1024` is a better starting point, depending on the available firewall RAM and kernel memory.

Periodic RAM Disk Data Backups
These options control how frequently data in RAM disks is backed up. If the firewall is rebooted unexpectedly, the last backup is restored when the firewall boots. The lower the value, the less data that will be lost in such an event, but more frequent backups write more to the disk.

RRD Data
The time, in hours, between periodic backups of RRD files containing graph data.

DHCP Leases
The time, in hours, between periodic backups of the DHCP lease databases.

Log Directory
The time, in hours, between periodic backups of the system log directory.

Captive Portal Data
The time, in hours, between periodic backups of Captive Portal databases.

> **Warning:** Aside from the points mentioned above, there are several items to be cautious about when choosing whether to use the RAM disk option. Used improperly, this option can lead to data loss or other unexpected failures.
>
> Utilize remote syslog to send the logs to another device on the network rather than risking losing data from unexpected outages.

Packages may not properly account for the use of RAM disks, and may not function properly at boot time or in other ways. Test each package, including whether it works immediately after a reboot.

These are RAM disks, so the amount of RAM available to other programs will be reduced by the amount of space used by the RAM disks. For example if the firewall has 2GB of RAM, and has 512MB for `/var` and 512MB for `/tmp`, then only 1GB of RAM will be available to the OS for general use.

Special care must be taken when choosing a RAM disk size, which is discussed in the following section.

RAM Disk Sizes

Setting a size too small for `/tmp` and `/var` can backfire, especially when it comes to packages. The suggested sizes on the page are an **absolute minimum** and often much larger sizes are required. The most common failure is that when a package is installed, and parts of the package touch places in both `/tmp` and `/var`, and it can ultimately fill up the RAM disk and cause other data to be lost.

For `/tmp`, a minimum of `40` MiB is required. For `/var` a minimum of `60` MiB is required. To determine the proper size, check the current usage of the `/tmp` and `/var` directories before making a switch. Check the usage several times over the course of a few days so it is not caught at a low point. Watching the usage during a package installation adds another useful data point.

Hard Disk Standby

The **Hard disk standby time** option activates power management for disk drives in the firewall. The drop-down field sets the number of minutes that the disk can be idle before going into standby mode.

Using standby mode is not necessary for SSD or flash media. For traditional spinning platter hard disks, it may result in power savings and can potentially lengthen the disk lifetime by saving wear, at a cost of slower disk access when resuming from an idle state. Actual results entirely depend on the hardware involved.

The default behavior is *Always On* which prevents the disk from entering standby mode.

Installation Feedback

When this option is set, the firewall will not send its Netgate Device ID when making requests to Netgate servers.

8.6.5 System Tunables

The **System Tunables** tab under **System > Advanced** provides a means to set runtime FreeBSD system tunables, also known as sysctl object identifiers (OIDs).

Tip: In most cases, the best practice is to leave these tunables at their default values.

Firewall administrators familiar with FreeBSD, or users acting under the direction of a developer or Netgate TAC representative, may want to adjust or add values on this page so that they will be set as the system starts.

Kernel State and Tunables

The `sysctl` facility on FreeBSD allows managing certain aspects of the kernel state through a "Management Information Base" (MIB) style tree composed of individual object identifiers (OIDs) containing components separated by periods. These individual `sysctl` OIDs are often referred to as "tunables" but not all of them can be changed.

See also:

This is a simplified description. The sysctl manual page contains more detail.

The most common types of FreeBSD tunables are:

Runtime Tunables
: The values of runtime tunables can be changed at any time while the system is running.

Loader Tunables
: The values of loader tunables can only be changed at boot in the loader, and they are read only afterward when the system is running.

Read Only Tunables
: The values of read only tunables can never be changed manually, they are typically for reference or statistical purposes.

Tunable OIDs and Values

There are many OIDs available from `sysctl`. The full list of OIDs and their possible values is outside the scope of this documentation, but for those interested in digging a little deeper, The sysctl manual page from FreeBSD contains detailed instructions and information.

To see the current values of all visible OIDs, run:

```
# sysctl -a
```

Managing Runtime Tunables

Persistent values for runtime tunables can be managed from within the GUI.

To create a new tunable:

- Navigate to **System > Advanced**, **System Tunables** tab
- Click  **New** at the top right of the list

To edit an existing tunable:

- Navigate to **System > Advanced**, **System Tunables** tab
- Locate the entry to edit
- Click  on its row

---

Note: The items on this page are different from **Loader Tunables**. For details on loader tunables, see *Managing Loader Tunables*.

When editing or creating a tunable, the following fields are available:

Tunable
The `sysctl` OID to set.

Value
The value to which the **Tunable** will be set.

Note: Some values have formatting requirements. Due to the vast number of sysctl OIDs, the GUI does not validate that the given **Value** will work for the chosen **Tunable**.

Description
An optional description for reference.

Click **Save** when the form is complete.

Managing Loader Tunables

Loader tunable values must be set before the kernel boots and user-defined loader tunables belong in `/boot/loader.conf.local`, which can be created or edited in several ways.

To determine loader tuneable values at boot the operating system first reads `/boot/defaults/loader.conf`, then `/boot/loader.conf`, and finally `/boot/loader.conf.local`. After the kernel boots, loader tunable values become read only.

These files each have a distinct purpose:

`/boot/defaults/loader.conf`
This file contains default values from FreeBSD and **must not be changed** as it will be rewritten during any upgrade.

`/boot/loader.conf`
This file contains loader values managed by pfSense software internally and **must not be changed**. It is rewritten each boot and when certain options are changed, and any manual modifications are discarded. Values in this file can override the operating system defaults.

`/boot/loader.conf.local`
Administrators can use this file to define custom loader tunable values. Since it is read last, it can override values from the OS default values as well as values set by pfSense internally.

This file does not exist by default, but can be created at any time.

This file is not backed up in `config.xml`, make a separate manual backup of its contents.

Note: Loader tunable values can also be defined for a single boot by setting them at the loader prompt from the boot menu.

Loader tunables are not currently manageable in the GUI in an integrated way, they must be manually managed by creating or editing the `/boot/loader.conf.local` file.

Users can create and edit that file in a variety of ways in the GUI or in the shell.

GUI File Editor

The file editor in the GUI can make changes to this file:

- Navigate to **Diagnostics > Edit File**
- Enter `/boot/loader.conf.local` in the **Path to file to be edited** box
- Click **Load** to load the existing content in the file if any exists

  If the file does not exist, the editor will print an error. This error can be ignored.
- Enter the loader tunable OIDs and values in the file, one per line.

  Comments can also be added by starting a line with #

  Example:

```
# Disable flow control on all ix interfaces
hw.ix.flow_control="0"
```

- Click **Save**

Reboot the firewall to activate the new tunable values.

Shell Editors

Similar to the above process, any text editor available in the shell can make changes to `/boot/loader.conf.local`. Available editors include `vi` and `ee` in the base system, along with `vim` and `nano` which are available to install via `pkg`. These changes must be made as the `admin` or `root` user, or by a user given sufficient access using the `sudo` package.

8.6.6 Notifications

The firewall can notify administrators of important events and errors by displaying an alert in the menu bar, indicated by the  icon.

In addition to GUI notifications, the firewall also offers the following notification methods:

- Local via LED indicators on compatible hardware (not configurable)
- Local via Sounds using a PC speaker
- Remote via E-mail using SMTP
- Remote via Telegram notification API
- Remote via Pushover notification API
- Remote via Slack notification API

General Settings

Certificate Expiration
When set, the firewall will issue notifications as CA and certificate entries approach their expiration date so that administrators can take corrective action to renew or replace them. Notifications are also sent for expired entries.

Expiration times are checked daily, and notifications are displayed in the GUI and sent remotely.

Ignore Revoked
When set, the firewall will not send notifications for expired certificate entries which have been revoked in at least one CRL.

Certificate Expiration Threshold
The value, in days, at which CA and certificate entries are considered to be approaching their expiration date.

The default value is currently `27` days. Certificates from Let's Encrypt (*ACME package*) typically renew when they have around 30 days before they expire. The default value is long enough that it does not notify unnecessarily, but with enough time left that problems can be corrected.

Tip: If certificates are imported into the firewall from third party sources which take longer to process, increase this value sufficiently to give administrators enough notice to obtain an updated replacement certificate before the expiration date.

SMTP E-mail

E-mail notifications are delivered by a direct SMTP connection to a mail server. The server must be configured to allow relaying from the firewall or accept authenticated SMTP connections.

Disable SMTP
When checked, the firewall will not send SMTP notifications. This is useful to silence notifications while keeping SMTP settings in place for use by other purposes such as packages that utilize e-mail.

E-mail server
The hostname or IP address of the e-mail server through which the firewall will send notifications.

SMTP Port of E-mail server
The port to use when communicating with the SMTP server. The most common ports are `25` and `587`.

In many cases, `25` will not work unless it is to a local or internal mail server. Providers frequently block outbound connections to port `25`, so use `587` (the Submission port) when possible.

Connection Timeout to E-Mail Server
The length of time, in seconds, that the firewall will wait for an SMTP connection to complete.

Secure SMTP Connection
When set, the firewall will attempt a direct SSL/TLS connection when sending e-mail. The server must accept SSL/TLS connections on the configured port.

> **Warning:** This option is not compatible with ports that utilize plain text and switch to TLS after using STARTTLS (e.g. 25, 587).

Note: When this option is **not** checked, the firewall will still automatically attempt to use START-TLS to set up a secure TLS communications channel on ports 25 and 587 when authentication is configured.

Validate SSL/TLS
When set, the certificate presented by the mail server is checked for validity against the root certificate authorities trusted by the firewall. Ensuring this validity is the best practice.

In some rare cases a mail server may have a self-signed certificate or a certificate that otherwise fails validation. Unchecking this option will allow notifications to be sent to these servers using SSL/TLS. In this case, communication is still encrypted, but the identity of the server cannot be validated.

From e-mail address
The e-mail address for the `From:` header in notification messages, which specifies the source. Some SMTP servers attempt to validate this address, so the best practice is to use a real address in this field. This is commonly set to the same address as **Notification E-mail address**.

Notification E-mail address
The e-mail address(es) for the `To:` header of the message, which is the destination where the notification e-mails will be delivered by the firewall.

Note: This field allows multiple addresses if they are separated by a comma, for example: `me@example.com,otheradmin@example.com`.

Notification E-Mail Auth Username
Optional. If the mail server requires a username and password for authentication, enter the username here.

Notification E-Mail Auth Password
Optional. If the mail server requires a username and password for authentication, enter the password here and in the confirmation field.

Notification E-mail Auth Mechanism
This field specifies the authentication mechanism required by the mail server. The majority of e-mail servers work with *PLAIN* authentication, others such as MS Exchange may require *LOGIN* style authentication.

Note: In 2022 Google phased out access to SMTP Submission and other similar services using the account username and password directly. To access these services Google has deemed "less secure", the user must enable 2-Step Verification for their Google account and then create an App Password which can authenticate with these services.

Click **Save** at the bottom of the page to store the settings before proceeding.

Click **Test SMTP Settings** to generate a test notification and send it via SMTP using the previously stored settings. Save settings before clicking this button.

Sounds

Console Bell

When checked, emergency log messages, such as from a GUI login, will trigger a bell in connected consoles including serial terminals. On devices with a speaker, such messages can trigger an audible beep.

Startup/Shutdown Sound

If the firewall hardware has a PC speaker, it will play a sound when startup finishes and again when a shutdown is initiated.

Check **Disable the startup/shutdown beep** to prevent the firewall from playing these sounds.

Telegram

The notification system can send messages using the Telegram API which can relay notifications to desktops and mobile devices, among others.

Note: Using the Telegram API requires a Telegram Bot and its associated API key.

Enable Telegram
When set, the firewall will attempt to send remote notifications using the Telegram API and the settings in this section.

API Key
Required. The Telegram Bot API key the firewall will use to authenticate with the Telegram API server.

Chat ID
The destination for the notifications. This can be a chat ID number for private notifications, or a channel `@username` for public notifications.

Click  **Save** at the bottom of the page to store the settings before proceeding.

Click  **Test Telegram Settings** to generate a test notification and send it using the Telegram API with the previously stored settings. Save settings before clicking this button.

Pushover

The notification system can send messages using the Pushover API which can relay notifications to desktops and mobile devices, among others.

Note: Using the Pushover API requires a Pushover account user key and API key (Pushover Registration).

Enable Pushover
When set, the firewall will attempt to send remote notifications using the Pushover API and the settings in this section.

API Key
Required. The Pushover API Key (Pushover Registration) the firewall will use to authenticate with the Pushover API server.

User Key
Required. The User Key (Pushover Registration) of the Pushover account to which the **API Key** belongs.

Notification Sound
The notification sound that the end user device (Phone, etc.) will play when notification messages are sent by the firewall.

See also:

For a list of sounds and audio, see the Pushover API Notification Sounds Documentation.

Message Priority
The message priority for firewall notifications.

Note: For more information about the priorities and their meanings, see the Pushover API Priority Documentation.

The following priorities are available:

Normal
Default setting. May trigger sound, vibration, and notification display depending on the user settings and client platform.

Lowest
No sound or vibration, but increases the notification count on some platforms.

Low
No sound or vibration. May trigger a notification display depending on the user settings and client platform.

High
Always play sound and vibrate. Bypasses pre-set quiet hours. Notification display is highlighted in red.

Emergency
Similar to *High* priority, but the notification is repeated until acknowledged by the user.

Emergency Priority Notification Retry Interval
The amount of time, in seconds, the Pushover servers will send the same notification for *Emergency* priority notifications until the notification is acknowledged.

This parameter must have a value of at least `30` seconds between retries. Default is `60` seconds (1 minute).

Emergency Priority Notification Expiration
The duration, in seconds, for which *Emergency* priority notifications will be retried until the notification is acknowledged. Notifications will be resent at intervals determined by the value of **Emergency Priority Notification Retry Interval**.

This parameter must have a maximum value of at most `10800` seconds (3 hours). Default is `300` seconds (5 minutes).

Click  **Save** at the bottom of the page to store the settings before proceeding.

Click  **Test Pushover Settings** to generate a test notification and send it using the Pushover API with the previously stored settings. Save settings before clicking this button.

Slack

The notification system can send messages using the Slack API which can relay notifications to Slack channels.

Note: Using the Slack API requires a Slack API key (Slack API Registration).

Enable Slack
When set, the firewall will attempt to send remote notifications using the Slack API and the settings in this section.

API Key
Required. The Slack API Key (Slack API Registration) the firewall will use to authenticate with Slack servers.

Slack Channel
The name of the Slack channel to which the firewall will send notifications (e.g. `#firewall`).

Click  **Save** at the bottom of the page to store the settings before proceeding.

Click  **Test Slack Settings** to generate a test notification and send it using the Slack API with the previously stored settings. Save settings before clicking this button.

8.7 Console Menu Basics

Basic configuration and maintenance tasks can be performed from the pfSense® system console. The console is available using a keyboard and monitor, serial console, or by using SSH. Access methods vary depending on hardware. Below is an example of what the console menu will look like, but it may vary slightly depending on the version and platform:

```
WAN (wan)       -> vmx0       -> v4/DHCP4: 198.51.100.6/24
                                 v6/DHCP6: 2001:db8::20c:29ff:fe78:6e4e/64
LAN (lan)       -> vmx1       -> v4: 10.6.0.1/24
                                 v6/t6: 2001:db8:1:eea0:20c:29ff:fe78:6e58/64

0) Logout (SSH only)                  9) pfTop
1) Assign Interfaces                 10) Filter Logs
2) Set interface(s) IP address       11) Restart GUI
3) Reset admin account and password  12) PHP shell + pfSense tools
4) Reset to factory defaults         13) Update from console
5) Reboot system                     14) Disable Secure Shell (sshd)
6) Halt system                       15) Restore recent configuration
7) Ping host                         16) Restart PHP-FPM
8) Shell
```

Page Contents

- *First Connection Behavior*
- *1) Assign Interfaces*
- *2) Set interface(s) IP address*
- *3) Reset admin account and password*
- *4) Reset to factory defaults*
- *5) Reboot system*
- *6) Halt system*
- *7) Ping host*
- *8) Shell*
- *9) pfTop*
- *10) Filter Logs*
- *11) Restart GUI*
- *12) PHP shell + pfSense tools*
- *13) Upgrade from console*
- *14) Enable/Disable Secure Shell (sshd)*
- *15) Restore recent configuration*
- *16) Restart PHP-FPM*

8.7.1 First Connection Behavior

On pfSense Plus software version 24.03 and later, during the first connection to the console or SSH after installation or resetting to factory defaults, the user is prompted to set a new password for the `admin` account.

This change is mandatory, however, it can also be performed in the GUI using the *Setup Wizard*, the *User Password Manager*, or the *User Manager*.

If the password has been changed in the GUI, press `Ctrl-C` to cancel the console password change prompt. The script will check the password again and if it has been changed, it will display the menu. If the password is still the default value, however, the user will be logged out.

8.7.2 1) Assign Interfaces

This option restarts the **Interface Assignment** task, which is covered in detail in *Assign Interfaces* and *Manually Assigning Interfaces*. This menu option can create VLAN interfaces, reassign existing interfaces, or assign new ones.

8.7.3 2) Set interface(s) IP address

The script to set an interface IP address can set WAN, LAN, or OPT interface IP addresses, but there are other useful features of this script:

- The firewall prompts to enable or disable DHCP service for an interface, and to set the DHCP IP address range if it is enabled.
- If the firewall GUI is configured for HTTPS, the menu prompts to switch to HTTP. This helps in cases when the SSL configuration is not functioning properly.
- If the anti-lockout rule on LAN has been disabled, the script enables the anti-lockout rule in case the user has been locked out of the GUI.

8.7.4 3) Reset admin account and password

This menu option invokes a script to reset the `admin` account and password.

The script takes a few actions to help regain access to the `admin` account:

- If the authentication source is set to a remote server such as RADIUS or LDAP, the script prompts to return the authentication source to the Local Database (User Manager).
- If the `admin` account has been removed, the script re-creates the account.
- If the `admin` account is disabled or expired, the script re-enables the account.

Once the `admin` account has been restored to a working state the script prompts to set and confirm a new password. This new password can then be used to log in to the `admin` account in the GUI, console, or SSH (if enabled).

Tip: This option can be used to change the password for the `admin` from the console instead of using the GUI.

Note: On previous versions of pfSense software this option reset the password to a default value (*Default Username and Password*). This is no longer the case as the best practice is to avoid using default passwords.

8.7.5 4) Reset to factory defaults

This menu choice restores the system configuration to factory defaults. It will also attempt to remove any installed packages.

This action is also available in GUI at **Diagnostics > Factory Defaults**.

See *Resetting to Factory Defaults* for more details about how this process works.

8.7.6 5) Reboot system

This menu choice cleanly shuts down the firewall and restarts the operating system. There are several options which control what the firewall will do when rebooting. The choices offered by the reboot option are explained in *Reboot Methods*.

See also:

This action is also available in GUI at **Diagnostics > Reboot**, see *Rebooting the Firewall* for details.

8.7.7 6) Halt system

This menu choice cleanly shuts down the firewall and either halts or powers off, depending on hardware capabilities.

> **Warning:** The best practice is to **never** cut power from a running system. Halting before removing power is always the safest choice.

See also:

This action is also available in GUI at **Diagnostics > Halt System**. See *Halting and Powering Off the Firewall* for additional details.

8.7.8 7) Ping host

This menu option runs a script which attempts to contact a host to confirm if it is reachable by the firewall through a connected network. The script prompts the user for an IP address, and then the script sends that target host three ICMP echo requests.

The script displays output from the test, including the number of packets received, sequence numbers, response times, and packet loss percentage.

The script uses `ping` when given an IPv4 address or a hostname, and `ping6` when given an IPv6 address.

This is only a basic ping test. For more options, see *Ping Host* to run a similar test from the GUI.

8.7.9 8) Shell

This menu choice starts a command line shell.

> **Warning:** A shell is very useful and very powerful, but also has the potential to be very dangerous.

Note: The majority of users do not need to touch the shell, or even know it exists.

Complex configuration tasks may require working in the shell, and some troubleshooting tasks are easier to accomplish from the shell, but there is always a chance of causing irreparable harm to the system.

Veteran FreeBSD users may feel slightly at home there, but there are many commands which are not present on pfSense software installations since unnecessary parts of the OS are removed for security and size constraints.

A shell started in this manner uses `tcsh`, and the only other shell available is `sh`. While it is possible to install other shells for the convenience of users, Netgate neither recommends nor offers assistance with other shells.

8.7.10 9) pfTop

This menu option invokes `pftop` which displays a real-time view of the firewall states, and the amount of data they have sent and received. It can help pinpoint sessions currently using large amounts of bandwidth, and may also help diagnose other network connection issues.

See also:

See *pfTop* for more information on how to use pfTop.

8.7.11 10) Filter Logs

The **Filter Logs** menu option displays firewall log entries in real-time, in their raw form. The raw logs contain much more information per line than the log view in the GUI (**Status > System Logs**, **Firewall tab**), but not all of this information is easy to read.

Tip: For a simplified console view of the firewall logs in real time with low detail, use the following shell command:

```
tail -F /var/log/filter.log | filterparser.php
```

8.7.12 11) Restart GUI

Restarting the webConfigurator will restart the system process that runs the GUI (`nginx`). In extremely rare cases the process may have stopped, and restarting it will restore access to the GUI.

If the GUI is not responding and this option does not restore access, invoke menu option 16 to **Restart PHP-FPM** after using this menu option.

8.7.13 12) PHP shell + pfSense tools

The PHP shell is a powerful utility that executes PHP code in the context of the running system. As with the normal shell, it is also potentially dangerous to use. This is primarily used by developers and experienced users who are intimately familiar with both PHP and the pfSense software code base.

See also:

See *Using the PHP Shell* for additional details and a list of available playback scripts.

8.7.14 13) Upgrade from console

This menu option runs the `pfSense-upgrade` script to upgrade the firewall to the latest available version. This is operationally identical to running an upgrade from the GUI and requires a working network connection to reach the update server.

This method of upgrading is covered with more detail in *Upgrading using the Console*.

8.7.15 14) Enable/Disable Secure Shell (sshd)

This option toggles the status of the Secure Shell Daemon, `sshd`. This option works the same as the option in the GUI to enable or disable SSH.

8.7.16 15) Restore recent configuration

This menu option starts a script that lists and restores backups from the configuration history. This is similar to accessing the configuration history from the GUI at **Diagnostics > Backup/Restore** on the **Config History** tab (*Configuration History*).

This script can display the last few configuration files, along with a timestamp and description of the change made in the configuration, the user and IP address that made the change, and the config revision. This is especially useful if a recent configuration error accidentally prevented access to the GUI.

8.7.17 16) Restart PHP-FPM

This menu option stops and restarts the daemon which handles PHP processes for `nginx`. If the GUI web server process is running but unable to execute PHP scripts, invoke this option. Run this option in conjunction with **Restart webConfigurator** for the best result.

8.8 Resetting to Factory Defaults

The firewall configuration can be reset back to defaults, a process which also attempts to remove any installed packages. This reset can be performed in the GUI from **Diagnostics > Factory Defaults**, by using the console menu, or in some cases by using a hardware button.

In each case, the firewall will automatically reboot with a default configuration after the reset, which may require console access to resolve.

Note: This process does not remove any changes made to the file system, it only resets the configuration.

If system files have been corrupted or altered in an undesirable way, the best practice is to make a backup and reinstall from installation media.

8.8.1 Factory Default from the GUI

To reset the configuration to factory defaults using the GUI:

- Navigate to **Diagnostics > Factory Defaults**
- Review the items on the page which will be affected by the reset
- Click  **Factory Reset**
- Click **OK** to confirm the action and start the reset process

8.8.2 Factory Default from the Console

To reset the configuration to factory defaults using the console:

- Access the console menu locally or via SSH with an admin-level account (`admin`, `root`, or another privileged account using `sudo`).
- Enter the menu option which corresponds with **Reset to factory defaults** (e.g. `4`)
- Press `Enter`
- Enter the `y` to confirm the action
- Press `Enter` to start the reset process

8.8.3 Factory Default using a Hardware Button

On some appliances from Netgate, the reset button can perform a factory reset. If the button is recessed, it may be depressed with a paperclip or other similar object. Some models require pressing and holding the button during the boot sequence, while others allow pressing and holding the button at any time.

> **Warning:** Reset button behavior varies by hardware. Some hardware performs an immediate hard reset and reboots when pressing the reset button.
>
> Check the Netgate Product Manuals to confirm button behavior before attempting this procedure.

The exact procedure varies by hardware and is explained in the Netgate Product Manuals for models with this capability.

One common procedure is:

- Apply power to the unit
- Depress the reset button after the initial POST sequence completes
- Hold the reset button in until the system LEDs turn off or the system reboots

Another common procedure is:

- Allow the unit to boot if it is not already running
- Depress the reset button and monitor the status LEDs
- Hold the reset button in and wait for all the lights to flash in unison
- Release the button and wait for the device to reboot

The unit will reset the configuration to factory defaults and reboot again with that default configuration.

8.9 XML Configuration File

pfSense® software stores its settings in an XML format configuration file. All configuration settings including settings for packages are held in this one file. Run-time configuration files for services and firewall behavior are generated dynamically based on the settings held within this XML configuration file.

Those familiar with FreeBSD and related operating systems have found this out the hard way, when their changes to system configuration files were repeatedly overwritten by the firewall before they came to understand that pfSense software handles everything automatically.

The configuration file is stored at `/conf/config.xml` on the firewall.

8.9.1 Manually editing the configuration

A handful of configuration options are only available by manually editing the configuration file, though this isn't required in the vast majority of deployments. Some of these options are covered in other parts of this documentation where they are relevant. Additionally, for advanced administrators in rare cases large-scale or tricky changes may be easier to make by directly editing the configuration file.

Warning: Even for seasoned administrators it is easy to incorrectly edit the configuration file. Always keep backups and be aware that breaking the configuration will result in unintended consequences.

Edit a Backup

The safest and easiest method of editing the configuration file is to make a backup, edit the backup, and then restore:

- Navigate to **Diagnostics > Backup/Restore** in the GUI
- Download and save backup file
- Open the file in a text editor that properly understands UNIX line endings, and preferably an editor that has special handling for XML such as syntax highlighting. Do not use `notepad.exe` on Windows.
- Make changes to the configuration and save
- Navigate to **Diagnostics > Backup/Restore** in the GUI
- Restore the edited configuration

The firewall will automatically reboot as a part of the restoration process, and the new settings will be active afterward.

Edit In Place

Editing the configuration in-place is also possible in a variety of ways. The general procedure is:

- Edit `/conf/config.xml`
- Run `rm /tmp/config.cache` to clear the configuration cache
- Reboot, or use the GUI to save/reload whichever part of the firewall utilizes the edited settings

From the console or ssh, administrators familiar with the vi editor can use the `viconfig` command to edit the running configuration, and this command automatically clears the cache file after saving and exiting.

Other editors are available on the firewall, such as `ee` or in the GUI under **Diagnostics > Edit File** (*Editing Files on the Firewall*). Clear the cache file manually after using one of these other methods, either using the shell or **Diagnostics > Command Prompt** (*Command Prompt*).

8.10 pfSense® Plus Software Registration

The pfSense Plus Software Registration page is located at **System > Register**. This page activates features in pfSense Plus software installations on hardware and virtual machines not purchased from Netgate. The page also activates older Netgate® hardware purchased with Factory Edition (FE) pfSense software before the Netgate Device ID (NDI) was introduced.

The registration process requires an activation token supplied by Netgate. This token is generated when purchasing pfSense Plus software or via Netgate TAC for older Netgate hardware.

For more information about pfSense Plus software, or to purchase pfSense Plus software, visit Netgate Store.

Note: Registration is free for hardware purchased from Netgate with pfSense Plus software or the older Factory Edition of pfSense software. Most hardware is pre-registered and does not require activation. To activate hardware which is not automatically recognized, submit a request to Netgate TAC along with the serial number and NDI for the device at https://go.netgate.com. The serial number and NDI are displayed on the dashboard in the **System Information** widget, and may also be on a sticker located on the bottom of the device.

The current registration status is shown on the dashboard in the **Netgate Services and Support** widget, and is also indicated on **System > Register**.

The text on the registration page varies depending on the current registration status and availability. The page also displays errors encountered during the activation process, such as not being able to contact the registration server.

8.10.1 Registration Process

To register an installation of pfSense Plus software with Netgate:

- Obtain a pfSense Plus software activation token from Netgate
- Navigate to **System > Register** on the firewall
- Enter the **Activation Token**
- Click **Register**

See also:

- *Basic Firewall Configuration Example*
- *Troubleshooting Clock Issues*
- *Troubleshooting*
- *Troubleshooting Access when Locked Out of the Firewall*
- *Troubleshooting Time Zone Configuration*

Most pfSense® software configuration is performed using the web-based GUI. There are a few tasks that may also be performed from the console, whether it be a monitor and keyboard, over a serial port, or via SSH.

8.11 Connecting to the GUI

To reach the GUI, follow this basic procedure:

- Connect a client computer to the same network as the LAN interface of the firewall. This computer may be directly connected with a network cable or connected to the same switch as the LAN interface of the firewall.

  By default, the LAN IP address of a new installation of pfSense software is `192.168.1.1` with a `/24` mask (`255.255.255.0`), and there is also a DHCP server running. If a client computer is set to use DHCP, it should obtain an address in the LAN subnet automatically.

- On the client computer, open a web browser such as Firefox, Safari, or Chrome and navigate to https://192.168.1.1.

  The GUI listens on HTTPS by default, but if the browser attempts to connect using HTTP, the firewall will redirect the browser to the HTTPS port.

- Enter the default credentials in the login page:

username
`admin`

password
`pfsense`

In some cases additional steps may be necessary before the client computer can reach the GUI.

Warning: If the default LAN subnet conflicts with the WAN subnet, the LAN subnet must be changed before connecting it to the rest of the network. Attempting to access the GUI in this situation is unpredictable and unlikely to work until the conflict is resolved.

The LAN IP address may be changed and DHCP may be disabled using the console:

- Open the console (VGA, serial, or using SSH from another interface)
- Choose option `2` from the console menu
- Enter the new LAN IP address, subnet mask, and specify whether to enable DHCP.
- Enter the starting and ending address of the DHCP pool if DHCP is enabled. This can be any range inside the given subnet.

Note: When assigning a new LAN IP address, it cannot be in the same subnet as the WAN or any other active interface. If there are other devices already present on the LAN subnet, it also cannot be set to the same IP address as an existing host.

If the DHCP server on the firewall is disabled, client computers on LAN must have a statically configured IP address in the LAN subnet, such as `192.168.1.5`, with a subnet mask that matches the one given to the firewall, such as `255.255.255.0`.

CHAPTER

NINE

NETGATE® NEXUS

Netgate Nexus is the multi-instance management (MIM) system for pfSense Plus software. Designed to address the growing complexity of managing multiple firewall instances across distributed environments, Netgate Nexus empowers network operators to securely manage hundreds of pfSense Plus instances through a unified, intuitive GUI and REST API.

The Netgate Nexus controller is included in pfSense Plus software versions 25.07 and later.

9.1 Netgate Nexus Controller Setup

Before instances of pfSense Plus software can be registered to the Netgate Nexus controller for tasks such as multi-instance management (MIM), there are several setup tasks to complete.

9.1.1 Enable Netgate Nexus Controller

The Netgate Nexus controller must be enabled and running before registering instances.

- Open the pfSense Plus software WebGUI on the designated controller
- Navigate to **System > Advanced**, **Netgate Nexus** tab
- Check **Enable**
- Configure any other options as needed (*Netgate Nexus Controller Configuration Options*)
- Click **Save**

9.1.2 Firewall rules for Netgate Nexus

The Netgate Nexus controller **does not** automatically add firewall rules for the Netgate Nexus GUI or external controller VPN connectivity. Firewall rules are necessary for instances to connect the VPN itself and for administrators to reach the Netgate Nexus GUI. Configure these firewall rules on the controller host in the pfSense Plus software WebGUI.

Note: The Netgate Nexus controller automatically passes traffic tunneled through its VPN between the instances and the controller. There is no need to manage rules for that internal communication.

Allowing Incoming Netgate Nexus VPN Connections

Add a rule on WAN to pass connections to the Netgate Nexus VPN port.

- Open the pfSense Plus software WebGUI on the designated controller
- Navigate to **Firewall > Rules**, **WAN** tab

> **Note:** WAN is used as an example. This could also be any other interface to which instances will connect.

- Click  to add a new rule at the top of the list:
- Configure the rule with the following options:

  Action
  Pass

  Protocol
  UDP

  Source
  Any

  > **Note:** This is acceptable if instances have dynamic addresses. If all instances are static, consider creating an alias to allow only those addresses.

  Destination
  This Firewall (self)

  > **Note:** This could also be the specific interface or IP address instances use when connecting.

  Destination Port

  From
  (Other)

  Custom
  `_nexus_vpn_port_`

  > **Note:** This is a built-in alias which automatically contains the random port the controller selected to use for incoming VPN connections.

- Click **Save**
- Click **Apply Changes**

Allowing Netgate Nexus GUI Access

Access to the Netgate Nexus GUI is also restricted by firewall rules. If local interfaces or VPNs are restricted, rules must be added there as well. The ports for those rules are configured in the Netgate Nexus options (*General Options*).

> **Danger:** **Do not expose this port to the Internet**. Limit access as much as possible. Use a VPN for remote access.
>
> As with the pfSense Plus software WebGUI, the best practice is to restrict access to specific management hosts, networks, or VPN clients.

9.1.3 Accessing the Netgate Nexus GUI

To access the Netgate Nexus GUI, follow the links in the Netgate Nexus status under **System > Advanced**, **Netgate Nexus** tab (*Viewing Netgate Nexus Status*).

Use the HTTPS link to securely access the Netgate Nexus controller.

Note: If the Netgate Nexus controller is using a self-signed TLS certificate, then it may be necessary to click through an error in the browser warning about the validity of the self-signed certificate.

9.1.4 Netgate Nexus Authentication

After following the link, the controller will display a login screen.

Tip: Bookmark this page for faster access.

The Netgate Nexus controller uses the pfSense Plus software *User Manager*, so the same credentials will work for the Netgate Nexus controller that work for the pfSense Plus software WebGUI.

Enter valid credentials and click **Sign In** to access the Netgate Nexus GUI.

9.2 Netgate Nexus Options in the pfSense Plus WebGUI

The pfSense Plus software WebGUI contains options to manage the Netgate Nexus controller itself. These options are at **System > Advanced** on the **Netgate Nexus** tab.

9.2.1 Enabling Netgate Nexus

The Netgate Nexus controller must be *enabled and running* before it can be accessed by browsers or other instances.

- Open the pfSense Plus software WebGUI
- Navigate to **System > Advanced**, **Netgate Nexus** tab
- Check **Enable**
- Configure any other options as needed (*Netgate Nexus Controller Configuration Options*)
- Click **Save**

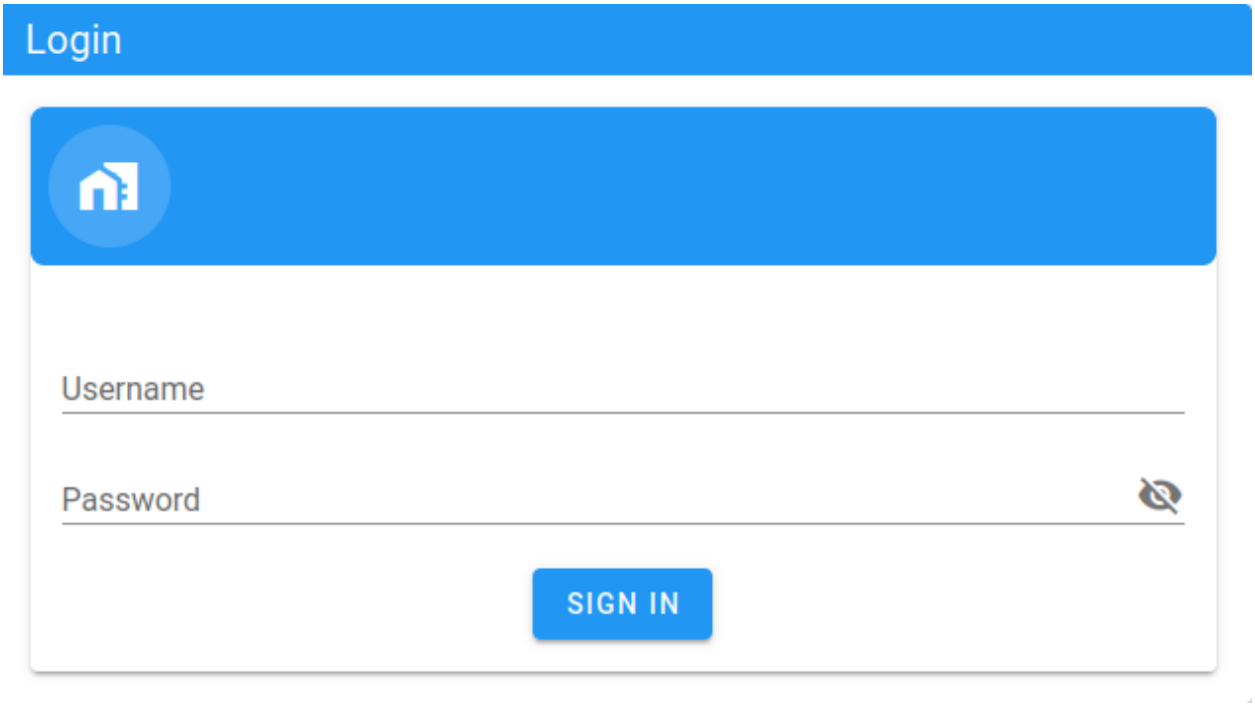


Fig. 1: Netgate Nexus Controller Login Screen

9.2.2 Netgate Nexus Controller Configuration Options

The default options are acceptable for most environments, but the behavior of the Netgate Nexus controller daemon can be fine-tuned by the options on the page at **System > Advanced**, **Netgate Nexus** tab.

These options are broken into two sections, **General Options** and **Advanced Options**.

General Options

Enable
Controls whether or not the Netgate Nexus controller daemon is enabled or disabled.

TLS Certificate
The TLS certificate the Netgate Nexus controller will use when acting as a TLS server (e.g. HTTPS).

This is typically the same certificate used by the pfSense Plus software WebGUI, but it can be a different certificate. This can be a local self-signed certificate, a globally trusted certificate imported into the WebGUI, or even a certificate managed by the *ACME package*.

Service Ports
The Netgate Nexus controller daemon uses multiple ports to accept connections:

API HTTP Port
The port used for unencrypted communication from browsers and API clients.

> **Danger:** Communication on this port is not encrypted. The best practice is to only use encrypted communication, so avoid using this port outside of local testing and development.

API HTTPS Port
The port used for encrypted communication from browsers and API clients.

Note: The controller picks a random port to use for the VPN. This port is available for use in firewall rules through the built-in alias named `_nexus_vpn_port_`. This port may be set to a specific value by setting **Listening Port** in the *Advanced Options*.

Advanced Options

Logging
Controls the behavior of messages logged by the Netgate Nexus controller daemon.

Level
The type of messages to log. Each level also includes messages from levels below it in the list.

Verbose
When checked, increases the verboseness of log messages.

Listening Address
A specific address the Netgate Nexus controller daemon will use to listen for incoming VPN connections from remote instances.

When set to *Any* (default), the controller listens on all available addresses.

Note: This does not control binding of the Netgate Nexus GUI service, only the VPN used for communication between instances and the controller.

Listening Port
A specific port the Netgate Nexus controller daemon will use to listen for incoming VPN connections from remote instances.

Advertised Addresses
One or more IP addresses or fully qualified domain names (FQDNs), optionally with a port number, which the controller will advertise to instances. When blank (default), the controller advertises all of its IP addresses for auto-discovery.

Tip: Using an FQDN here in combination with some form of Dynamic DNS can enable instances to reach the controller in cases where the controller host does not have a static IP address on any WAN interface.

JWT Session Expiry
Sets the expiration time of session tokens. Lower values are more secure, but require more frequent requests for new tokens.

Custom Options
Custom controller options. Each option must be on a separate line.

Warning: Do not use this section unless requested to do so by developers or TAC.

9.2.3 Viewing Netgate Nexus Status

The Netgate Nexus settings tab also includes status information for the controller daemon while the daemon is running. This includes links to the Netgate Nexus GUI as well as an entry indicating the VPN port upon which the controller is listening.

To view this status information:

- Open the pfSense Plus software WebGUI on the designated controller device
- Navigate to **System > Advanced**, **Netgate Nexus** tab
- Look in the bottom section of the **General Options** area

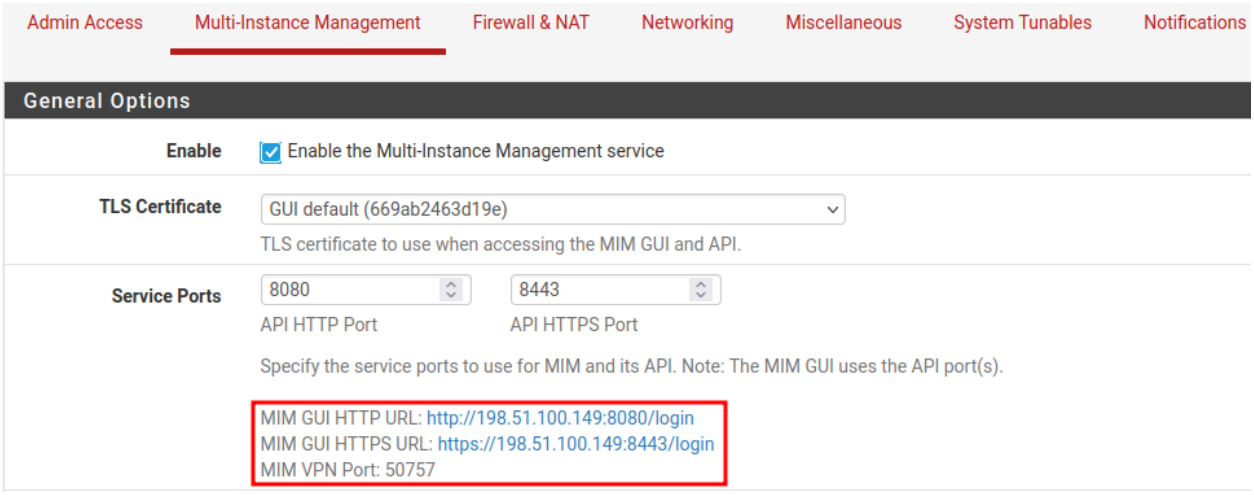


Fig. 2: Netgate Nexus Controller Daemon Status

9.3 Netgate Nexus Licensing

A Netgate Nexus license resides on the instance which performs the role of the controller. Customers can purchase entitlements from the Netgate Store which they can then apply to this license. At any given time, there will only be **one** active license (with its various entitlements) in use on the controller.

This document explains how to obtain an initial license, as well as how to apply tokens to update the entitlements of the license.

Note: This guide assumes the user has already enabled Multi-Instance Management as described in *Netgate Nexus Controller Setup*.

9.3.1 Licensing Requirements

Netgate Nexus requires pfSense Plus 25.11 or later running on Netgate® appliances, or on AWS and Azure virtual instances. While Netgate Nexus may work on third-party enterprise class hardware, Netgate cannot guarantee it will function and strongly cautions customers against using such hardware.

To check device eligibility, enable Netgate Nexus and visit the **License** menu on the devices. If Netgate Nexus cannot run on the installation, it will display an error message.

Controlled instances must be able to access the Controller. Air-gapped, offline, virtual and some hardware systems without serial numbers or other identifiers are not capable of participating in Netgate Nexus and cannot be managed.

9.3.2 Initializing the License

To start using the licensing system, the first step is to initialize a license.

To initialize a license:

- *Open the |nexus| GUI* on the designated controller and login
- Click the **License** menu item on the left side to open the license system

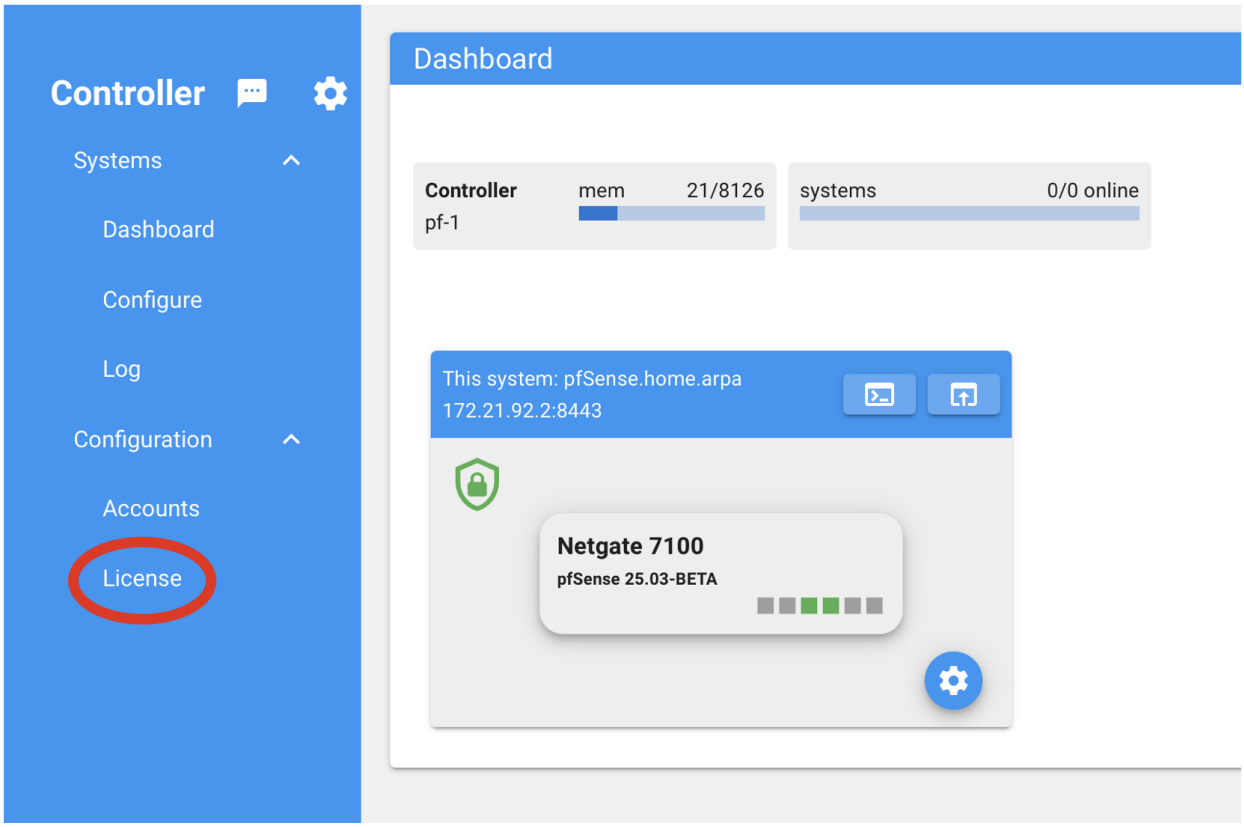


Fig. 3: Netgate Nexus License Menu Entry

- Inspect the **Current License** area of the screen to ensure the table is empty, as depicted in *Netgate Nexus License System with Add button marked*

Note: An empty table indicates no license has been initialized on this controller.

- Click  to create a license request, which initializes the license

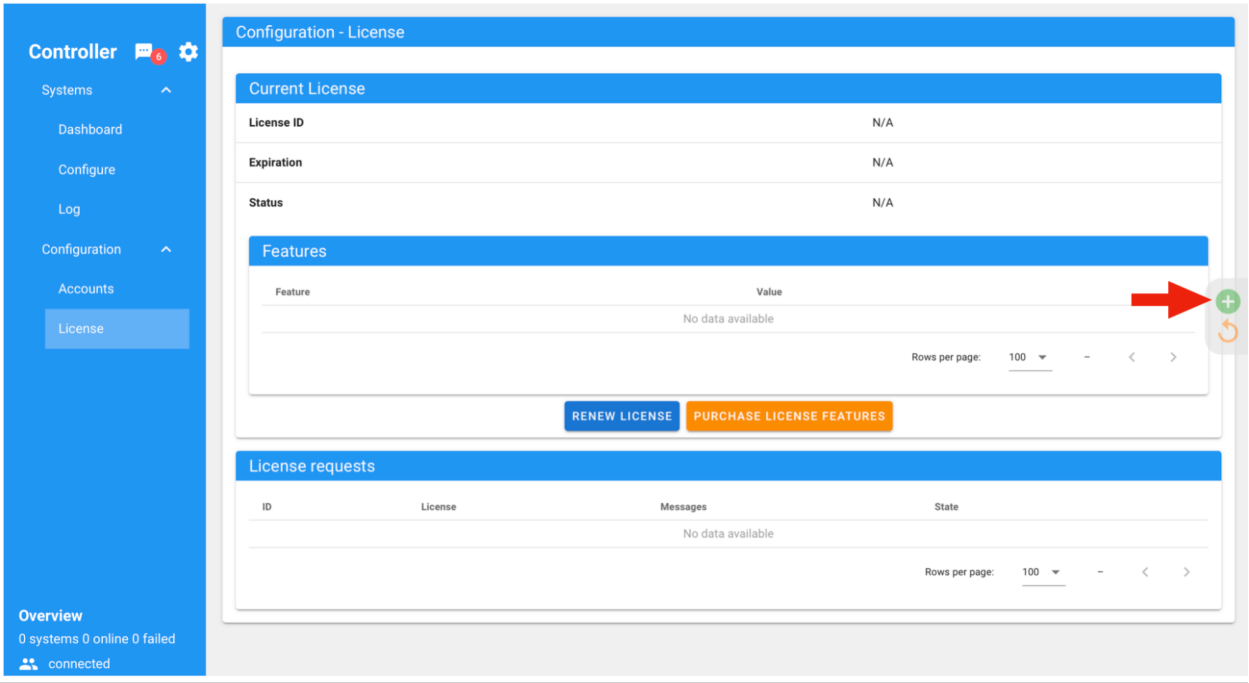


Fig. 4: Netgate Nexus License System with Add button marked

- Fill in the **Name** and **Email** fields on the **Create License Request** form

Note: Leave the **License Request Type** set to **License Signing Request**.

- Click **OK**

This initializes the license, which the GUI then displays in the license system screen.

9.3.3 License Details

When viewing the license, the License System displays the following details on the screen:

License ID
: A four-part phrase (e.g. `word1-word2-word3-word4`) which uniquely identifies the license.

Expiration
: The expiration date of the license, typically a one-year period.

Status
: The current status of the license, e.g. **Valid**.

Features
: A list of entitlements applied to this license.

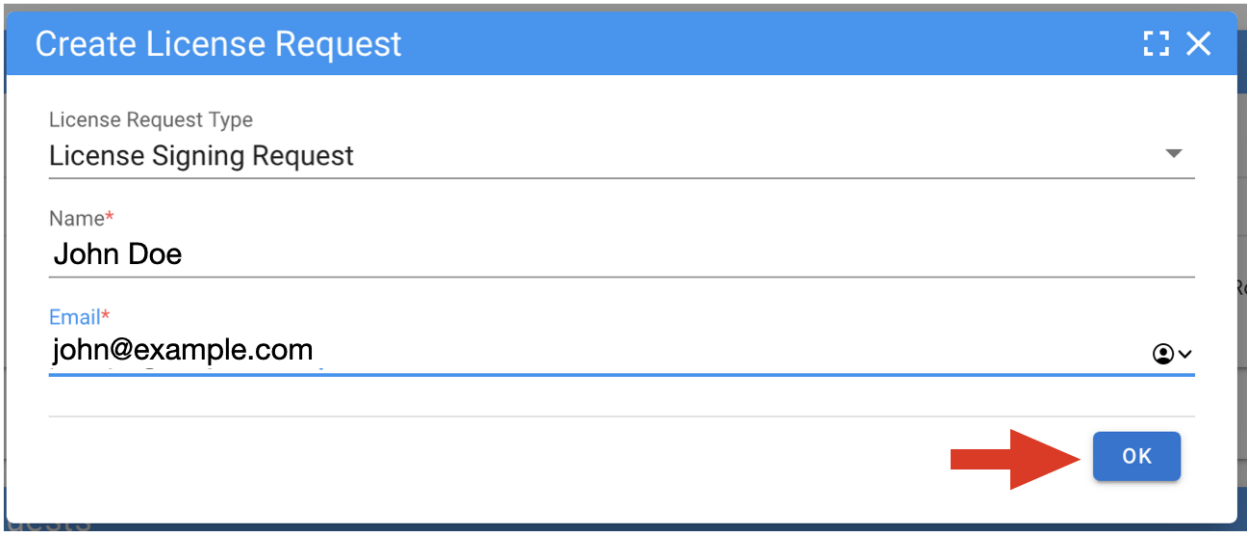


Fig. 5: Netgate Nexus License

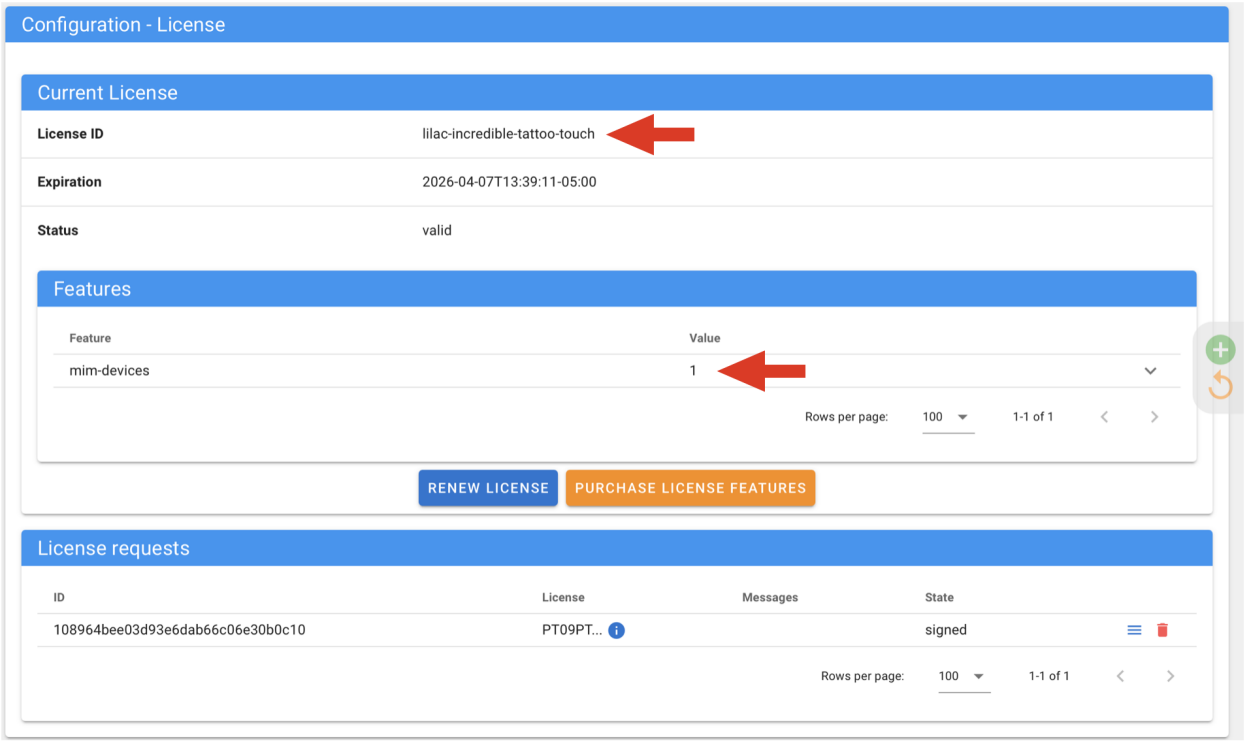


Fig. 6: Netgate Nexus License Details

mim-devices
The number of devices which can be controlled by multi-instance management from this device. The initial value is 1 by default, which means the current license may control **one** instance of pfSense Plus software, which is the instance upon which the controller itself is running.

License Requests
A list of license requests associated with the current license and entitlements.

9.3.4 Increasing the Managed Instances

To increase the number of instances which can be managed by this controller, customers can purchase entitlements from the Netgate Store.

On the Multi-Instance Management product page, choose the quantity of additional instances to control, then complete the purchase.

The store sends several email messages after a purchase. One of these email messages contains the title "Netgate License Tokens" and looks like *Netgate Nexus License Token Text*:

Hello!

Thank you for your Netgate Purchase!
Your order number is: SO25-425733

Below is the license information for your purchase.
Please follow the instructions in our documentation to install your token.
Netgate Nexus Documentation

SKU: LIC.NEX-MIM-D1
Quantity: 6
Token Uses: 1
Expires: 2025-06-06 18:41:27.7068 UTC
Token:

eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9.eyJleHBpcmVzIjoiMjAyNS0wNi0wNiAxODo0MT
oyNy43MDY4IiwiaXNzdWVkIjoiMjAyNS0wNC0wNyAxODo0MToyNy43MDYgKzAwMDAgVVRDIiwib
3JkZXIiOiJTTzI1LTQyNTczMyIsInJlZmVyZW5jZSI6MTkwMzUyOTE2Mzc1Njc2OTE0NiwidmVu
ZG9yIjoic2hvcGlmeSJ9.ixQBFi4VDyVIbKoO63NJl3MDUHEFLI9KE_1fSQl0kg_8Ui8j6rEPsn
VZYKyx8o9GaG_WZZRCHpgG-ZLkUgXf-B-piJ-4lmaazIb-LQ-6legSF0wNvYY8lXIA5oaJB9JCN
2YO2UwwHpyQ96ADn5qKTnhXDaIq40dGMJsyEBLxQG1Ac1xPnvgoezGvgbKVsI-4o2L0dueVhL1a
nfi7lSkMtQ-CWFi4oztKlJhQrMPh37KYRdP8dN2OcA03OIU2m-P5jTKbOF_kYsXejf8WDtTmGmH
Iap_CyvEAA4wsPcHLLh9skuqO1g3A24oNISLuQBkBLqcqKlrdORszAFavsdEdYpX7LFcEr2pHt1
QixtYSRw6HyCQMTHAjhryZMRrb2bv8PNmIrKG4h1GIp_nhMd3UQyz8TeBP5yEhXOnVyMgcsfE6l
IgKotrAWTDdAcB_-o5b9bdx36M4OeHJTwLDLRUkbodNdE5CfqvrhtAEjQUrgrSRbNHc4a-ygeT0
I0GO27nR0LjGOfIyqR8QA8xcOb8u7O-eai8XFyp-MTv8sjVHTVIUBjvMGBmscpAMdmqJtACUktg
oNJnsqoAj5iLnSVxWdR8LeI0sU95BVk1OzDW2YoDW3wT9uwimKcJ3CrMbjKHypRT1eGuAlbeZm9
au1xbB85F5PomJ6YOmIzVhhOsyMdA

Best regards,
Your Netgate team

Please don't reply to this email.
You received this email because you purchased a subscription from Netgate. If you were not expecting this email, please Contact Netgate
Reddit | YouTube | Forum
Privacy Policy

Fig. 7: Netgate Nexus License Token Text

The example purchase depicted in *Netgate Nexus License Token Text* contains entitlements for **six** additional managed devices.

To add this entitlement to the license:

- Select the token text in the email message, as depicted in *Netgate Nexus License Token Text*
- Copy the text to the clipboard
- *Open the |nexus| GUI* on the designated controller and login
- Click the **License** menu item on the left side to open the license system
- Click  to create a license request
- Paste the token text into the **License Token** field



Fig. 8: Adding a Netgate Nexus License Token

- Click **OK**

The license details now contain the new **License Request** in the list at the bottom of the screen and the **mim-devices** count reflects the new total of device entitlements: 7:

Customers can add more managed instances to a license at any time by following the above procedure again.

See also:

Customers with questions can reach out to Netgate Sales or Netgate TAC.

9.4 Netgate Nexus GUI

The web-based Netgate Nexus graphical user interface (GUI) has multi-instance management (MIM) functionality which allows administrators to manage pfSense Plus software installations registered with the controller.

From within the Netgate Nexus GUI, administrators can view status information, configure instances, use remote consoles, and more.

Note: This document assumes the Netgate Nexus controller is enabled (*Netgate Nexus Options in the pfSense Plus WebGUI*), configured (*Netgate Nexus Controller Setup*), and has instances registered (*Instance Registration*).

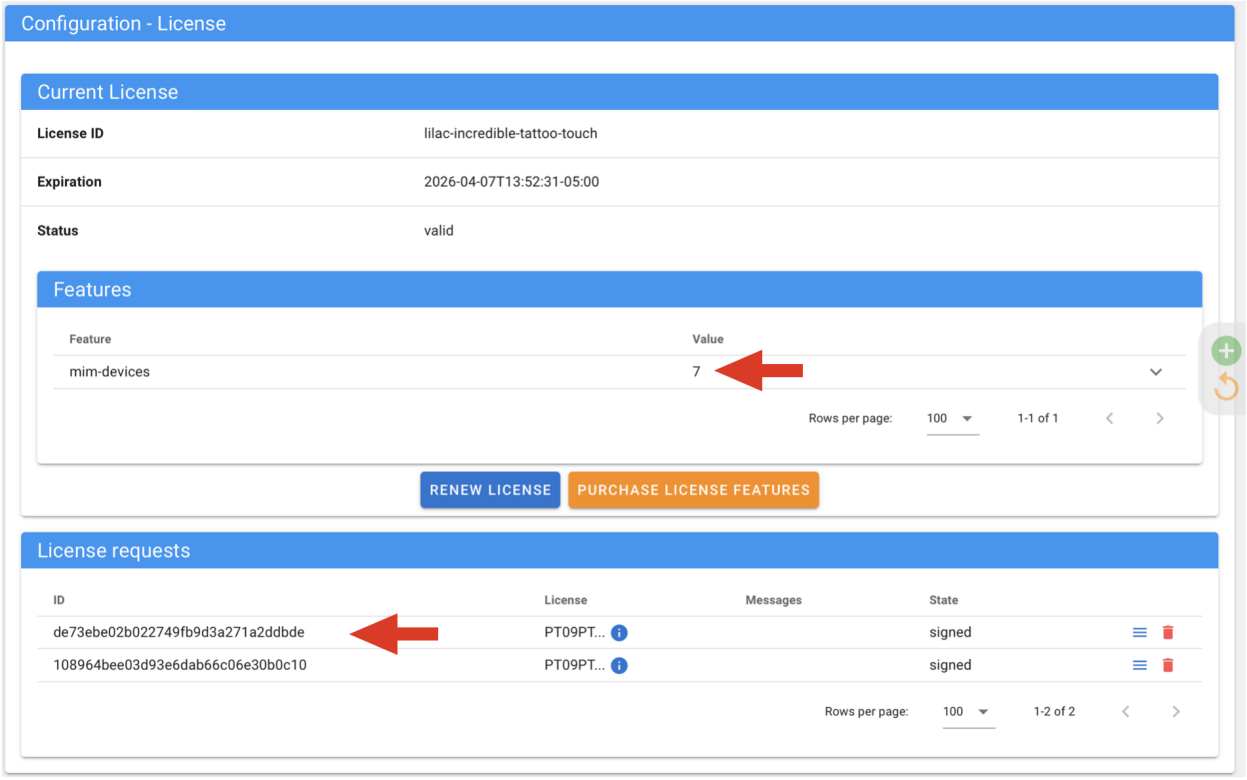


Fig. 9: Netgate Nexus License with Entitlements

9.4.1 Netgate Nexus Dashboard

The Netgate Nexus GUI dashboard contains an overview of all instances managed by this controller (*Multi-Instance Management*).

Access the Netgate Nexus GUI

Before proceeding, open the Netgate Nexus GUI as described in *Accessing the Netgate Nexus GUI* and login.

Tip: Accessing the Netgate Nexus GUI on the controller will allow administrators to manage any instance registered with the controller. The Netgate Nexus GUI on instances can only manage the local instance.

After logging in, the Netgate Nexus GUI displays its dashboard.

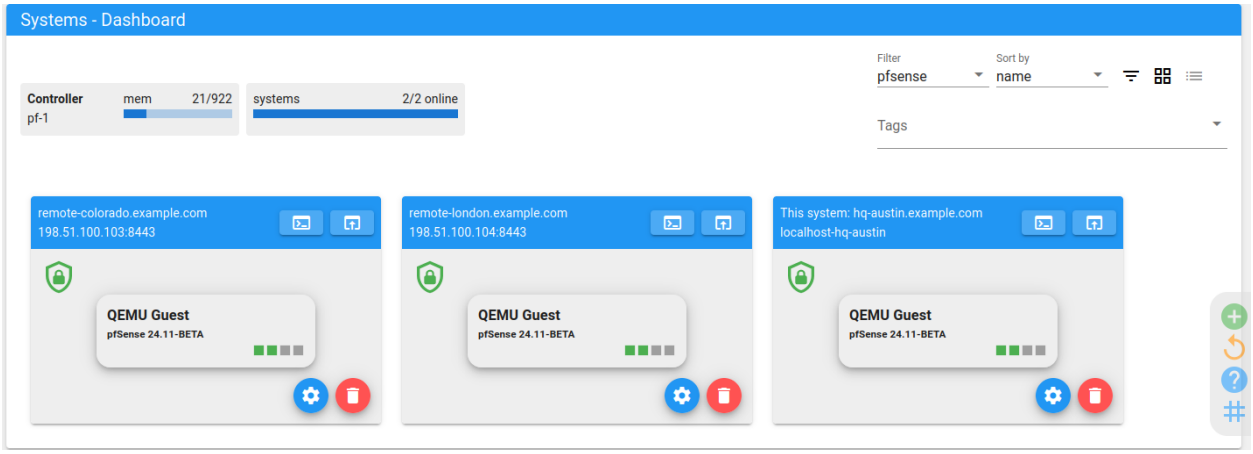


Fig. 10: Netgate Nexus Controller Dashboard

Status Summary

The dashboard contains status summary items in the upper left.

The first status item reports the current controller memory usage, both in text form and as a gauge.

The second status item reports the total number of registered instances and their online status.

Note: The controller host itself is not included in the systems summary as it would be redundant – if the controller was offline, the dashboard could not be loaded.

Instance List Entries

The instance list in the dashboard contains an entry for every instance registered with the controller, plus the controller itself.

Each list entry contains several items indicating the status of the instance, management controls, and options.

Hostname
: The hostname of the instance is printed in the bar at the top of the instance entry at the upper left.

Address
: The IP address and Netgate Nexus GUI port is printed in the bar at the top of the instance entry under the hostname.

Management Controls
: The upper right area of the bar at the top of the instance entry contains two icons to manage the instance:



Opens a *Remote Console* on the instance.



Enters the *Remote Configuration* interface for the instance.

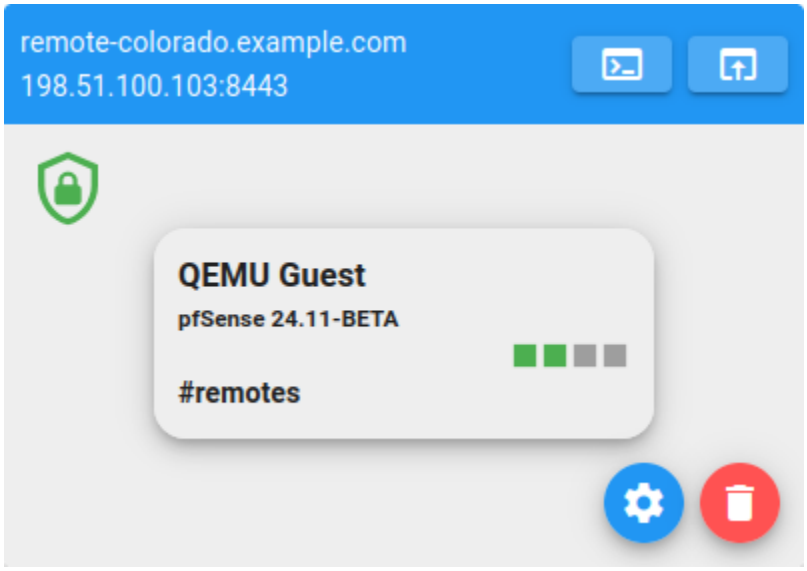


Fig. 11: Instance List Entry

Status Icon
In the upper left of the main area of the instance entry is an icon indicating its status. Possible states include:

- Online: 
- Error: 

Device Type
The first line of text in the instance body is the type of device. This can be a hardware model, hypervisor type, cloud provider, etc.

Version Information
The current software version is printed under the device type.

Interface Status
The instance entry contains a row of squares indicating network interface status for the device. Connected interfaces are green, disconnected or disabled interfaces are gray. Hover the mouse cursor over a square to see its name and link speed.

Tags (If present)
Entries can be classified via tags (*Tag Manager*). If an instance has assigned tags, they are printed here.

Entry Options
The  icon opens a modal dialog window with general options for this instance entry.

Remove Entry
The  icon deregisters (removes) this entry from the controller.

Instance List Filter & Sort

The list of instances on the dashboard can be filtered and sorted to locate and manage items easier. The options to filter and sort the list are in the upper right area of the dashboard.



Fig. 12: Instance List Filter & Sort Controls

In that area, the following options are available:

Filter
: This drop-down menu allows filtering the list by item type.

Sort By
: This drop-down menu sorts the items on the dashboard by a specific attribute.

  The current sort options are:

  Name
  : Sorts by instance hostname.

  State
  : Sorts by online/error status.

  Address
  : Sorts by IP address.

  Type
  : Sorts by device type.

Tags
: Enter a tag name to filter the list so it only shows items matching a specific tag or multiple tags.

Sort Order

: The sort order icon toggles between ascending ( ) and descending ( ) sort order.

Tile View

: Clicking the icon switches the dashboard to tile view, which is the default.

Detail View

: Clicking the icon switches the dashboard to a detailed list view.

Dashboard ToolBox

The Netgate Nexus Dashboard ToolBox is located on the right side of the dashboard screen and contains several options to manage dashboard items.

This ToolBox can appear in two different ways as determined by the *Controller Options*.

- It can be a floating bar pinned to the right side of the screen:



Fig. 13: Dashboard ToolBox - Right Side Style

- It can be collapsed behind a  hamburger menu button:

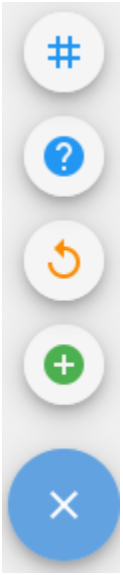


Fig. 14: Dashboard ToolBox - FAB Style

The available actions are the same in both styles.

The actions available in the ToolBox are:



Register an instance with the controller (*Instance Registration*).



Reload dashboard data.



Open the Help menu which also contains the current controller data for use in the registration process (*Instance Registration*).



Opens the tag manager which assigns tags to dashboard list entries (*Tag Manager*).

Tag Manager

Dashboard items can be tagged to allow filtering entries based on custom keywords. To assign tags to instances, open the Tag Manager from the *Dashboard ToolBox*.

The Tag Manager can add or remove tags from one or more instances as follows:

- Click **Select Devices**
- Pick devices from the list, or click **Select All**
- Click in the **Tags** field
- Enter a tag name
- Press `Tab`, press `Enter`, or click out of the **Tags** field
- Click the desired action:  to add the tag to the selected instances, or  to remove the tag from the selected instances.

9.4.2 Nexus Controller

The topics in this section cover areas of the Netgate Nexus GUI which apply to the controller itself.

Nexus Menu

The Nexus Menu is on the left side of the Netgate Nexus GUI.

At the top right of the menu are icons for *controller messages* () and *controller options* ().

At the bottom of the menu is an **Overview** area with a text summary of instances and their status, along with a count of users currently logged into the controller.

The main body of the menu consists of several areas of interest, including:

Dashboard
: The Netgate Nexus dashboard (*Netgate Nexus Dashboard*)

Configure
: Configure an instance (*Remote Configuration*).

  While configuring an instance, the menu lists its address in the menu.

Log
: View Netgate Nexus controller logs.

Software
: Netgate Nexus image and package management.

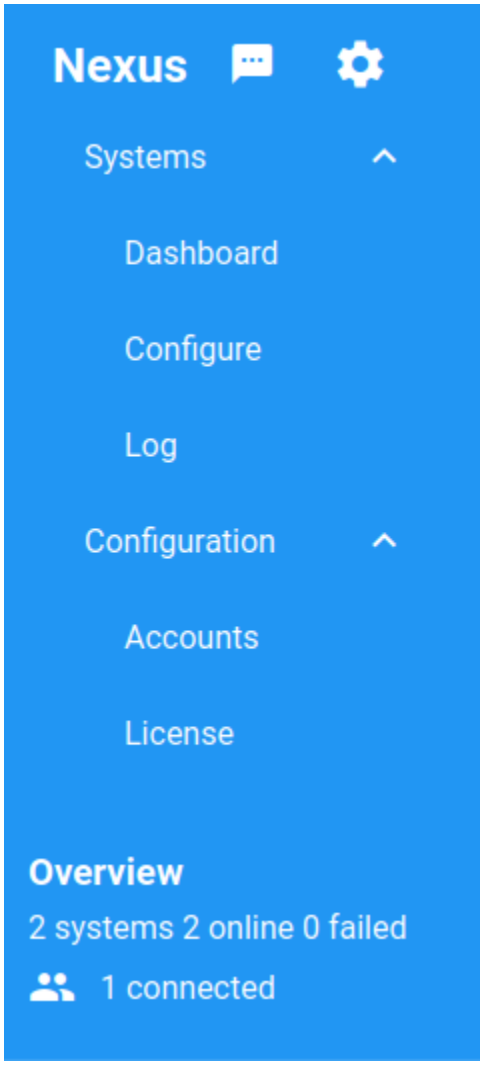


Fig. 15: Nexus Menu

Accounts
User authentication configuration. This configuration is linked to the pfSense Plus software User Manager on the controller host.

Controller Messages

Click in the *Nexus Menu* to view messages and alerts from the Netgate Nexus controller.

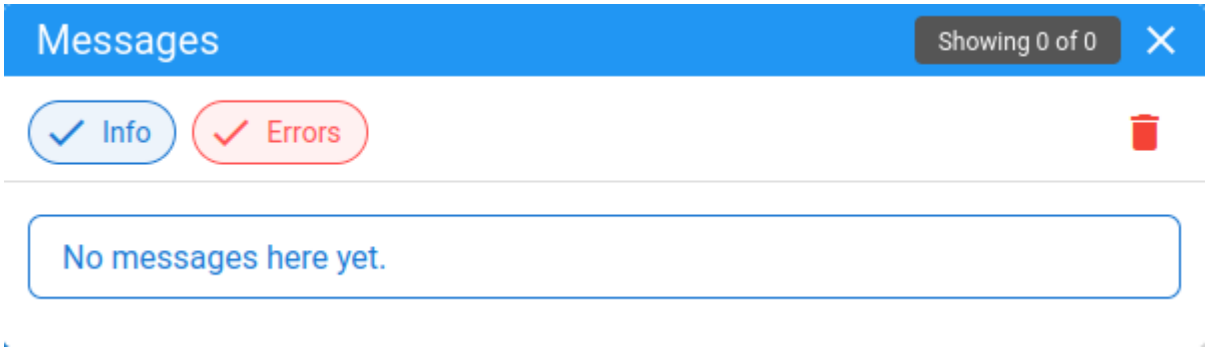


Fig. 16: Messages List

The messages can be filtered by type. Click **Info** or **Errors** to hide/show that particular type of message. If the type has a check mark icon, messages of that type will be included in the list. The window title bar contains a count of messages displayed in the current view, plus a total.

Click at the top right of the window to clear all messages.

Controller Options

Click in the *Nexus Menu* to configure options which affect the behavior of the controller itself and the GUI. This dialog window also contains some status information.

Options

The top area of the dialog contains options which affect the appearance of the controller GUI.

Theme Color
Select one of the colors to change the GUI theme to a scheme based on the chosen color.

Dark Mode
Toggles the GUI theme between light and dark mode.

ToolBox View
Changes the appearance of the *ToolBox* on the right side of the page.

Right-side ToolBox
Presents the ToolBox as a free-floating menu always visible on the right side of the page (*Dashboard ToolBox - Right Side Style*).

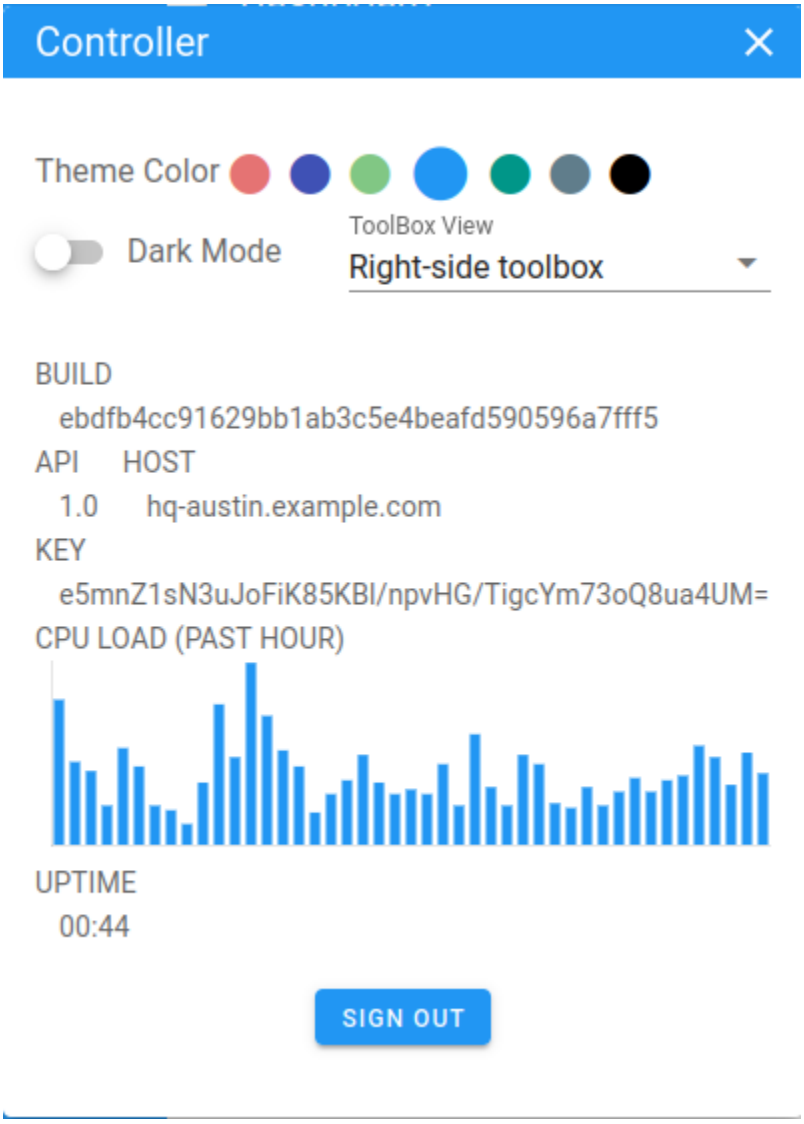


Fig. 17: Controller Options Dialog

FAB Toolbox

Hides the ToolBox behind a  hamburger menu button which must be clicked to see the individual menu entries (*Dashboard ToolBox - FAB Style*).

Status Information

The middle portion of the window contains status information about the controller.

Build
The current build of the controller software.

API Version
The current API version used by the controller.

Host
The hostname of the device running the controller.

Key
The API key used by the controller.

CPU Load
A graph of CPU load on the device running the controller from the last hour.

Uptime
The amount of time the controller has been running.

Actions

At the bottom of the window are buttons to perform actions on the controller.

Sign Out
Logs the current user out of the controller and returns to the login screen.

Restart Controller
Restarts the controller software.

9.4.3 General Netgate Nexus GUI Usage

Window Controls

Certain functions in the Netgate Nexus GUI present modal dialog windows which have window controls.



Fig. 18: Modal Window Controls

These functions include:



Expand the current window so it fills the browser window.



Close the current window and return to the previous page.

Managing Lists of Items

Some areas in the Netgate Nexus GUI contain lists of items. For example, when configuring the controller itself or instances. Lists of items are managed similarly no matter where they are located, though some lists have more options than others.

Per-item Actions

List entries may contain icons to act on individual rows.



Create a new entry and add it to the list immediately below this row.



Create a new separator bar below this row.



Create a new item using the settings from this item as starting values.



Edit this list item.



Click and drag using this icon to move an item. Drop it in the new location.



View settings which control this list entry.

Selecting List Items

Some lists allow selecting one or more items to perform actions. Certain lists also have a “Select All” toggle to change the selection status of all entries.



This item is not selected. Click the icon to select the item.



This item is selected. Click the icon to deselect the item.



Some items in this list are selected, and some are deselected. Click the icon to select all items.

ToolBox List Actions



Fig. 19: List Management - ToolBox for Firewall Rules

When managing lists, the ToolBox may show one or more of the following icons.



Create a new item and add it to the list.



Reload data for this list from the controller or instance.



Add a new separator row.



Remove the selected item(s).



Toggle the state of the selected item(s). For example, disable or enable the item.



Copy the selected item(s) to another location. For example, copy firewall rules to another interface.



View help information for the page or items.



Open the tag manager.

9.5 Multi-Instance Management

A core feature of Netgate Nexus is multi-instance management (MIM). The Netgate Nexus controller is capable of managing multiple installations of pfSense Plus software from a single host. The Netgate Nexus controller includes a web-based graphical user interface (GUI) and an API.

Using the Netgate Nexus GUI MIM functionality, administrators can:

- See quick at-a-glance status of local and remote pfSense Plus software installations.
- Configure the local pfSense Plus software installation using a new single page application interface.
- Configure multiple remote pfSense Plus software installations paired with the controller using a new single page application interface.
- Open remote interactive consoles similar to an SSH connection.

9.5.1 Designing a Multi-Instance Management Configuration

The Netgate Nexus controller host has a critical role at the center of a Multi-Instance Management (MIM) configuration, so choosing the most appropriate place to host the controller is an important decision.

MIM Behavior Overview

Instances of pfSense Plus software are registered with the controller *using a manual process*. After registration, instances connect a special-purpose VPN back to the controller host which allows them to communicate privately and securely. All communication between the controller and instances is encrypted using this VPN. The only port on the controller which must be exposed to the Internet is the port for inbound VPN connections from registered instances.

Note: This VPN is **only** for use by the Netgate Nexus controller, it **does not** enable connectivity between networks behind the other instances and the controller.

Designate a Controller Host

Selecting a controller host typically involves a few factors:

- The controller works best with a static IP address, but can also advertise a hostname which can be used in combination with Dynamic DNS if necessary. Instances can have dynamic addresses.
- Instances must be able to make connections back to the controller.
- The controller host requires additional CPU and memory resources.
- The controller should reside in a location with stable power and connectivity.

Considering these factors, designate one host as the Netgate Nexus controller and the remaining devices will be instances.

9.5.2 Instance Registration

After setting up the Netgate Nexus controller as described in *Netgate Nexus Controller Setup*, it is possible to start registering instances of pfSense Plus software with the controller for multi-instance management (MIM).

This is a two-way process: Instances must be registered with the Netgate Nexus controller and the controller must be activated on instances.

Repeat the process in this document for each instance of pfSense Plus software to be registered with the controller.

Register Instance with Controller

The first task is to register the instance with the Netgate Nexus controller.

First, enable Netgate Nexus on the instance and copy the registration data:

- Open the pfSense software GUI on the instance
- Navigate to **System > Advanced**, **Netgate Nexus** tab
- Check **Enable**
- Click **Save**
- Click **Copy** to the right of the **Registration Data** field

Next, use that registration data to add the device to the controller

- Open the pfSense Plus GUI on the designated controller (*Accessing the Netgate Nexus GUI*)
- Click **Systems > Dashboard** in left menu
- Click  in the toolbox on the right side of the page to add a new item

---

Note: Depending on the controller configuration, this may first require a click on the  (hamburger menu) button to expand the right menu.

- Set **Create Method** to **Quick**

  This is the default option.
- Click in the **Device info JSON** text field
- Paste the clipboard contents
- Click **OK**

The device will now appear on the Netgate Nexus GUI dashboard, but with incomplete information since the registration is not yet complete.

Activate MIM on pfSense Plus Software Instance

To complete the registration, the controller must also be activated on the instance.

First, copy the activation data from the Netgate Nexus controller:

- Open the Netgate Nexus GUI on the designated controller (*Accessing the Netgate Nexus GUI*)
- Click  in the toolbox on right

---

Note: Depending on the controller configuration, this may first require a click on the  (hamburger menu) button to expand the right menu.

- Click the **JSON** tab
- Click  **Copy**
- Click  to close the window

Next, enter that activation data on the instance:

- Open the pfSense Plus GUI on the instance
- Navigate to **System > Advanced**, **Netgate Nexus** tab
- Click in the **Activation Data** field
- Paste the clipboard contents
- Click **Save**

Now the registration is complete, and the instance can be managed from the controller.

9.5.3 Managing pfSense Plus Software Instances

The web-based Multi-Instance Management (MIM) graphical user interface (GUI) enables administrators to manage individual instances of pfSense Plus in multiple ways.

Note: This document assumes the controller is already enabled (*Netgate Nexus Options in the pfSense Plus WebGUI*), configured (*Netgate Nexus Controller Setup*), and has instances registered (*Instance Registration*).

Before proceeding, open the MIM GUI as described in *Accessing the Netgate Nexus GUI* and log in, after which the browser will be on the *Netgate Nexus Dashboard*.

Remote Configuration

Click the  icon on an instance entry from the *Netgate Nexus Dashboard* to open the remote configuration view.

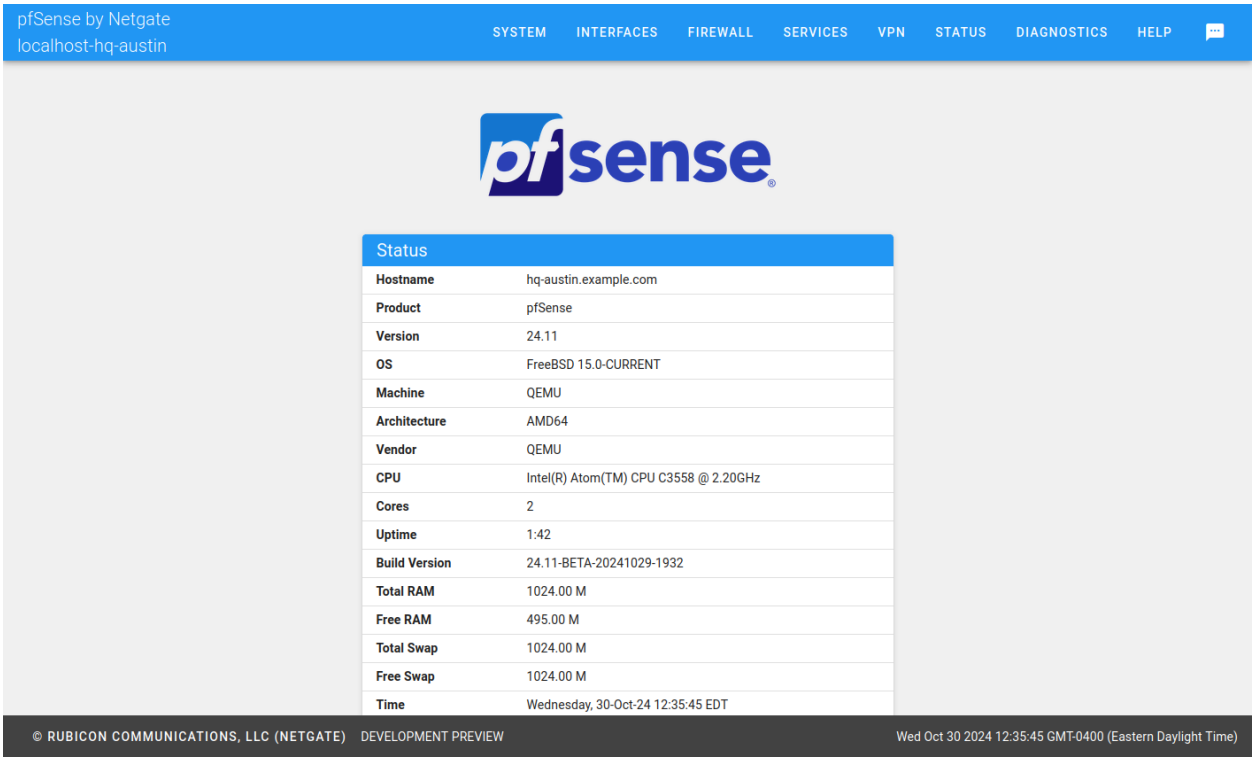


Fig. 20: Remote Configuration View

Managing individual instances of pfSense Plus software using the MIM GUI works similarly to managing those instances using their *pfSense software GUI* directly. The look and feel of the MIM GUI configuration screens are different, but functions in both GUIs are under the same menu locations and offer the same options.

Areas where the behavior of the MIM GUI differs from the pfSense software GUI are noted in this document.

Instance Status / pfSense Software Dashboard

When opening the configuration view for an instance, the MIM GUI presents a status page for the instance. Though this page shares some similarity with the pfSense Software GUI dashboard, the instance status page in the MIM GUI does not yet have the same functionality or widgets available.

Managing Interfaces

Managing interfaces via the MIM GUI is different from using the pfSense software GUI. Rather than having individual menu entries, interfaces are managed from the assignment list.

To edit interface settings:

- Open the remote device management view for an instance
- Navigate to **Interfaces > Assignments**
- Find the interface in list
- Click  at the end of its row to edit the interface settings

Remote Console

Click the  icon on an instance entry from the *Netgate Nexus Dashboard* to open the remote console.

The remote console appears in a *modal window* which has the same features and functionality as the *console menu on pfSense software* as if it were accessed directly via serial console, video console, or SSH.

CHAPTER

TEN

BACKUP AND RECOVERY

10.1 Making Backups in the GUI

Making a backup in the GUI is simple:

- Navigate to **Diagnostics > Backup & Restore**
- Set any desired options, or leave the options at their default values.
- Click **Download Configuration as XML** (Figure *GUI Backup*).

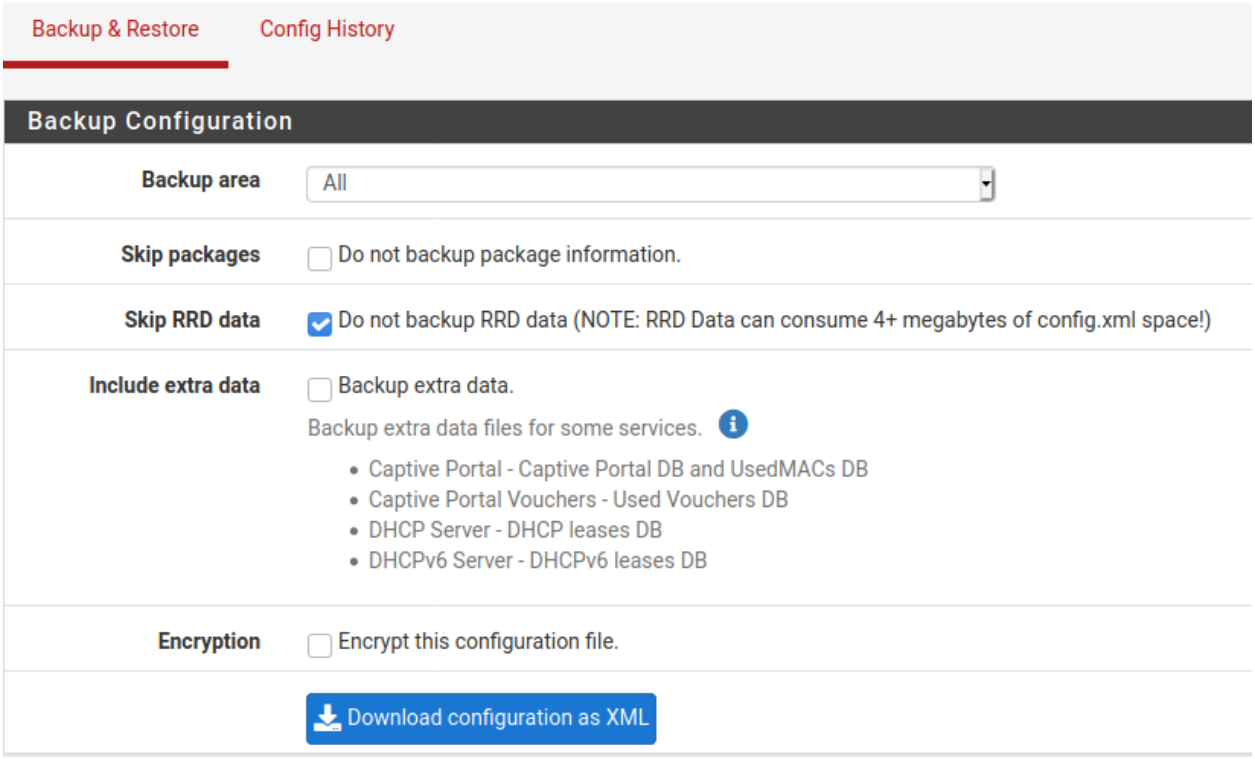


Fig. 1: GUI Backup

The web browser will then prompt to save the file somewhere on the PC being used to view the GUI. It will be named `config-<hostname>-<timestamp>.xml`, but that may be changed before saving the file.

10.1.1 Backup Options

When performing a backup, GUI options are available to control what is contained within the backup file.

Backup Area
Limits the backup contents to a single configuration area, rather than a complete configuration backup.

The default behavior is to include all areas in the backup.

Note: When restoring a configuration containing only a single area, the **Restore area** value must be set to match.

Skip Packages
Controls whether the backup will contain installation data and settings for packages. Omitting this data from a backup can be a useful way to quickly remove all traces of packages from a configuration when troubleshooting.

> **Warning:** After restoring a configuration without package data all packages must be reinstalled and reconfigured.

The default is *unchecked* so that all package data is included in the backup.

Skip RRD Data
Controls whether the backup will contain an exported copy of data used to generate *monitoring graphs*. When restoring a backup containing RRD data, the graph data is also restored.

The default is *checked* which omits the RRD data from the backup as it significantly increases the size of backup files.

Include Extra Data
Controls whether the backup file will include additional optional data. This includes Captive Portal databases and DHCP lease databases. These databases are volatile. While the data can be useful for transferring to new hosts or for frequent backups, it is not as useful for long-term backups.

The default is *unchecked* which omits this extra data from the backup as it can significantly increase the size of backup files.

Backup SSH Keys
Controls whether the backup file will include a copy of the SSH host keys.

Clients use these keys to uniquely identify the firewall, so preserving the keys when restoring makes it easier for clients to recognize the firewall after reinstalling or restoring to new hardware. Additionally, AutoConfigBackup uses the SSH host keys to identify the firewall when creating and restoring backups, so preserving the keys allows the firewall to maintain a consistent backup history after a reinstallation.

Encryption
Controls whether the backup file is encrypted before download.

When set, the GUI presents **Password** and confirmation fields, the contents of which are used by pfSense® software to encrypt the backup file with AES-256.

The default behavior is *unchecked* which creates clear text XML backup files.

10.2 Automatic Configuration Backup Service

Automatic Configuration Backup, also called AutoConfigBackup or ACB for short, is an encrypted remote backup service. AutoConfigBackup provides secure off-site backups for devices running pfSense® software without user intervention beyond the initial configuration.

Netgate® provides the AutoConfigBackup service free of charge for all users of pfSense software, both Plus and CE. This functionality is available as a core component of pfSense software included in the base installation.

This feature is located in the pfSense software GUI at **Services > Auto Config Backup**.

Tip: Read this entire section thoroughly before enabling the service.

10.2.1 How AutoConfigBackup Works

AutoConfigBackup performs several actions for each backup:

- AutoConfigBackup is triggered by configuration change (default), scheduled backup, or manual entry (*Backup Frequency*).
- AutoConfigBackup encrypts a backup copy of the pfSense® software configuration with *a user-supplied password*.
- AutoConfigBackup prepares to upload the backup by placing the encrypted backup and its metadata in a staging area.
- Once per minute, AutoConfigBackup checks for staged backups pending upload.
- AutoConfigBackup uploads each staged backup entry over HTTPS to the AutoConfigBackup service at Netgate® which stores it using *a device-specific key* to identify the owner of the backup.

Note: The AutoConfigBackup service retains the most recent **100** encrypted configuration backups for each device key. The server removes any backups beyond this number when AutoConfigBackup uploads new entries for a device key.

10.2.2 AutoConfigBackup Configuration

To configure the *Automatic Configuration Backup Service*, navigate to **Services > Auto Config Backup**, **Settings** tab in the pfSense® software GUI.

This document describes each option available on the **Settings** tab.

Enable ACB

When checked, AutoConfigBackup is active and will make automatic configuration backups as determined by the other settings on this page.

Backup Frequency

This option controls when AutoConfigBackup creates backups.

On Every Configuration Change
When selected, AutoConfigBackup will create a backup on every significant configuration change.

Note: Some minor configuration changes are safely ignored if they do not impact functionality.

On a Regular Schedule
Enables *Schedule* controls to create periodic backups at specific times instead of creating a backup on every change. This can be more efficient on systems with many frequent changes.

Tip: The default schedule creates a backup once per day. In most cases it should not be set more frequently than that, or at most a 2-4 times per day. If a device requires more frequent backups, backing up each change is likely a better practice.

Note: If the configuration has not changed since the previous scheduled backup time, AutoConfigBackup will not make a new backup.

Schedule

Controls the **Minute** of the hour, **Hours** of the day, **Day** of the month, **Month** of the year, and **Day of the week** on which backups are performed using standard cron format.

The value of **Minute** is randomized until the page is saved. The default value of **Hours** is `0`. The default for **Day**, **Month**, and **Day of Week** are all `*`. This results in a default schedule of a backup taken each day during the midnight hour at a random minute.

Note: This control is only visible when **Backup Frequency** is set to *On a Regular Schedule*.

Device Key

The AutoConfigBackup service requires a unique identifier for each device, this is called the **Device Key**. AutoConfigBackup uses the device key to identify entries belonging to a specific device when uploading or retrieving backups, similar to a username.

AutoConfigBackup displays the device key on each AutoConfigBackup tab for reference.

Warning: Treat this key as a secret!

Anyone who has this key can manipulate the backups for this key.

The device key value defaults to a randomly generated string for new configurations as well as upgraded configurations with AutoConfigBackup disabled. This randomized device key is stored in the configuration. Previous versions of AutoConfigBackup used the SHA256 hash of the SSH public key on the device to generate the device key, these older style keys are now called "legacy keys".

Note: Upgraded AutoConfigBackup configurations which are enabled and using a legacy key will continue using the legacy key until the key is changed manually or AutoConfigBackup is disabled.

Next to the **Device Key** reference field there is a  **Change Key** button. This button navigates to the *Change Device Key* page which allows administrators to generate or enter a new **Device Key**.

Warning: Keep a careful record of this Device Key!

Each AutoConfigBackup tab has a download icon to save a local copy of the device key (). Store this key in a safe and secure location.

If an administrator loses the **Device Key**, there is a chance Netgate TAC can help recover the key if AutoConfigBackup has a unique *Hint/Identifier*. If the hint is distinct, Netgate TAC *may* be able to use it to recover the device key. Do not count on this though!

Encryption Password

AutoConfigBackup encrypts the backup using the AES-256-CBC algorithm and this **Encryption Password** before it transmits the configuration to the AutoConfigBackup servers hosted by Netgate®.

Warning: This password is an encryption key, so the best practice is to use a long and complex string to ensure the backup contents are securely encrypted.

AutoConfigBackup **only** uses this password locally for encryption and decryption, it does not use or transmit the password as a login/credential.

When an administrator restores, views, or downloads an entry from the list of available remote backups, AutoConfigBackup fetches the remote backup entry and decrypts it with this password.

Danger: Keep a careful record of the encryption password! The encryption password is private and only known to the local device.

If the administrator loses the Encryption Password the backup contents cannot be recovered.

Danger: Changing this password will make AutoConfigBackup encrypt new backups using the new password, but old backups using a different password will no longer be readable by AutoConfigBackup.

Hint/Identifier

AutoConfigBackup sends the contents of the optional **Hint/Identifier** field as plain text metadata along with the encrypted configuration. This value is not visible to users, it is only visible to Netgate TAC.

If this hint is unique, it may allow Netgate TAC to locate a missing device key.

Manual Backup Limit

This setting controls the number of manual backups AutoConfigBackup will retain.

AutoConfigBackup can retain up to `50` manual backups, which it will not overwrite with automatic backups. Manual backups still count against the `100` backup limit. When the amount of manual backups exceeds this limit, AutoConfigBackup will remove older manual backups.

Descending Date Order

When set, AutoConfigBackup will sort the list of backup entries on the *Restore tab* in descending order by date (newest first) instead of the default order, which is oldest first.

10.2.3 Changing the AutoConfigBackup Device Key

The *Device Key* for the *Automatic Configuration Backup Service* can be changed at any time.

Changing the device key involves invoking the **Change Key** function from *AutoConfigBackup settings* and then completing the process using this *Change Device Key Page*.

> **Danger:** Changing the **Device Key** disconnects AutoConfigBackup from all previous backups stored using the old key.
>
> Administrators can still access the old backups so long as they know the old device key and encryption password. Securely store a backup copy of the current device key(s) before changing the device key.

Device Key Change Procedure

To change the *Device Key*:

- Open the pfSense® software GUI in a web browser
- Navigate to **Services > Auto Config Backup**, **Settings** tab
- Click  **Change Key**
- Use the  icons in each section to download and securely store copies of the **Current Device Key** and, if present, the **Legacy Device Key**
- Click  **Generate New Key** or manually enter a **New Device Key**
- Read the **Warning** section
- Check the confirmation box

- Click  **Update Key**
- Click **OK** on the confirmation dialog

For an explanation of each field on this page, refer to the next section, *Change Device Key Page*.

Change Device Key Page

The **Change Device Key** page allows administrators to change the AutoConfigBackup device key. It contains numerous warnings and explanations which guide users through the process.

> **Danger:** Changing the **Device Key** disconnects AutoConfigBackup from all previous backups stored using the old key.
>
> Administrators can still access the old backups so long as they know the old device key and encryption password. Securely store a backup copy of the current device key(s) before changing the device key.

Current Device Key

The **Current Device Key** section displays the current randomized device key stored in the device configuration.

The GUI will not display this section if AutoConfigBackup is enabled and using a legacy device key.

The section also includes a count of backup entries on the AutoConfigBackup service associated with the current device key.

Tip: Use the  icon in this section to download a copy of the current device key, then store it in a safe and secure location.

If there are **zero** hosted backups for the current device key, then it may not be necessary to store this key as there isn't anything left to access using this device key.

Legacy Device Key

The **Legacy Device Key** section displays the legacy style device key. This older key style is based on the SHA256 hash of the SSH public key on the device.

The GUI will not display this section if the device does not contain any SSH host keys. For example, on a fresh installation where SSH has never been enabled.

The section also includes a count of backup entries on the AutoConfigBackup service associated with the legacy device key.

Tip: Use the  icon in this section to download a copy of the legacy device key, then store it in a safe and secure location.

If there are **zero** hosted backups for the legacy device key, then it may not be necessary to store this key as there isn't anything left to access using this device key.

New Device Key

This field sets a new AutoConfigBackup device key. The key must be exactly `64` hexadecimal characters in length (`0` through `9`, `a` through `f`).

Use the  **Generate New Key** button next to the text field to generate a new randomized key in the correct format.

It is also possible to enter an existing device key in this field to make AutoConfigBackup use a device key from a previous installation.

> **Warning: Treat this key as a secret!**
>
> Anyone who has this key can manipulate the backups for this key.

Warning

This checkbox serves as a confirmation that the administrator has read the warnings and acknowledges the consequence that changing the device key will disconnect AutoConfigBackup from backups stored under the old device key.

After reading and understanding the warnings and saving the old keys, check the box in this section to enable the **Update Key** button.

Update Key

Clicking the  **Update Key** button will store the new device key in the configuration, replacing the current device key in the process.

Note: This button is disabled by default. Check the *Warning* checkbox to enable this button.

After clicking the button the GUI presents one more confirmation box using a JavaScript alert. Click **OK** to the confirmation dialog to complete the process.

If the **New Device Key** is not valid, the page will display an error without taking any action.

10.2.4 Restoring from AutoConfigBackup

Restoring a backup entry from the *Automatic Configuration Backup Service* is a straightforward process, but there are several other actions and techniques which may be useful as well.

The AutoConfigBackup **Restore** tab contains a list of backups and controls to act on those backups as well as showing backups for other devices.

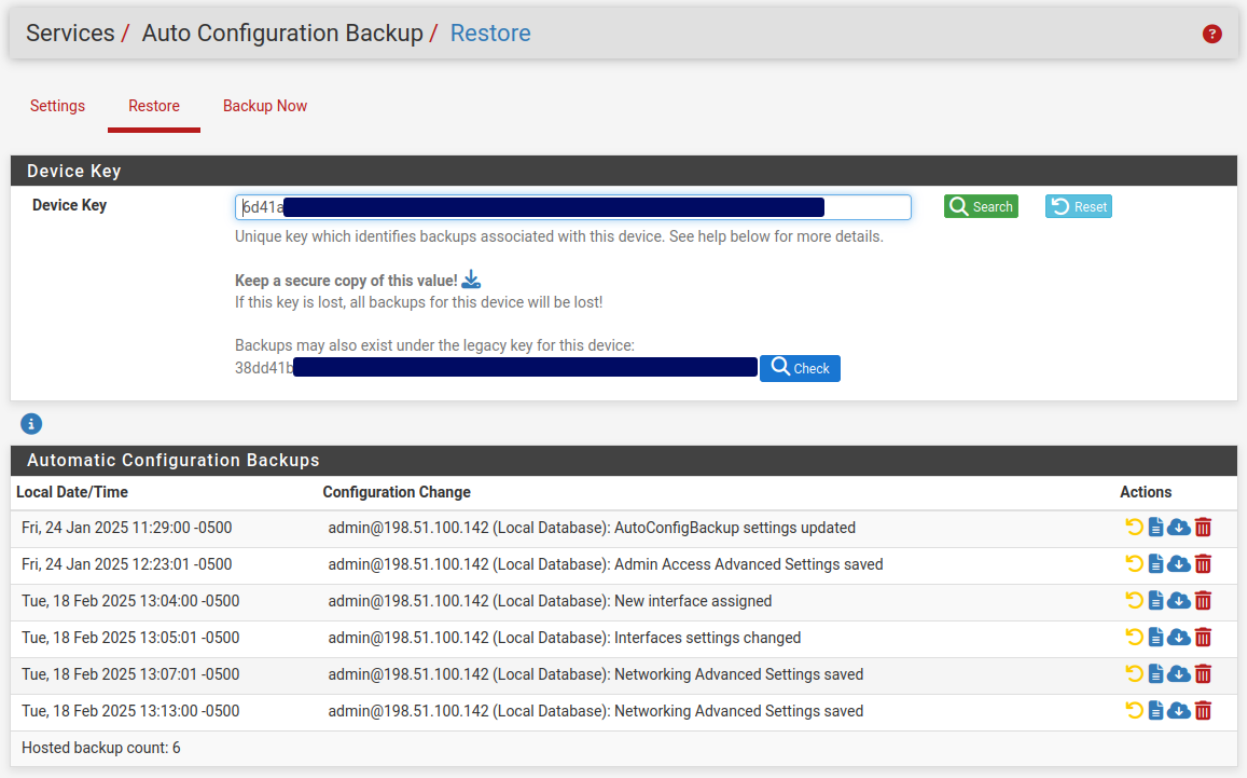


Fig. 2: AutoConfigBackup **Restore** tab, with Device Key and backup list visible

Restoring an AutoConfigBackup Entry

This procedure is a brief overview of how to restore a remote configuration backup from the AutoConfigBackup service using the GUI:

- Navigate to **Services > Auto Config Backup**
- Click the **Restore** tab at the top
- Locate the desired entry in the *Backup List*
- Click  to the right of the configuration row
- Click OK to confirm the restore action

AutoConfigBackup will download the configuration specified from the AutoConfigBackup service, decrypt it with the *Encryption Password*, and restore the configuration.

> **Warning:** By default, the restore process **will not** initiate a reboot. Depending on the configuration items restored, a reboot may not be necessary. For example, firewall and NAT rules are automatically reloaded after restoring a configuration, but interface configurations are not.

After restoring, the GUI presents a prompt offering to reboot. If the restored configuration changes anything other than the NAT and firewall rules, the device requires a reboot to fully activate the changes:

- Click **Yes** in the save message, which navigates to the *Reboot page*

- Click  **Submit** to reboot the device

For more details about other capabilities of the **Restore** tab, continue reading this document.

Device Key

This section displays the *Device Key* AutoConfigBackup used to retrieve the backups on the page. This could be a randomized device key stored in the configuration or a *legacy device key*.

Tip: Click  to save a copy of this device key.

To see backups stored using another key, paste the key into the **Device Key** field and click  **Search**.

When viewing another key, use the  **Reset** button to return the page to the original device key.

The lower section contains a  **Check** button to easily swap to another known key. This button behavior changes depending on the current view:

- When viewing a current device key, if the device contains a legacy device key, this button will show backups stored under the legacy key.
- When viewing backups from the legacy key, it changes to a button which displays backups from the current device key.
- If there is no alternate key, the button does not display.

Backup List

The section of the page titled **Automatic Configuration Backups** displays a list of remote backup entries stored on the AutoConfigBackup service for a given *Device Key*.

The list contains the following columns for each backup entry:

Local Date/Time
: The date and time the backup was created, with the time zone adjusted to be local to this device.

  Note: The AutoConfigBackup service stores entries with a UTC timestamp.

Configuration Change
: A brief description of the configuration change. This is typically an automatic string containing who made a change and what change they made, but may also contain a *manual backup reason*.

Actions
: A list of actions which can be taken on a backup entry:

  -  : Restore this AutoConfigBackup backup entry. Prompts for confirmation. Read *Restoring an AutoConfigBackup Entry* for additional details and warnings.

- : *Views the details of the AutoConfigBackup backup entry*, including the configuration itself.
- : Downloads the AutoConfigBackup backup entry as an XML file.
- : Deletes the AutoConfigBackup backup entry. Prompts for confirmation.

At the bottom of the list are two lines:

- A current count of hosted backups for this device key on the AutoConfigBackup service.
- A count of configuration backups staged for upload which have not yet been processed. The page only includes this line when there are staged entries waiting. Staged configuration backup entries are uploaded once per minute. See *How AutoConfigBackup Works* for details.

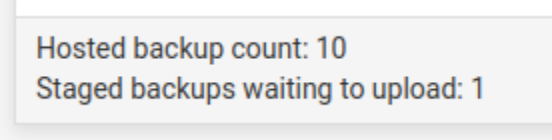


Fig. 3: Bottom of AutoConfigBackup **Restore** tab backup list showing backup count and number of backups staged waiting to upload.

Viewing AutoConfigBackup Entry Details

Clicking the icon for an entry on the *Backup List* opens up the **Revision** tab for the AutoConfigBackup backup entry.

This page contains the following fields:

Service Date/Time
The date and time in UTC when this backup revision was stored on the AutoConfigBackup service.

Local Date/Time
The same timestamp converted to the local time zone of the device.

Revision Reason
A brief description of the configuration change. This is typically an automatic string containing who made a change and what change they made, but may also contain a *manual backup reason*.

SHA256 Summary
A SHA256 hash of the configuration data, used to confirm that the contents are correct.

Encrypted config.xml
The encrypted blob containing the configuration data. This can be copied and pasted and decrypted manually as described in *Encrypted Configuration files*.

Decrypted config.xml
The decrypted contents of the configuration. This can be copied and pasted to a file and saved, similar to downloading the entry.

Tip: The  **Download this revision** button accomplishes this much easier, but some users may wish to copy/paste the contents for other purposes.

Restore Button

Clicking the  **Restore** button restores this AutoConfigBackup backup entry. Prompts for confirmation. Read *Restoring an AutoConfigBackup Entry* for additional details and warnings.

Download this revision Button

Clicking the  **Download this revision** button Downloads the AutoConfigBackup backup entry as an XML file.

Bare Metal Restoration from AutoConfigBackup

If the disk in the firewall fails or if the device key changes, the AutoConfigBackup service can restore a backup from the previous installation as long as the *Device Key* and the *Encryption Password* of the previous installation are both known.

- Install pfSense® software on the device
- Connect to the GUI and login
- Click the logo at the top left to skip the wizard

  If the WAN requires special configuration, use the wizard or configure it manually.
- Navigate to **Services > Auto Config Backup**, **Settings** tab
- Set the **Encryption Password** to match the previous installation
- Navigate to the **Restore** tab
- Paste the old device key into the **Device Key** field
- Click the  **Search** button

  This temporarily displays a list of backups for an alternate *Device Key*.
- Locate the desired entry in the *Backup List*
- Click  to the right of the configuration row
- Click **OK** to confirm the restore action
- Click **Yes** in the save message, which navigates to the *Reboot page*
- Click  **Submit** to reboot the device

When the device boots back up it will be running the restored configuration.

10.2.5 Manual AutoConfigBackup Entry

Manual backups using the *Automatic Configuration Backup Service* are a best practice before and after an upgrade or a series of significant changes.

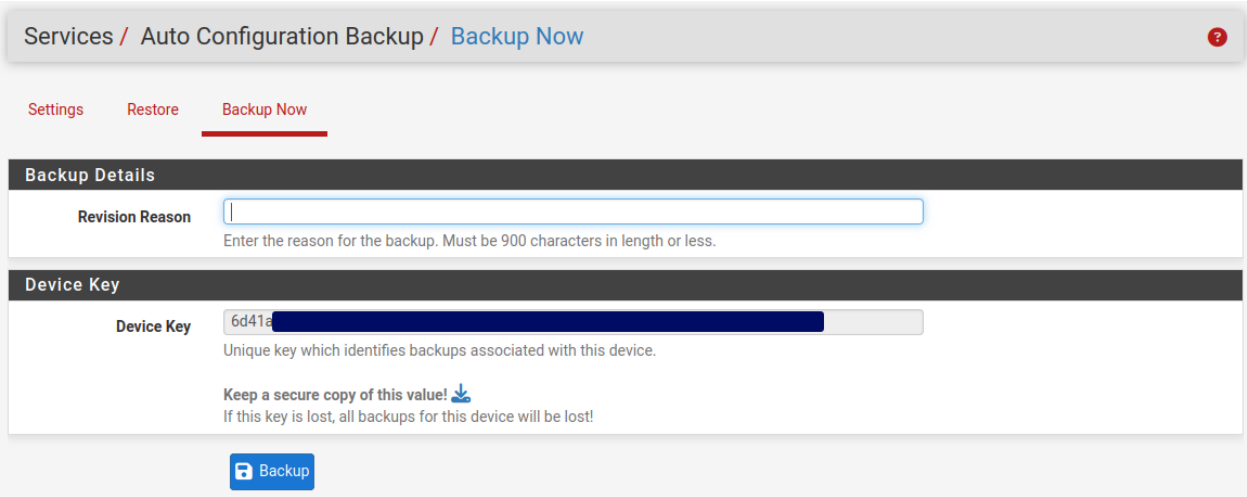


Fig. 4: **Backup Now** tab which creates manual AutoConfigBackup backup entries

This page creates manual AutoConfigBackup entries using a custom backup description provided by the administrator. The custom backup description makes it easy to identify and restore the manual backup entry.

Note: This page is not available when AutoConfigBackup is in a disabled state.

When AutoConfigBackup creates backups on every configuration change, a series of changes can make it difficult to know where the process started or ended. When using scheduled backups, the time of the scheduled backup may not have happened at a key point in the process. In either case, having manual entries at the start and end provides important remote backup entries which otherwise may not exist.

To create a manual backup of the configuration using AutoConfigBackup:

- Navigate to **Services > Auto Config Backup**
- Click the **Backup Now** tab at the top
- Enter a **Revision Reason**
- Confirm the **Device Key** is correct as it is the key AutoConfigBackup will use when uploading the backup
- Click **Backup**
- Wait *at least one minute* before *checking if the backup succeeded*

10.2.6 AutoConfigBackup Testing and Status

After *configuring* the *Automatic Configuration Backup Service*, the best practice is to perform tests to confirm backups are working properly.

Testing AutoConfigBackup

Testing backups is a vital part of every backup strategy. Without testing the entire backup and restore process, administrators can never know if a backup is viable. As such, it is crucial to test creating backups and restoring backups when first configuring AutoConfigBackup and periodically after.

Danger: Do not skip backup testing!

Without knowing if the entire backup and restore process works, any problems in the process may not be discovered until it is too late during a critical moment.

The only good backup is a tested backup!

Test Manual Backup

The following steps cover the general tasks involved in testing a *manual backup*. For specifics on each step, follow the cross-reference links.

- *Configure and enable* AutoConfigBackup
- Create a *Manual AutoConfigBackup Entry* entry
- Wait *at least one minute*
- *Check if the backup succeeded*
- Make a change, such as editing and saving a firewall or NAT rule
- *Restore the manual backup* and reboot
- Check if the change made after the manual backup is present

The restored configuration should not contain the change made after the manual backup.

Test Configuration Change Backup

The following steps cover the general tasks involved in testing AutoConfigBackup when it is configured to make a new backup on each configuration change. For specifics on each step, follow the cross-reference links.

- *Configure and enable* AutoConfigBackup using a **Backup Frequency** of *Automatically backup on every configuration change*
- Make a change to trigger a configuration backup, such as editing and saving a firewall or NAT rule
- Wait *at least one minute*
- *Check if the backup succeeded*
- Make a second change, such as editing and saving a **different** firewall or NAT rule
- *Restore the first backup* and reboot
- Check if the second change is present

The restored configuration should not contain the second change.

Test Scheduled Backups

The following steps cover the general tasks involved in testing AutoConfigBackup when it is configured to create backups on a schedule. For specifics on each step, follow the cross-reference links.

- *Configure and enable* AutoConfigBackup using a **Backup Frequency** of *Automatically backup on a regular schedule* with an appropriate schedule

  Tip: For ease of testing, configure the schedule to back up every 5 or 10 minutes, then configure a more reasonable schedule after testing, such as once per day.

- Make a change to trigger a configuration backup, such as editing and saving a firewall or NAT rule
- Wait until after the scheduled backup time passes, plus *at least one additional minute*
- *Check if the backup succeeded*
- Make a second change, such as editing and saving a **different** firewall or NAT rule
- *Restore the scheduled backup* and reboot
- Check if the second change is present

The restored configuration should not contain the second change.

Check Restore Tab

Administrators can check the status of an AutoConfigBackup run by reviewing the *Backup List* on the *Restore tab*.

AutoConfigBackup fetches its list of backup entries from the AutoConfigBackup service. If the backup is in the list on the **Restore** tab, the backup was successful and is present on the server.

Depending on the *Backup Frequency*, a backup may not appear right away. For example, a scheduled backup won't appear until after the scheduled time passes. For manual backups or backups on each configuration change, new backups may not appear for *at least one minute* as the uploads are queued and then processed out of that queue.

If there are backups pending upload, the backup list will contain a note with a count of staged backup entries.

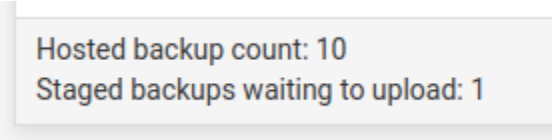


Fig. 5: Bottom of AutoConfigBackup **Restore** tab backup list showing backup count and number of backups staged waiting to upload.

If a backup fails, AutoConfigBackup logs an alert and displays a notice in the GUI.

View and/or download a backup then confirm it has the expected content.

AutoConfigBackup Logs

AutoConfigBackup creates log entries which describe its successes and failures. These can be found in the **System/General** sub-tab of the *System Logs*.

This is a sample of log entries created when creating a *Manual AutoConfigBackup Entry*:

```
/services_acb_backup.php: Staging AutoConfigBackup encrypted configuration backup for␣
↪deferred
          upload to https://acb.netgate.com
/usr/local/sbin/acbupload.php: Starting upload of staged AutoConfigBackup encrypted␣
↪configuration
          backups to https://acb.netgate.com
/usr/local/sbin/acbupload.php: Completed AutoConfigBackup encrypted configuration backup␣
↪upload
          to https://acb.netgate.com (success)
```

Check the logs when creating backups to ensure there are no errors which were not also printed in notices in the GUI.

Some AutoConfigBackup logs are informational and do not indicate a problem. For example, AutoConfigBackup skips certain redundant or frequent and not critical configuration changes. When AutoConfigBackup ignores a change, it logs the following message:

```
Skipping staging AutoConfigBackup entry for ignored reason: <reason>
```

10.3 ZFS Boot Environments (Plus Only)

ZFS Boot Environments make upgrades and major changes safer by taking snapshots of key filesystem areas, allowing the firewall to be rolled back to an earlier known good state if the user encounters problems with an upgrade, configuration change, or other potentially problematic situation.

The upgrade process automatically creates a new ZFS Boot Environment and administrators can create them manually as well. Administrators can then select a previous ZFS Boot Environments using the GUI or even the bootloader menu which makes quickly recovering from unforeseen issues a breeze.

Warning: ZFS Boot Environments are available only in pfSense Plus software version 22.05 and later. They are not available on pfSense® CE software.

10.3.1 How Boot Environments Work

A ZFS Boot Environment is a snapshot of the filesystem at a specific point in time, plus a clone of that snapshot. Snapshots are read only views of the filesystem at a given point, whereas clones are read/write.

Each snapshot and clone consumes some disk space but the exact amount varies based on how much the current contents of the filesystem have diverged from the contents when the entries were created.

Note: For most users tracking periodic updates or creating occasional ZFS boot environments the disk usage will be moderate over time. Users tracking development snapshots with frequent updates may see much larger amounts of space consumed by ZFS Boot Environments from snapshots. See *Boot Environment Disk Space Usage* for details.

When an administrator triggers the upgrade process the firewall creates a new ZFS Boot Environment before the upgrade begins. This preserves the current state of the firewall as it was before the upgrade.

What happens next has changed over time. From pfSense Plus software version 22.05 until 23.09.1, the upgrade process then activates the new ZFS Boot Environment so that when the upgrade proceeds and reboots, it reboots into the new environment to complete the upgrade.

Starting with pfSense Plus software version 24.03, this changed to a more efficient and robust procedure: The upgrade process creates a new Boot Environment and performs the upgrade inside that entry before rebooting. It makes sure the upgrade succeeded and then reboots into the newly upgraded environment. It detects any errors during boot and if there is a problem it can automatically roll back to the previous Boot Environment.

Either way, if there is a problem, an administrator can manually activate a pre-upgrade ZFS Boot Environment and reboot the firewall, then it will return to its state before the upgrade happened.

See also:

For an overview of the upgrade process, see *Upgrade Process Overview*.

10.3.2 Boot Environment Requirements

- pfSense Plus software version 22.05 or later
- The firewall must be using ZFS

Note: If the firewall is using UFS, it must be reinstalled with ZFS.

- ZFS requires 64-bit hardware (amd64, arm64)
- Certain ZFS dataset layout changes may require a fresh install, though many existing ZFS installations will work

10.3.3 Managing Boot Environments in the GUI

The GUI page to manage ZFS Boot Environments is **System > Boot Environments**.

Note: If the **Boot Environment** menu entry is missing, the installation is not compatible with ZFS Boot Environments.

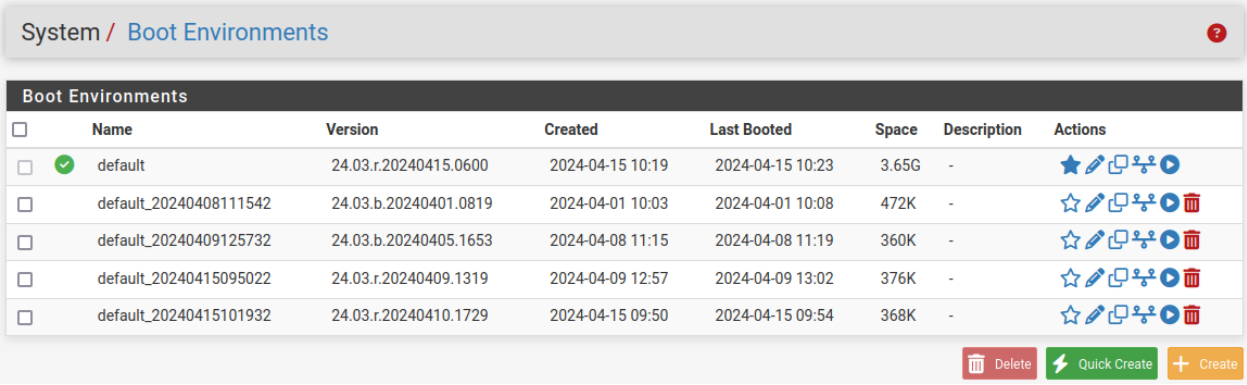


Fig. 6: ZFS Boot Environment list in the GUI

The Boot Environments page lists all existing ZFS Boot Environments with the following fields, as shown in *ZFS Boot Environment list in the GUI*:

Checkbox
Selection checkbox for batch actions which affect multiple Boot Environment entries.

Name
The name of the ZFS Boot Environment.

Automatic entries, such as those created by the upgrade process, are prefixed by `auto-` and include the timestamp at which they were created.

Base Version
The version of pfSense® software contained within the ZFS Boot Environment.

Created
The time at which the ZFS Boot Environment was created.

Last Booted
The time at which the firewall last booted into the ZFS Boot Environment.

Space
The amount of disk space consumed by the ZFS Boot Environment.

Description
The longer text description of the ZFS Boot Environment.

Actions
Actions the administrator can take on the ZFS Boot Environment.

-  : Indicates the ZFS Boot Environment the firewall will use for the next boot
-  : Persistently activate the entry as the next ZFS Boot Environment
-  : Edit the ZFS Boot Environment
-  : Clone the ZFS Boot Environment
-  : View the configuration history within this Boot Environment (*Configuration History*)
-  : Temporarily activate the ZFS Boot Environment one time and reboot

  There is an additional confirmation prompt to reboot after selecting this option.

-  : Delete the ZFS Boot Environment

Creating a new Boot Environment

Administrators can create new ZFS Boot Environments in several different ways.

> **Warning:** While boot environments are helpful, they do not remove the need for off-device backups. Take separate *configuration backups* before starting any potentially disruptive set of changes, including upgrades.

Automatic During Upgrade

By default, the firewall automatically creates a new ZFS Boot Environment before performing an upgrade. The upgrade is performed inside the new Boot Environment and the firewall automatically reboots and verifies that the Boot Environment is working properly. This behavior can be altered using the options under *Boot Environments*.

Quick Create

Clicking **Quick Create** from the ZFS Boot Environment list will clone the current default ZFS Boot Environment. The resulting entry will be named `quick-` followed by the current timestamp.

Create / Clone

Clicking **Create** from the ZFS Boot Environment list opens a form to create a new ZFS Boot Environment with custom options, including:

Name
: Short name to briefly indicate purpose, must only contain characters from the set `a-z`, `A-Z`, `0-9` and `_`.

Clone From
: The existing ZFS Boot Environment to use as the basis for this new entry.

Description
: A longer description for the ZFS Boot Environment without formatting restrictions.

Click **Save** to create the new ZFS Boot Environment.

The  entry from the ZFS Boot Environment list works identically but it pre-selects the chosen entry in the **Clone From** field.

Editing an existing Boot Environment

Clicking  on the row for a ZFS Boot Environment opens a form to edit the **Name** and **Description** of the entry. The clone source cannot be changed after the entry has been created.

Selecting Boot Environments in the GUI

There are multiple ways in the GUI to select which ZFS Boot Environment the firewall will use next.

From the ZFS Boot Environment at **System > Boot Environments** there are two methods:

- Click  to select the ZFS Boot Environment persistently
- Click  to select the ZFS Boot Environment for a single boot only and reboot. This is not persistent and the next boot after will return to the default.

From **Diagnostics > Reboot**, select a **Boot Environment** from the list and reboot. This is not persistent and the next boot after will return to the default.

Removing Boot Environments

ZFS Boot Environments can be removed individually or in batches.

Note: The current Boot Environment cannot be removed!

To remove individual ZFS Boot Environments:

- Navigate to **System > Boot Environments**
- Locate the entry to remove in the list
- Click the  icon at the end of the row
- Click **OK** on the confirmation dialog to remove the Boot Environment

To remove multiple Boot Environments:

- Navigate to **System > Boot Environments**
- Select each entry to remove by clicking the checkbox to the left of each entry or by clicking anywhere in the row *except* the icons in the **Actions** column.

  Note: To quickly select all entries except the current Boot Environment, use the checkbox in the header row of the list.

- Click the  **Delete** button below the list
- Click **OK** on the confirmation dialog to remove all the selected Boot Environment entries.

10.3.4 Selecting Boot Environments in the Loader Menu

At boot, pfSense® software briefly displays the loader menu with a logo and several options to control the boot behavior.

This loader menu will contain an option for ZFS **Boot Environments**, typically option `8` but may vary depending on the platform.

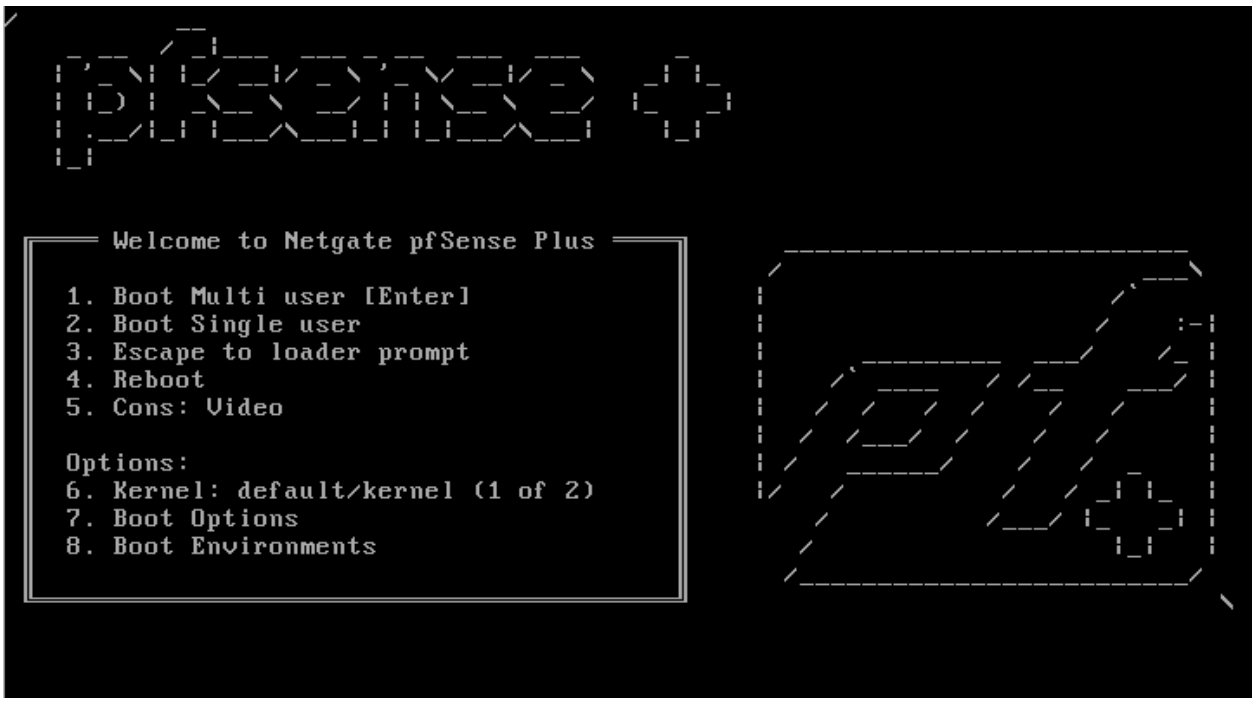


Fig. 7: Loader Menu - Enter the number for the Boot Environments option

Press the option for **Boot Environments** and the loader will display a new menu with ZFS Boot Environment options.

From this menu:

- Press option `2` to cycle through all available boot environments. Stop when the desired ZFS Boot Environment name is shown.
- Press option `3` to change the `bootfs` location if it is not correct

  This is unnecessary in the vast majority of cases as it likely only has one option.
- Press the `Enter` key to boot the selected Boot Environment or press `1` to return to the previous menu and change other options.

Note: This change is not persistent, and the next boot after will return to the default ZFS Boot Environment. To make this change persist, select the entry in the GUI using  .

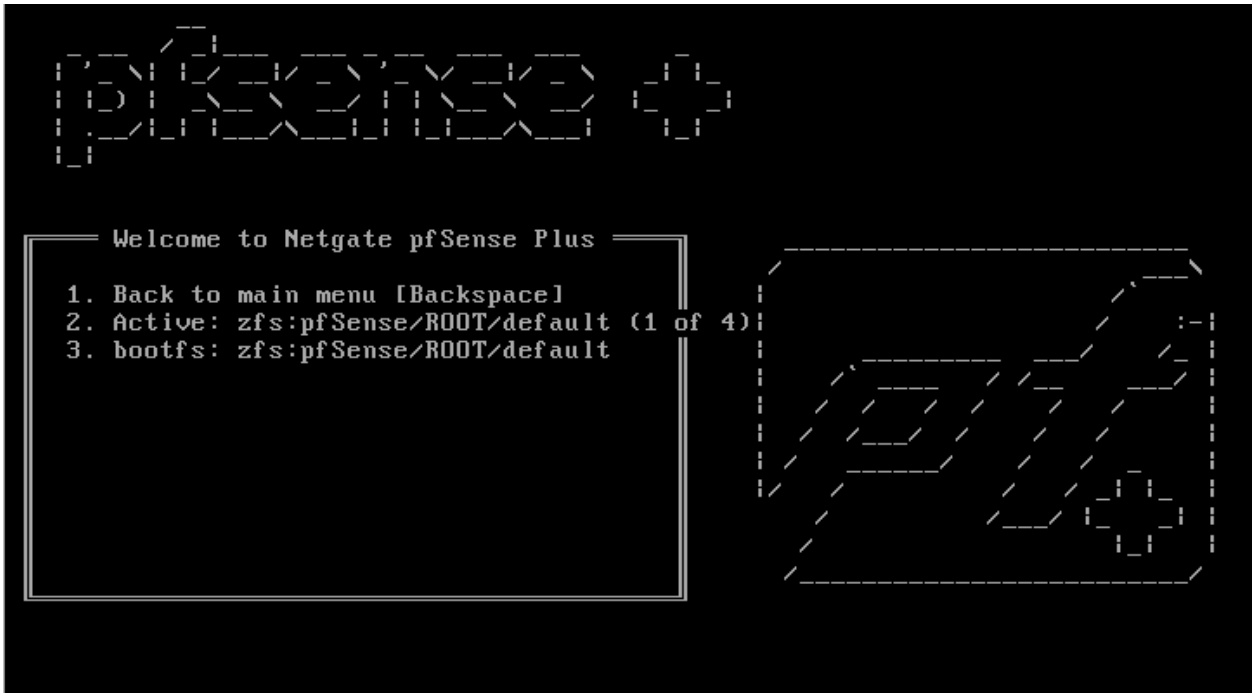


Fig. 8: Boot Environment Selection Menu

10.3.5 Boot Environment Status

The System Information widget on the Dashboard contains a **Boot Environment** section which prints the current ZFS Boot Environment and what the next ZFS Boot Environment will be.

Note: If the **Boot Environment** section of the widget is missing, the installation is not compatible with ZFS Boot Environments.

On **System > Boot Environments** the list of environments has an icon at the start of the row indicating the active and next ZFS Boot Environment.

-  : The firewall booted from this entry.
-  : The firewall will boot from this entry next.

  If this icon is not present the firewall will boot from the entry indicated by  .

10.3.6 Boot Environment Disk Space Usage

ZFS Boot Environment snapshots consume an increasing amount of disk space over time as the contents of the disk diverge from when it was created compared to the current state of the disk.

A ZFS Boot Environment snapshot taken before upgrading to a new version of pfSense® Plus software can consume several gigabytes of space as those updates will rewrite the entire base system and all the other components including packages as they are all reinstalled. Updating between development snapshots will cause a ZFS Boot Environment to consume about 500MB of disk space, give or take, based on what changed in the snapshot.

Warning: Frequent upgrades between development snapshots can cause ZFS Boot Environments to consume a lot of disk space!

The operating system reflects this usage as a change in the capacity of the disk. The size of a disk will appear to decrease proportionate to the snapshot usage, and this change is reflected on the dashboard **Disks** widget and in utilities such as `df`.

Removing older ZFS Boot Environments that are no longer necessary will free the space and make it available again. While the system will attempt to clean up older automatically created ZFS Boot Environments, ultimately it is up to the administrator to decide which ZFS Boot Environments are necessary.

Examples

The following are examples of space usage for numerous ZFS Boot Environments.

GUI

This figure shows the Dashboard **Disks** and **ZFS** widgets on a firewall with a 12GB disk and 12 ZFS Boot Environments from snapshot upgrades.

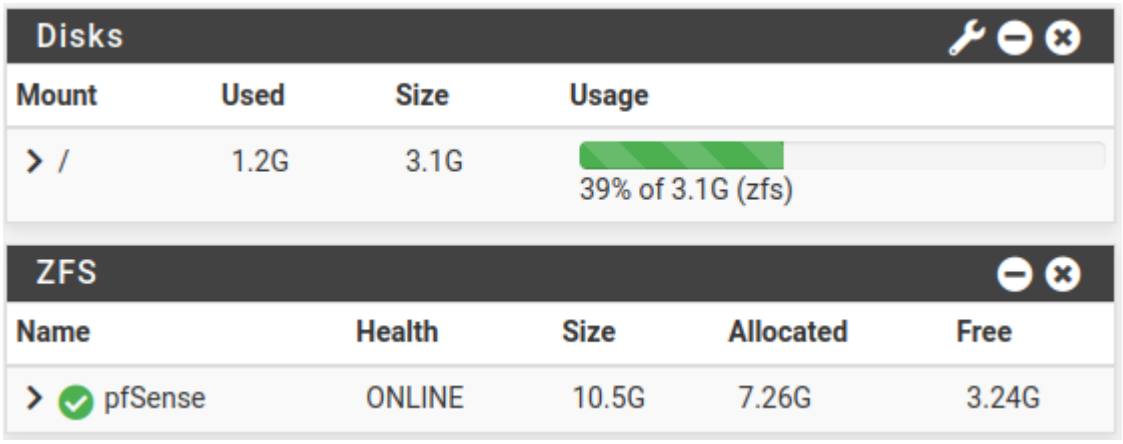


Fig. 9: Dashboard Disk Usage with 12 Boot Environments

Note that the disk size is listed as being only about 3GB when it should be significantly larger.

The next figure is the same system with the older Boot Environments removed so that only the default and one previous entry remain:

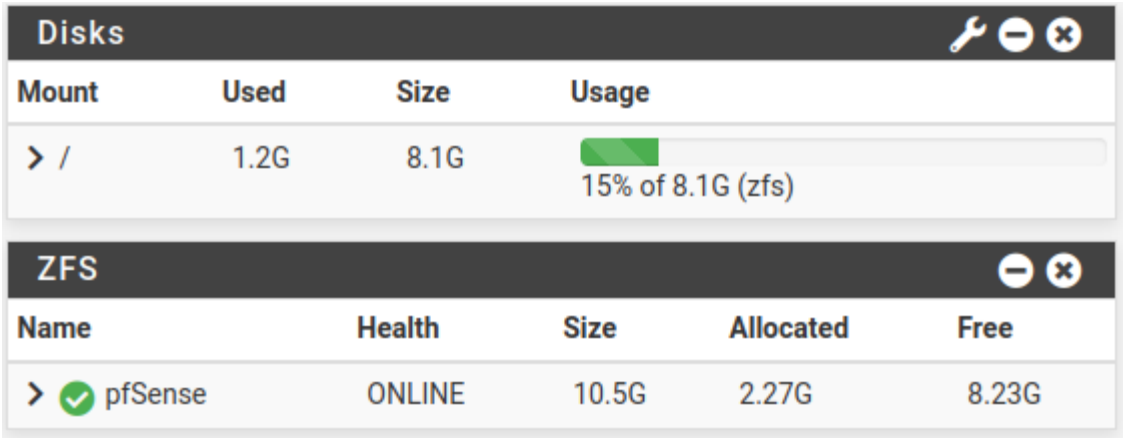


Fig. 10: Dashboard Disk Usage with 1 Boot Environment

Shell

Similar to the above example, this is the same firewall but with the disk usage checked at the shell instead of the GUI.

With 12 Boot Environments:

```
: df -h /
Filesystem              Size    Used   Avail Capacity  Mounted on
pfSense/ROOT/default    3.1G    1.2G    1.9G    39%    /

: zfs list /
NAME                   USED  AVAIL  REFER  MOUNTPOINT
pfSense/ROOT/default  6.93G  1.90G  1.20G  /
```

With the default plus one automatic Boot Environment:

```
: df -h /
Filesystem              Size    Used   Avail Capacity  Mounted on
pfSense/ROOT/default    8.1G    1.2G    6.9G    15%    /

: zfs list /
NAME                   USED  AVAIL  REFER  MOUNTPOINT
pfSense/ROOT/default  1.96G  6.89G  1.20G  /
```

10.3.7 Boot Environment Tips & Tricks

Reboot to Roll Back

- Create a new ZFS Boot Environment before making potentially disruptive changes to the firewall. This represents the current known-good state of the firewall.

> **Warning:** While boot environments are helpful, they do not remove the need for off-device backups. Take separate *configuration backups* before starting any potentially disruptive set of changes.

- Activate the new ZFS Boot Environment persistently with 
- Proceed to make the changes and monitor the firewall state.

If the changes caused a problem:

- Reboot and the firewall will restart from the ZFS Boot Environment with the known-good state.

If the changes are OK:

- Activate the default ZFS Boot Environment to continue using the new changes on future reboots.

10.4 Alternate Remote Backup Techniques

The easiest method to make secure and encrypted remote backups of the pfSense® software configuration is the free *Automatic Configuration Backup Service*. Rest easy knowing it is taking care of handling remote backups automatically without needing to worry. Sit back, have a cup of coffee, and read on for alternate techniques.

The other techniques in this document perform backups remotely, but each method has its own security issues which may rule out their use. For starters, several of these techniques do not encrypt the configuration, which may contain sensitive information. This can result in the raw configuration being transmitted over an unencrypted, untrusted link. If one of these techniques must be used, it is best to do so from a non-WAN link (LAN, DMZ, etc.) or across a VPN. Access to the storage media holding the backup must also be controlled, if not encrypted.

10.4.1 Pull

Pulling the configuration means to use a remote client to "pull" the configuration off of the firewall. The methods in this section accomplish the same goal using different utilities.

Pull with wget

The `wget` utility can retrieve the configuration from a remote firewall. This process can be scripted with `cron` or by other means to automate the process.

Warning: Even when using HTTPS, this is not a truly secure transport mode since certificate checking is disabled to accommodate self-signed certificates, enabling man-in-the-middle attacks. When running backups with `wget` across untrusted networks, use HTTPS with a certificate that can be verified by `wget`.

The `wget` command must be split into multiple steps to handle the login procedure and backup download while also accounting for CSRF verification.

For a firewall running HTTPS with a self-signed certificate, the commands are as follows:

- Fetch the login form and save the cookies and CSRF token:

```
$ wget -qO- --keep-session-cookies \
  --save-cookies cookies.txt \
  --no-check-certificate \
  https://192.168.1.1/diag_backup.php \
  | grep "name='__csrf_magic'" \
  | sed 's/.*value="\(.*\)".*/\1/' > csrf.txt
```

- Submit the login form along with the first CSRF token and save the second CSRF token (can't reuse the same file) – now the script is logged in and can take action:

```
$ wget -qO- --keep-session-cookies --load-cookies cookies.txt \
  --save-cookies cookies.txt --no-check-certificate \
  --post-data "login=Login&usernamefld=admin&passwordfld=pfsense&__csrf_magic=$(cat␣
↪csrf.txt)" \
  https://192.168.1.1/diag_backup.php \
  | grep "name='__csrf_magic'" \
  | sed 's/.*value="\(.*\)".*/\1/' > csrf2.txt
```

- Submit the download form along with the second CSRF token to save a copy of `config.xml`:

```
$ wget --keep-session-cookies --load-cookies cookies.txt --no-check-certificate \
  --post-data "download=download&donotbackuprrd=yes&__csrf_magic=$(head -n 1 csrf2.
↪txt)" \
  https://192.168.1.1/diag_backup.php -O config-router-`date +%Y%m%d%H%M%S`.xml
```

Note: The behavior of variable expansion and other aspects of the commands may vary by shell. This example uses `bash` for the client shell.

Replace the username and password with the credentials for the firewall, and the IP address is whichever IP address is reachable from the client performing the backup, and using HTTP or HTTPS to match the firewall GUI.

There are additional parameters which can control the contents of the backup in several ways:

- To back up the RRD files, remove the `&donotbackuprrd=yes` parameter from the post data string on the last command.
- To include extra data such as DHCP leases and captive portal databases, add `&backupdata=yes` to the post data string on the last command.
- To include the SSH keys for the firewall, add `&backupssh=yes` to the post data string on the last command.

The client performing the backup will also need access to the GUI, so adjust the firewall rules accordingly. Performing this type of backup over an Internet-connected WAN is not secure. At a minimum, use HTTPS and restrict access to the GUI to a trusted set of public IP addresses. A better practice is to do this locally or over a VPN.

Using cURL

The same task can be accomplished using cURL instead of `wget`:

- Fetch the login form and save the cookies and CSRF token:

```
$ curl -L -k --cookie-jar cookies.txt \
    https://192.168.1.1/ \
    | grep "name='__csrf_magic'" \
    | sed 's/.*value="\(.*\)".*/\1/' > csrf.txt
```

- Submit the login form to complete the login procedure:

```
$ curl -L -k --cookie cookies.txt --cookie-jar cookies.txt \
    --data-urlencode "login=Login" \
    --data-urlencode "usernamefld=admin" \
    --data-urlencode "passwordfld=pfsense" \
```

(continues on next page)

(continued from previous page)

```
    --data-urlencode "__csrf_magic=$(cat csrf.txt)" \
    https://192.168.1.1/ > /dev/null
```

Now the script is logged in and can perform actions!

- Fetch the target page to obtain a new CSRF token:

```
$ curl -L -k --cookie cookies.txt --cookie-jar cookies.txt \
    https://192.168.1.1/diag_backup.php  \
    | grep "name='__csrf_magic'"   \
    | sed 's/.*value="\(.*\)".*/\1/' > csrf.txt
```

- Download the backup:

```
$ curl -L -k --cookie cookies.txt --cookie-jar cookies.txt \
    --data-urlencode "download=download" \
    --data-urlencode "donotbackuprrd=yes" \
    --data-urlencode "__csrf_magic=$(head -n 1 csrf.txt)" \
    https://192.168.1.1/diag_backup.php > config-router-`date +%Y%m%d%H%M%S`.xml
```

Note: The behavior of variable expansion and other aspects of the commands may vary by shell. This example uses `bash` for the client shell.

There are additional parameters which can control the contents of the backup in several ways:

- To back up the RRD files, remove the `--data-urlencode "donotbackuprrd=yes" \` parameter from the last command.
- To include extra data such as DHCP leases and captive portal databases, add `--data-urlencode "backupdata=yes" \` to the last command.
- To include the SSH keys for the firewall, add `--data-urlencode "backupssh=yes" \` to the last command.

10.4.2 Push with SCP

The `scp` command can push the configuration file from the firewall to another host. Using `scp` to push a one-time backup by hand can be useful, but using it in an automated fashion carries risks. The command line for `scp` varies depending on the system configuration, but will be close to the following:

```
$ scp /cf/conf/config.xml \
    user@backuphost:backups/config-`hostname`-`date +%Y%m%d%H%M%S`.xml
```

Pushing the configuration in an automated manner requires the firewall administrator to generate an SSH key without a passphrase. Due to the insecure nature of a key without a passphrase, generating such a key is left as an exercise for the reader. This adds risk due to the fact that anyone with access to that file has access to the designated account, though because the key is kept on the firewall where access is restricted, it isn't a considerable risk in most scenarios. Ensure the remote user is isolated and has little to no privileges on the destination system.

A chrooted `scp` environment may be desirable in this case. The `scponly` shell is available for most UNIX platforms which allows SCP file copies but denies interactive login capabilities. Some versions of OpenSSH have chroot functionality built in for sftp (Secure FTP). These steps greatly limit the risk of compromise with respect to the remote server, but still leave the backed up data at risk. Once access is configured, a cron entry could be added to the firewall to invoke `scp`.

A summary of the setup is as follows:

- Generate an ssh key for the root user on the firewall *without a passphrase*. (Warning: dangerous!)
- Add a user to a remote system, and add the new public key to its `~/.ssh/authorized_keys` file
- Create a cron job on the firewall that would copy `/cf/conf/config.xml` to the remote system with `scp`

10.4.3 Basic SSH backup

Similar to the `scp` backup, there is another method that will work from one UNIX system to another. This method does not invoke the SCP/SFTP layer, which in some cases may not function properly if a system is already in a failing state:

```
$ ssh root@192.168.1.1 cat /cf/conf/config.xml > backup.xml
```

When executed, that command will yield a file called `backup.xml` in the current working directory that contains the remote firewall configuration. Automating this method using cron is also possible, but this method requires an SSH key without as passphrase on the host performing the backup. This key will enable administrative access to the firewall, so it must be tightly controlled. (See *Secure Shell (SSH)* for details.)

10.5 Restoring from Backups

Backups are not useful without a means to restore them, and by extension, test them. Several means for restoring configurations are available in pfSense® software. Each method has the same end result: a running firewall identical to when the backup was made.

10.5.1 Backup Compatibility

The version of pfSense Plus or pfSense CE software is not as important as the **Configuration Revision** number when determining backup compatibility. Differences in the configuration revision number indicate changes in the format of the configuration data which makes them not directly compatible.

See also:

There is a list of software versions and their corresponding configuration revision numbers at *Versions of pfSense software and FreeBSD*.

Backups using the same configuration revision can be restored as-is, both for complete configuration backups and partial (section-based) backups.

Complete backups with a lower configuration revision **can** be restored to a current version. The upgrade code will adjust the values in the configuration to convert it into a current format.

Partial (section-based) backups **cannot** be restored if they were taken on a version with a different configuration revision, as there is no mechanism for the upgrade code to handle partial backups.

Backups with a higher configuration revision **cannot** be restored to an older version. There is no mechanism to downgrade a configuration as the older version will have no knowledge of changes which happened in future versions of the software.

Restoring between pfSense CE and pfSense Plus or vice versa may work in many cases, but results depend upon the target hardware and version. For example, restoring to pfSense Plus on hardware with an integrated Ethernet switch may require manual adjustments. Contact Netgate TAC for specific guidance.

10.5.2 Restoring with the GUI

The easiest way for most users to restore a configuration is by using the GUI:

- Navigate to **Diagnostics > Backup & Restore**
- Locate the **Restore Backup** section (Figure *GUI Restore*).
- Select the area to restore, or leave at the default selection for a complete backup.

Note: This value must match the **Backup area** chosen when creating the backup.

- Click **Browse**
- Locate the backup file on the local PC
- Click **Restore Configuration**

The firewall will then apply the configuration and reboot with the settings obtained from the backup file.

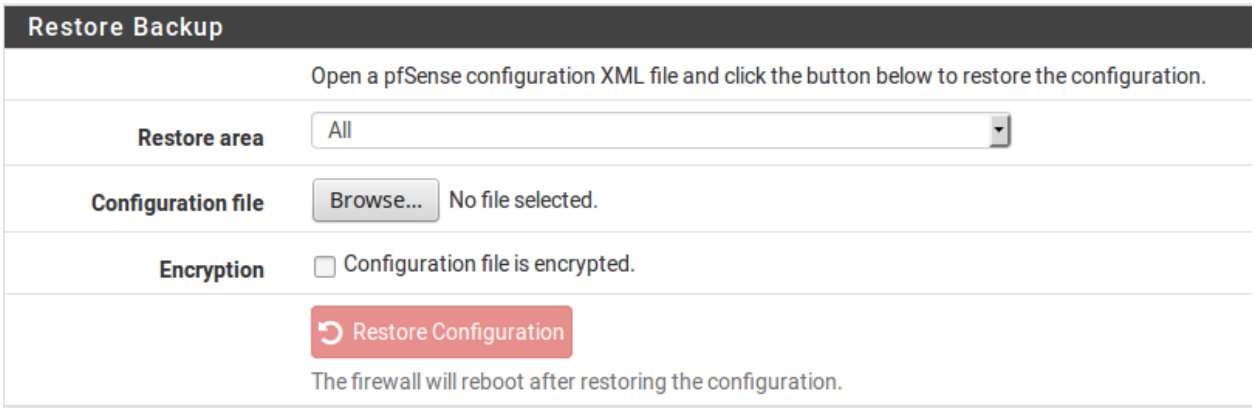


Fig. 11: GUI Restore

While easy to work with, this method has prerequisites when dealing with a full restore to a new installation. First, it would need to be done after the new target system is fully installed and running. Second, it requires an additional PC connected to a working network or crossover cable behind the firewall being restored.

Restore Options

Restore Area
Restores a backup containing only a single configuration area, rather than a complete configuration backup.

> **Warning:** Restoring a single area does not trigger a reboot nor does it cause any part of the configuration to be reapplied. To ensure the restored configuration area is active, issue a reboot or manually refresh the configuration for the relevant area after restore (e.g. edit/save/apply on a page, issue a filter reload, etc.).

> **Warning:** When restoring a single area, the area being restored must be from the **same** version. Upgrade code cannot run on single areas of the configuration, and thus cannot be adjusted if the format of the area changed from a previous version.

> **Warning:** This does not restore one area from a full backup, the backup file must only contain the area to restore.

Note: This value must match the **Backup area** chosen when creating the backup.

Configuration File
A **Browse** button to select a backup file to upload and restore.

Preserve Switch Configuration
This option is available on Netgate hardware with integrated switches. When set, the current active switch configuration will be copied into the restored configuration, preserving it for later use. This makes it easier to restore a configuration from hardware without an integrated switch.

Note: This only copies the integrated switch configuration, and does not copy VLAN or LAGG interface entries which may be relevant to using the switch. This behavior is safer, as the configuration being restored may also contain important configuration data in those areas.

Encryption
When set, a **Password** field is presented, the contents of which is used by the firewall to decrypt the contents of the backup file before restoring the configuration.

10.5.3 Configuration History

For minor problems, using one of the internal backups on the firewall is the easiest way to back out a change. By default, the previous `30` configurations are stored in the **Configuration History**, along with the current running configuration. The configuration history is found at **Diagnostics > Backup & Restore** on the **Config History** tab. The amount of entries available in the configuration history is configurable (*Configuration Backup Cache Settings*).

Configuration History List

Each row in the configuration history contains the following items:

Diff Selectors
Two columns of radio selectors which pick entries for viewing the differences between them. See *Config History Diff* for details.

Checkbox (Plus Only)
Selection checkbox for deleting multiple entries from the history. Select one or more entries and then click the  **Delete** button under the list.

Date
The date and time at which the configuration file was made.

Configuration Change
An identifier declaring which system or user made a configuration change along with a brief description of what changed in the configuration.

Version/Revision
The configuration revision of the entry, as described earlier in this document in *Backup Compatibility*.

See also:

- *Versions of pfSense software and FreeBSD*

Size
The amount of disk space consumed by this backup.

Actions
Icons which take action on each entry individually, including:

-  : Restore this configuration from the history.
-  : Download this configuration file from the history.
-  : Delete this configuration file from the history.

Note: There is no need to delete entries by hand to save space; the firewall automatically deletes old configuration backups when it creates new entries.

However, it is a good practice to remove known-bad configuration changes to ensure that they are not accidentally restored.

Note: The order of these columns differs between pfSense Plus and pfSense CE software.

Restoring from Configuration History

To restore a configuration from the history:

- Navigate to **Diagnostics > Backup & Restore**
- Click the **Config History** tab (Figure *Configuration History on pfSense Plus Software*)
- Locate the desired backup in the list
- Click  to restore that configuration file

Warning: Restoring a configuration with this method does not initiate an automatic reboot. Minor changes do not require a reboot, though reverting some major changes will.

If a change was only made in one specific section, such as firewall rules, trigger a refresh in that area of the GUI to apply the changes. For firewall rules, a filter reload would be sufficient. For OpenVPN, edit and save the VPN instance. The necessary actions to take depend on the changes in the restored configuration, but the best way ensure that the full configuration is active is to reboot.

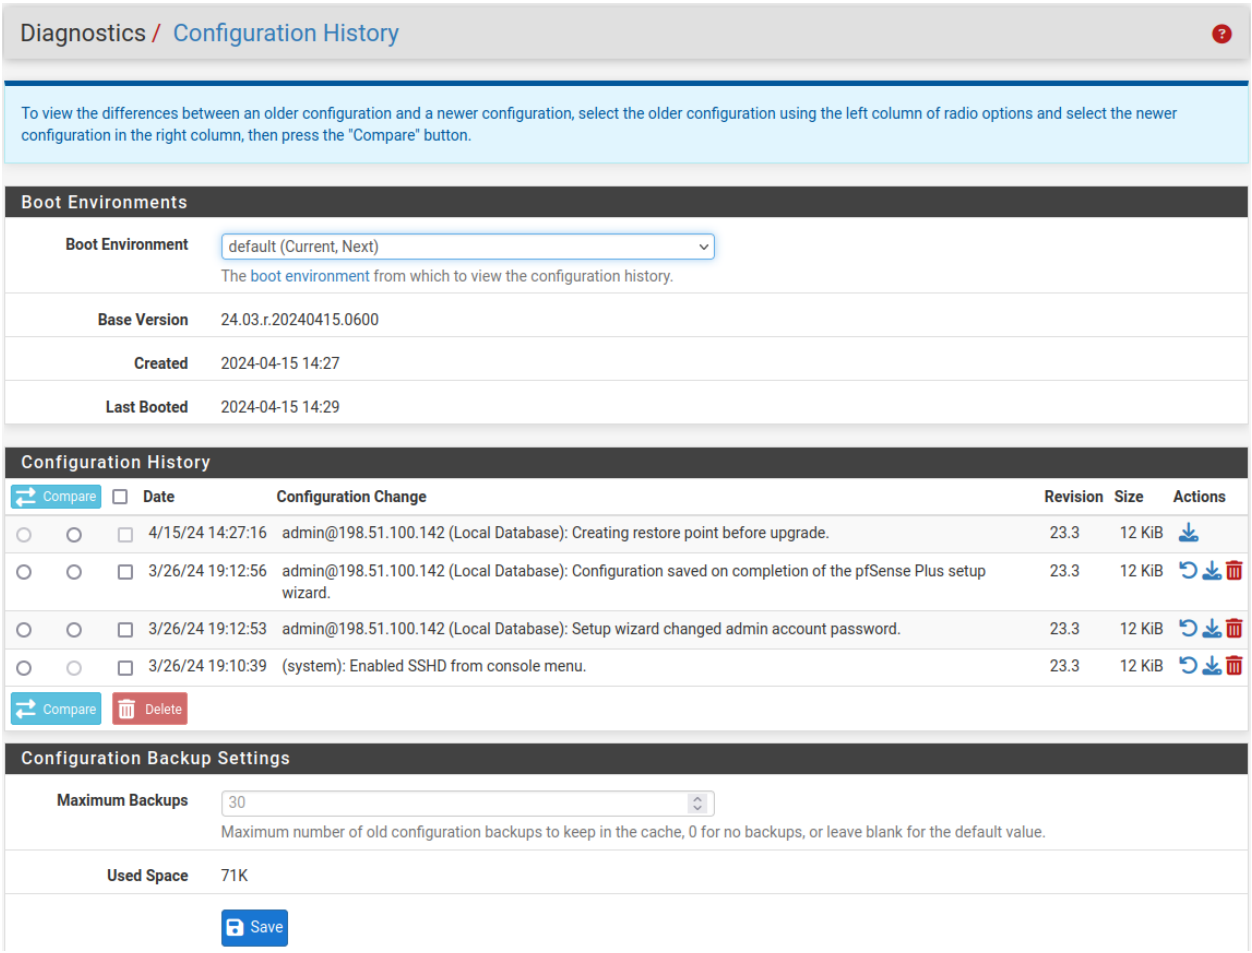


Fig. 12: Configuration History on pfSense Plus Software

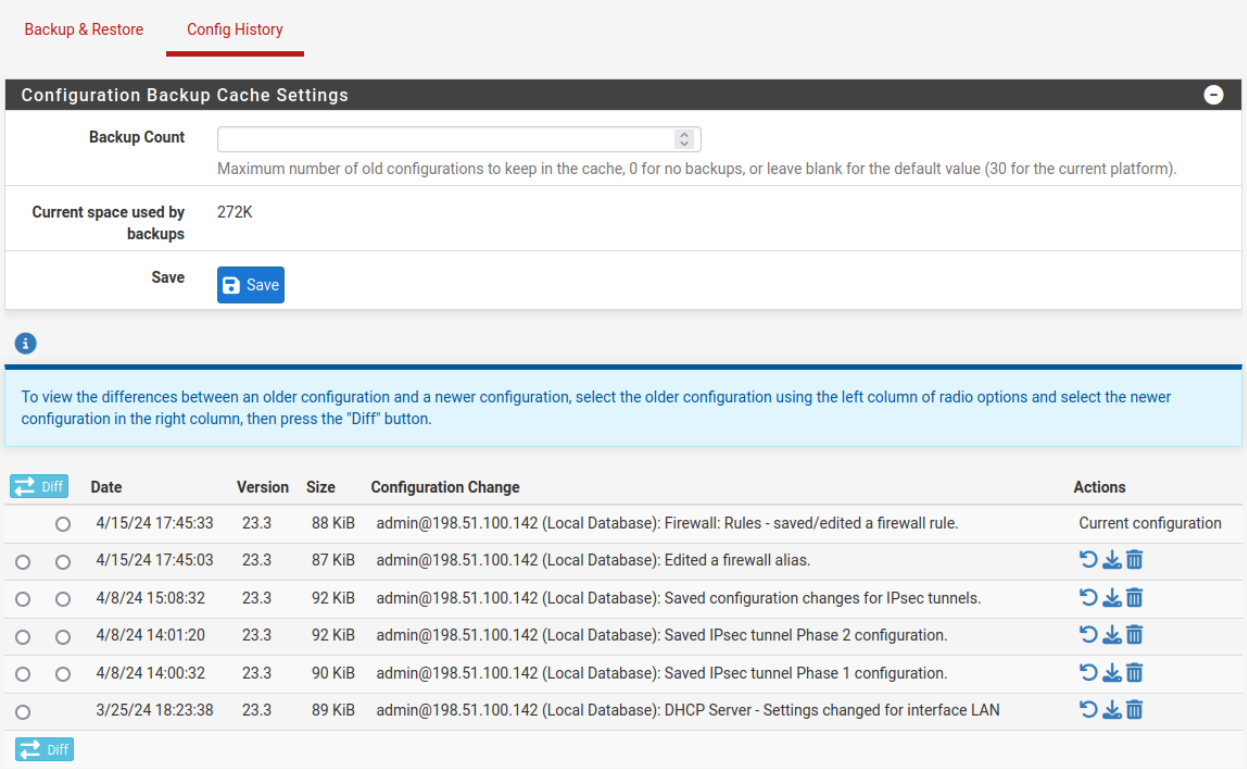


Fig. 13: Configuration History on pfSense CE Software

If necessary, reboot the firewall with the new configuration by going to **Diagnostics > Reboot System** and click **Yes**.

ZFS Boot Environments (Plus Only)

On pfSense Plus software installations utilizing *ZFS Boot Environments*, this page contains an additional **Boot Environments** section. Each Boot Environment contains its own Configuration History, and the option in this section controls which set of Configuration History entries the page displays.

The **Boot Environments** section contains the following fields:

Boot Environments
: A drop-down selection list containing all the available ZFS Boot Environments. Selecting a different entry refreshes the page to display Configuration History entries from that Boot Environment.

Base Version
: Read-only field indicating the version of pfSense software installed in the selected Boot Environment.

Created
: Read-only field indicating the date that the selected Boot Environment was created.

Last Booted
: Read-only field indicating the date that the selected Boot Environment was last booted.

Configuration Backup Cache Settings

The amount of backups stored in the configuration history may be changed if needed.

- Navigate to **Diagnostics > Backup & Restore**
- Click the **Config History** tab
- Click  at the right end of the **Configuration Backup Cache Settings** bar to expand the settings (If present)
- Enter the new number of configurations to retain in the **Backup Count** field
- Click **Save**

Note: On pfSense Plus software the settings are at the bottom of the page and always visible. On pfSense CE software the settings are above the list, but collapsed by default.

Along with the configuration count, the page also displays the amount of space consumed by the backup cache.

Config History Diff

The differences between any two configuration files may be viewed in the **Config History** tab.

On the left side of the configuration history list there are two columns of radio buttons. Use the leftmost column to select the *older* of the two configuration entries, and then use the right column to select the *newer* of the two entries. Once both entries have been selected, click **Compare** (Plus) or **Diff** (CE) at either the top or bottom of the column.

Console Configuration History

The configuration history is also available from the console menu as option `15`, **Restore Recent Configuration**. The menu selection will list recent configuration files and offer to restore one. This is useful if a recent change has locked administrators out of the GUI or taken the firewall off the network.

10.5.4 Restoring by Mounting the Disk

Attaching the disk from an installation of pfSense software to a computer running FreeBSD enables the drive to be mounted by the FreeBSD host and a new configuration may be copied directly onto the installed system, or a configuration file from a failed system may be copied off.

Note: This can also be performed on a separate installation of pfSense in place of a computer running FreeBSD, but do not use an active production firewall for this purpose. Instead, use a spare or test firewall.

The `config.xml` file is kept in `/cf/conf/`, but the difference is in the location where this directory resides. This is part of the root slice, typically `da0p2`. The drive and partition name will vary depending on disk type and position in the host.

10.5.5 Encrypted Configuration files

The GUI can automatically determine the correct decryption method when restoring an encrypted configuration backup file, whether it's from a current version or an older version. When restoring an encrypted configuration file, check **Configuration file is encrypted** then enter the password in the **Password** field, and restore as usual from there.

Encrypted configuration files can be manually decrypted using the correct password for offline inspection.

The method used to encrypt configuration files has changed in recent versions, so use the method appropriate for the version which generated the encrypted configuration file.

In any of the following cases, replace <PASSWORD> with the appropriate password string, and change the filenames as needed.

Plus 22.05 and CE 2.7.0 and later

These versions use secure options with high iterations for increased security:

```
$ openssl enc -d -a -aes-256-cbc \
  -in config-encrypted.xml -out dencryptedfile.xml \
  -pass pass:<PASSWORD> -salt -md sha256 -pbkdf2 -iter 500000
```

Warning: If the password contains any non-shell-safe characters such as `!`, `` `,\ ``, `;` or similar, escape the characters or quote the string.

These versions also include a *PHP shell script* which can encrypt and decrypt configurations from a shell on the firewall itself:

```
$ pfSsh.php playback cryptconfig \
  decrypt /root/config-encrypted.xml /root/dencryptedfile.xml
```

The script will prompt for the decryption password.

Plus 21.02 through 22.01 / CE 2.5.x through CE 2.6.x

These versions used more secure parameters than the older options, but with the default iteration count:

```
$ openssl enc -d -a -aes-256-cbc \
  -in config-encrypted.xml -out dencryptedfile.xml \
  -pass pass:<PASSWORD> -salt -md sha256 -pbkdf2
```

Warning: If the password contains any non-shell-safe characters such as `!`, `` `,\ ``, `;` or similar, escape the characters or quote the string.

Older versions

Versions before the ones stated previously used older legacy options:

```
$ openssl enc -d -a -aes-256-cbc \
  -in config-encrypted.xml -out dencryptedfile.xml \
  -pass pass:<PASSWORD> -salt -md md5
```

Warning: If the password contains any non-shell-safe characters such as `!`, `` `,\ ``, `;` or similar, escape the characters or quote the string.

10.6 Automatically Restore Configuration During Installation

In addition to restoring through the GUI, pfSense® software offers methods which restore a configuration to a new setup without going through all the trouble of setting up a client and restoring using a web browser.

These methods are significantly easier than reconfiguring the LAN and restoring via the network, especially in complex environments. The firewall will start up using the restored configuration immediately without needing intermediate steps.

- *Recover config.xml From Existing Installation*
- *Restore Configuration from Media During Install*
- *Restore using the External Configuration Locator (ECL)*

10.6.1 Recover config.xml From Existing Installation

The installer has a *Configuration Restore* option which can read configuration files and other key data (SSH host keys, DHCP leases) from an existing installation before it starts, then restore those files to the new installation at the end of the process.

This is useful for upgrades, filesystem changes, loader changes, or any other situation requiring a reinstallation on the same disk.

Note: The **Configuration Restore** option works on installations using either UFS or ZFS.

See *Configuration Restore* for information on how to utilize this feature during installation.

The firewall will boot off the target disk with the configuration restored by the installer already in place. The firewall will reinstall packages automatically in the background.

10.6.2 Restore Configuration from Media During Install

The *Configuration Restore* feature will look for files named `config.xml` anywhere on a FAT or FAT32 partition. Selecting one of these files will copy it into the target installation automatically during the installation process.

The configuration may include additional data from options on the backup page, such as RRD, SSH keys, DHCP lease databases, and captive portal data. The configuration may also be encrypted, the installer will prompt for the password to decrypt the configuration if necessary.

> **Warning:** This feature cannot read drives formatted with exFAT, only FAT or FAT32.
>
> For this feature to work correctly, the USB drive must contain a partition table, and it must not be formatted as a raw device.

Tip: The pfSense software memstick installation image contains a FAT partition which the installer can use for this purpose. If the partition is not visible on the workstation which wrote the memstick image, remove and reinsert the USB drive.

This feature works with any FAT or FAT32 partition the installer can mount during the installation process. This can be a USB thumb drive/memory stick or an optical disk/virtual drive.

- Connect a USB drive formatted with a FAT or FAT32 partition
- Copy a backup configuration file to the drive
- Rename the backup to `config.xml`

  Example: If the USB drive is `E:`, the full path would be `E:\config.xml`

  Note: The installer looks for `config.xml` in any directory on the drive, there are no restrictions on where the file must be located.

- Unmount/eject the USB drive, remove it, then plug it into the firewall

See *Configuration Restore* for information on how to utilize this feature during installation.

10.6.3 Restore using the External Configuration Locator (ECL)

pfSense software also includes a feature called the External Configuration Locator, or ECL for short. The ECL process runs at boot time to, as the name implies, locate configuration files on external storage. If the ECL finds a configuration file, it copies that file to the firewall disk, replacing any existing configuration.

Note: The ECL runs on every boot, so its use is not limited to fresh installations.

This procedure is nearly identical to the method in *Restore Configuration from Media During Install*, but the USB disk containing the configuration does not need to be present during the installation. The same warnings from that procedure also apply here.

- On a FAT, FAT32, or UFS formatted USB drive, make a directory called `config`
- Copy a backup configuration file to the `config` directory
- Rename the backup to `config.xml`

Example: If the USB drive is `E:`, the full path would be `E:\config\config.xml`.

Note: The ECL also looks for `config.xml` in the root directory of the drive, but the best practice is to place the file in the `config` directory.

- Unmount/eject and remove the USB drive
- Install pfSense software as usual

  This is optional, since the ECL runs on existing installations.
- Reboot the firewall
- Insert the USB drive containing the configuration while the firewall boots and the ECL will read in the configuration file from there

---

Note: USB drives which only contain files can be inserted before the firewall boots. Bootable USB drives, such as the installation memstick, should not be inserted until after the firewall has started to boot from its own disk. This behavior will vary by target device and its boot preferences. Monitor the console to find the appropriate timing.

Timing is also affected by the speed of the device. Slower systems may not mount the USB drive before the ECL runs.

- Wait for the firewall to complete the boot process
- Check that the configuration was loaded properly

  If the configuration did not load as expected, check the file location and name on the USB drive, and check the timing of when the USB drive was present during the boot process, then start over. Monitor the console for details.
- Remove the USB drive once the correct configuration file is in place

If this is the first boot post-installation, then this process also triggers reinstallation of packages listed in the restored configuration.

Warning: This procedure will copy the `config.xml` file from the USB drive to the target drive at **every** boot. However, the running firewall **will not** copy its own configuration back to the USB drive. Thus, leaving the drive inserted in the firewall will result in losing **all** configuration changes not present in the configuration file on the USB drive.

10.7 Restoring a Configuration File to a Different Version

Configurations are specific to a given version of pfSense® software. The configuration is the same on all platforms and architectures using the same version of pfSense software. The version of FreeBSD used is not relevant.

Generally speaking, a complete older configuration version can always be restored to a newer release of pfSense software. The firewall will upgrade the configuration, if necessary, provided that it has the **entire** configuration and not a partial copy.

A newer configuration **cannot** be restored to an older release that had a different configuration version. Certain releases of pfSense software had the same configuration version, and restoring between those releases is possible, but still not

recommended. See *Versions of pfSense software and FreeBSD* to see which configuration versions were used on specific releases.

A configuration **section** or partial configuration cannot be restored between different configuration versions. It may work by pure luck, but often there are configuration format differences that require changes to be made to the older configuration. These changes are automatic if a complete configuration is restored. If a partial restore is required, perform a full upgrade in a test VM or lab and then copy the needed section out of the resulting `config.xml` post-upgrade.

10.8 Caveats and Gotchas

While the configuration XML file kept by pfSense® software includes all the settings, it does not include any changes that may have been made to the system by hand, such as manual modifications of source code. Additionally, some packages require extra backup methods for their data.

The configuration file may contain sensitive information such as VPN keys or certificates, and passwords (other than the admin password) in plain text. Some passwords must be available in plain text during run time, making secure hashing of those passwords impossible (*Password Storage Security Policies*). Hence, backup copies of these files must also be protected in some way. If they are stored on removable media, take care with physical security of that media and/or encrypt the drive.

If the GUI must be used over the WAN without a VPN connection, at least use HTTPS. Otherwise, a backup is transmitted in the clear, including any sensitive information inside that backup file. We strongly recommend using a trusted network or encrypted connection.

10.9 Password Storage Security Policies

Sensitive data such as PPPoE/PPTP client, PPTP VPN, DynDNS passwords as well as remote authentication servers RADIUS (shared secret), LDAP (bind user password), and IPsec shared secrets, among others, appear in plain text or with reversible Base64 encoding in the pfSense® software configuration file, `config.xml`. This is a deliberate design decision in m0n0wall that has been carried over here.

Since the firewall cannot prompt the user for a password each time it is required, the implementations of affected areas require plain text passwords to operate. pfSense software could, of course, use some snake oil encryption on those passwords, but that would only create a false sense of security. Any encryption applied to the passwords could be reversed by anyone with access to the source code (i.e. everybody). Hashes like SHA256 cannot be used where the plain text password is needed at a later stage, unlike for the system password, which is only stored as a hash.

By leaving the passwords in plain text, it is very clear that `config.xml` deserves to be stored in a secure location (and/or encrypted with one of the countless programs out there). Any sort of hashing used would not be secure, and would be dangerous because it would give the impression of security where none exists.

See also:

- *Backup Files and Directories with the Backup Package*

Thanks to the XML-based configuration file used by pfSense® software, backups are a breeze. All the settings for the system are held in one single file (see *XML Configuration File*). In the vast majority of cases, this one file can be used to restore a system to a fully working state identical to what was running previously. There is no need to make an entire system backup, as the base system files are not modified by a normal, running, system.

Note: In rare cases, packages may store files outside of config.xml, check the package documentation for additional information and backup suggestions.

10.10 Backup Strategies

The optimal backup strategy can be summarized in the following points:

- Take frequent backups
- Keep multiple copies of backups in a safe location off the firewall
- Periodically test backups

The remainder of this section expands on these points.

The best practice is to make a backup after each minor change, and both before and after each major change or series of changes. Typically, an initial backup is taken in case the change being made has undesirable effects. An after-the-fact backup is taken after evaluating the change and ensuring it had the intended outcome. Periodic backups are also helpful, regardless of changes, especially in cases where a manual backup may be missed.

pfSense software makes an internal backup upon each change, and the best practice is to download a manual backup as well. The automatic backups made on each change are useful for reverting to prior configurations after changes have proven detrimental, but are not good for disaster recovery as they are on the system itself and not kept externally. As it is a fairly simple and painless process, administrators should make a habit of downloading a backup now and then and keeping it in a safe place. Backups may be handled easily and automatically using the free **AutoConfigBackup** service.

Tip: Backup files can contain sensitive information, so carefully consider security measures for backups kept off the firewall. If they are on other network file shares, ensure access is restricted. For offline backups, consider physical security measures such as keeping media containing backups in a fire safe and at a remote secure location such as a second office or bank safe deposit box.

If changes have been made to system files, such as custom patches or code alterations, those changes must be backed up manually or with the backup package described in *Backup Files and Directories with the Backup Package*, as they will not be backed up or restored by the built-in backup system. This includes alterations to system files mentioned elsewhere in the documentation, such as `/boot/device.hints`, `/boot/loader.conf.local`, and others.

Note: Custom patches should be handled using the **System Patches** package, which is backed up with `config.xml`, rather than saving manually patched files.

In addition to making backups, **backups must also be tested**. Before placing a system into production, backup the configuration, wipe the disk, and then attempt some of the different restoration techniques in this chapter. The best practice is to periodically test backups on a non-production machine or virtual machine. The only thing worse than a missing backup is an unusable backup!

RRD graph data can optionally be held in the XML configuration file backup. This behavior is disabled by default due to the resulting size of the backup file. There are other ways to ensure this data is backed up safely. See *Backup Files and Directories with the Backup Package* later in this chapter.

CHAPTER

ELEVEN

INTERFACE TYPES AND CONFIGURATION

11.1 WAN vs LAN Interfaces

pfSense® software treats interfaces differently based on whether they act as a WAN type interface (e.g. connection to an upstream network) or a LAN type interface (e.g. connection to an internal network). Most traditional interfaces will fall into one of the two categories, with VPN interfaces being more of a gray area.

Note: The NAT portions of this document only refer to IPv4 behavior, not IPv6.

11.1.1 Choosing between WAN and LAN Types

The **IPv4 Upstream Gateway** and **IPv6 Upstream Gateway** options on the *interface configuration* control whether the firewall considers an assigned interface as a WAN or LAN type interface.

If an interface has a gateway selected the firewall treats it as a WAN type interface. If an interface **does not** have a gateway selected the firewall treats as a LAN type interface.

There is no way to change the default behavior of dynamic interface types such as DHCP, PPP, and most assigned VPN interfaces. The GUI hides the gateway options on the interface configuration for these types of interfaces. The behavior of these interfaces is noted in the remainder of this document where relevant.

No matter how the firewall treats an interface by default the firewall behavior can almost always be adjusted through the use of options in the GUI.

11.1.2 WAN Type Interface

A WAN type interface is an interface through which the Internet can be reached, directly or indirectly. The firewall treats any interface with a gateway selected on its *interface configuration* as a WAN type interface. Dynamic IP address interfaces such as DHCP and PPP receive a dynamic gateway automatically and the firewall always considers them WAN interfaces.

For example, a static IP address WAN (e.g. **Interfaces > WAN**) would typically have a gateway selected such as *WAN_GW*. If this gateway selection is **not** present the firewall will treat the interface as a LAN type interface instead.

The firewall behavior changes in several ways for WAN type interfaces:

- The firewall performs outbound NAT on traffic **exiting** a WAN type interface when using *Automatic* or *Hybrid* outbound NAT modes.
- The firewall **will not** perform outbound NAT for traffic **originating** from the subnet(s) directly attached to a WAN type interface when using *Automatic* or *Hybrid* outbound NAT modes.

- The firewall includes a WAN type interface in the count of WAN interfaces for Multi-WAN features. Some functions are hidden unless the firewall has more than one WAN type interface.
- The firewall adds `reply-to` to firewall rules on a WAN type interface which returns packets for connections coming in through that WAN back out via the same WAN where possible.

Note: This behavior can be overridden on a per-rule basis using the option on firewall rules, or it can be disabled globally on **System > Advanced**, **Firewall & NAT** tab.

- The firewall adds `route-to` to automatic firewall rules for outbound traffic on a WAN type interface which ensures outbound traffic on the interface is sent to the configured gateway.
- The traffic shaper wizard treats a WAN type interface as a WAN.
- The DNS Resolver will not allow queries from the subnet(s) on a WAN type interface without a manual ACL entry.

11.1.3 LAN Type Interface

A LAN type interface is an interface which connects to a local network, for example a LAN, DMZ, management network, guest network, and so on. Typically, this also includes site-to-site links used to reach other local or internal networks, such as VPNs and private or dedicated circuits.

The firewall treats any assigned interface **without** a gateway selected on its *interface configuration* as a LAN type interface.

Warning: Do not select a gateway on the **Interfaces** menu entry for local interfaces such as LAN or for site-to-site VPNs.

Local and other interfaces may have a gateway defined under **System > Routing** so long as that gateway **is not** selected on its interface configuration.

The firewall behavior changes in several ways for LAN type interfaces:

- The firewall will perform outbound NAT for traffic **originating** from the subnet(s) directly attached to a LAN type interface when that traffic exits a WAN type interface and *Automatic* or *Hybrid* outbound NAT mode is active.
- If NAT reflection is active the firewall will create NAT reflection rules which allow clients on LAN type interfaces to access port forwards from behind the firewall.

Note: This behavior can be changed on a per-rule basis using the option on NAT rules, or it can be controlled globally on **System > Advanced**, **Firewall & NAT** tab.

- The firewall **will not** perform outbound NAT on traffic **exiting** a LAN type interface when using *Automatic* or *Hybrid* outbound NAT mode.
- The firewall **does not** add `reply-to` or `route-to` to firewall rules on a LAN type interface.
- The traffic shaper wizard treats a LAN type interface as a LAN.
- The DNS Resolver automatically **allows** queries from the subnet(s) on a LAN type interface.

11.1.4 VPN Interfaces

Assigned IPsec VTI and OpenVPN interfaces are treated differently than traditional interfaces. Most, but not all, of these points also apply to assigned GRE and GIF tunnel interfaces.

VPNs have numerous use cases which are similar to both LAN and WAN type interfaces, and in some cases both. For example a VPN could be for site-to-site links, remote access for mobile clients, or for connecting to the Internet through a VPN provider. The default behavior of the firewall attempts to balance the most common user needs and expectations when handling assigned VPN interfaces.

Note: WireGuard interfaces act similar to traditional interfaces when assigned, so their behavior primarily depends upon whether a gateway is selected in their interface configuration.

- The firewall treats an assigned VPN interface as a LAN type interface for NAT. This means that the firewall lists the subnets on these interfaces as traffic sources for outbound NAT, and it does not perform outbound NAT on traffic exiting these interfaces.

  In most cases a user does not expect the firewall to perform NAT on VPN traffic by default. Outbound NAT rules in *Hybrid* or *Manual* outbound NAT modes can make the firewall perform outbound NAT if a use case requires NAT.

- The firewall treats an assigned VPN interface as a WAN type interface for traffic shaping if a VPN interface is capable of using ALTQ traffic shaping.

- The firewall treats an assigned VPN interface as a WAN interface for firewall rule attributes such as `reply-to` and `route-to`. This ensures that traffic entering the firewall over a specific VPN connection returns back through the same VPN.

- The DNS Resolver treats an assigned VPN interface as a LAN interface and allows queries from subnet(s) configured on the VPN.

Note: Firewall features such as per-interface rules, NAT, and `reply-to` do not work with IPsec VTI interfaces by default. The **IPsec Filter Mode** setting can allow IPsec VTI interfaces to utilize these features. See *Advanced IPsec Settings*.

11.1.5 Verifying an Interface Type

There are a couple ways to confirm if the firewall is treating an interface as a WAN or a LAN.

The interface status page (**Status > Interfaces**) is useful for determining the interface type. For non-VPN interfaces the presence of the **Gateway IPv4** and/or **Gateway IPv6** attribute on an interface indicates that the firewall considers it as a WAN type interface.

The next easiest method is to check the outbound NAT settings at **Firewall > NAT**, **Outbound** tab. Check the **Automatic Rules** section if the mode is set to *Automatic* or *Hybrid*. WAN type interfaces will have rules in the list with their name in the **Interface** column. LAN type interfaces have their subnets listed in the **Source** column of each rule.

Note: If the outbound NAT mode is *Automatic* or *Hybrid* and there are **no** entries in the **Automatic Rules** list, that generally indicates that the firewall has either no WAN type interfaces or no LAN type interfaces. Check the gateway settings on each assigned interface and ensure that **all** WAN interfaces have a gateway selected and that **no** LAN interfaces have a gateway selected.

Another method is to start a traffic shaper wizard (**Firewall > Traffic Shaper**, **Wizards** tab) and step through until the wizard lists the interfaces. From there, check if an interface is present in either the LAN or WAN interface selection lists.

Note: This method only works with interface types capable of using ALTQ traffic shaping.

11.2 Interface Configuration

To assign a new interface:

- Navigate to **Interfaces > Assignments**
- Pick the new interface from the **Available network ports** list
- Click  **Add**

The newly assigned interface will be shown in the list. The new interface will have a default name allocated by the firewall such as OPT1 or OPT2, with the number increasing based on its assignment order. The first two interfaces default to the names WAN and LAN, but they can be renamed. These OPTx names appear under the **Interfaces** menu, such as **Interfaces > OPT1**. Selecting the menu option for the interface will open the configuration page for that interface.

11.2.1 General Configuration

The following options are available for all interface types.

Description
The name of the interface. Interface names may only contain letters, numbers and the only special character that is allowed is an underscore (_).

This changes the name of the interface on the **Interfaces** menu, on the tabs under **Firewall > Rules**, under **Services > DHCP**, and elsewhere throughout the GUI. Using a custom name makes it easier to remember the purpose of an interface and to identify an interface for adding firewall rules or choosing other per-interface functionality.

IPv4 Configuration Type
Configures the IPv4 settings for the interface. Details for this option are in the next section, *IPv4 Configuration Types*.

IPv6 Configuration Type
Configures the IPv6 settings for the interface. Details for this option are in *IPv6 Configuration Types*.

MAC address
The MAC address of an interface can be changed ("spoofed") to mimic a previous piece of equipment, depending on the type of interface.

> **Warning:** The best practice is to not force a specific MAC address. The old MAC address will generally be cleared out by resetting the equipment to which this firewall connects, or by clearing the ARP table, or waiting for the old ARP entries to expire. Changing the MAC address is a long-term solution to a temporary problem.

Spoofing the MAC address of the previous firewall can allow for a smooth transition from an old router to a new router, so that ARP caches on devices and upstream routers are not a concern. It can also be used to fool a piece of equipment into believing that it's talking to the same device that it was talking to before, as in cases where a certain network router is using static ARP or otherwise filters based on MAC address. This is common on cable modems, where they may require the MAC address to be registered if it changes.

Note: ARP cache problems tend to be very temporary, resolving automatically within minutes or by power cycling other equipment.

One downside to spoofing the MAC address is that unless the old piece of equipment is permanently retired, there is a risk of later having a MAC address conflict on the network, which can lead to connectivity problems.

If the old MAC address must be restored, this option must be emptied out and then the firewall *must* be rebooted. Alternately, enter the original MAC address of the network card and save/apply, then empty the value again.

MTU (Maximum Transmission Unit)
The Maximum Transmission Unit (**MTU**) size field can typically be left blank, but can be changed when required. Some situations may call for a lower MTU to ensure packets are sized appropriately for an Internet connection. In most cases, the default assumed values for the WAN connection type will work properly. It can be increased for those using jumbo frames on their network.

On a typical Ethernet style network, the default value is `1500`, but the actual value can vary depending on the interface configuration.

MSS (Maximum Segment Size)
Similar to the MTU field, the **MSS** field "clamps" the Maximum Segment Size (MSS) of TCP connections to the specified size in order to work around issues with Path MTU Discovery.

Speed and Duplex
The default value for link speed and duplex is to let the firewall decide what is best. That option typically defaults to *Autoselect*, which negotiates the best possible speed and duplex settings with the peer, typically a switch.

The speed and duplex setting on an interface must match the device to which it is connected. For example, when the firewall is set to *Autoselect*, the switch must also be configured for *Autoselect*. If the switch or other device has a specific speed and duplex forced, it must be matched by the firewall.

Switch Port
Netgate Appliances with an integrated switch have an option on this page which controls the link state for this interface by having it mirror the state of a switch port. In this way, a firewall interface configured as a VLAN which maps to a switch port can be set to follow the status of the physical switch port. Otherwise, since it is a VLAN attached to an internal uplink, the status would always show as up.

Consult the Netgate Product Manuals for more information on switch configuration.

11.2.2 Reserved Networks

Block Private Networks
When **Block private networks** is active, the firewall inserts a rule automatically which prevents any RFC 1918 networks (`10.0.0.0/8`, `172.16.0.0/12`, `192.168.0.0/16`) and loopback (`127.0.0.0/8`) from communicating on that interface.

This option is typically only desirable on WAN type interfaces to prevent the possibility of privately numbered traffic coming in over a public interface.

Block bogon networks
When **Block bogon networks** is active, the firewall will block traffic from a list of unallocated and reserved networks. This list is periodically updated by the firewall automatically.

> **Warning:** This option should only be used on external interfaces (WANs). It is not necessary on local interfaces, and it can potentially block required local traffic.

See *Block Bogon Networks* for more details on how this feature works.

11.3 IPv4 Configuration Types

Once an interface has been assigned, in most cases it will require an IP address. For IPv4 connections, the following choices are available in the **IPv4 Configuration Type** selector on an interface page (e.g. **Interfaces > WAN**):

- None
- Static IPv4
- DHCP
- PPP
- PPPoE
- PPTP
- L2TP

Each of these is described in this document.

11.3.1 None

When **IPv4 Configuration Type** is set to *None*, IPv4 is disabled on the interface. This is useful if the interface has no IPv4 connectivity or if the IPv4 address on the interface is being managed in some other way, such as for a VPN or tunnel interface.

11.3.2 Static IPv4

With **Static IPv4**, the interface contains a manually configured IPv4 address. When chosen, three additional fields are available on the interface configuration screen:

IPv4 Address
The IPv4 address for the interface (e.g. `192.168.1.1`).

CIDR Subnet Mask
The CIDR Subnet Mask determines the size of the subnet to which the **IPv4 Address** belongs. This must match the value used by other hosts in the same subnet.

IPv4 Upstream Gateway
An upstream gateway for IPv4 traffic, if any. Selecting a gateway here will cause the firewall to treat this interface as a WAN-type interface for NAT and related functions. See *WAN vs LAN Interfaces* for more information.

> **Warning:** Do not set a gateway for internal interfaces such as a LAN or DMZ. Only select a gateway on externally-connected interfaces such as a WAN or a private site-to-site link which the firewall should consider a WAN.
>
> Gateways may still be used on internal interfaces for the purpose of static routes without selecting an **IPv4 Upstream Gateway** here.

The **IPv4 Upstream Gateway** field is pre-populated with existing IPv4 gateways defined under **System > Routing** (*Gateways*).

The  **Add a new gateway** button is a shortcut to create a new gateway if one does not already exist. Clicking that button displays a modal form to add the gateway without leaving this page. Fill in the details requested on the new form:

Default Gateway
If this is the only WAN or will be a new default WAN, check this box. The default IPv4 and IPv6 gateways work independently of one another. The two need not be on the same interface. Changing the default IPv4 gateway has no effect on the IPv6 gateway, and vice versa.

Gateway Name
The name used to refer to the gateway internally, as well as in places like Gateway Groups, quality graphs, and elsewhere.

Gateway IPv4
The IPv4 address of the gateway. This address must be inside the same subnet as the Static IPv4 address when using this form.

Description
A bit of text to indicate the purpose of the gateway.

11.3.3 DHCP

When an interface is set to **DHCP**, the operating system will attempt automatic IPv4 configuration of this interface via DHCP. This option also activates several additional fields on the page. Under most circumstances these additional fields may be left blank.

Hostname
Some ISPs require the **Hostname** for client identification. The value in the **Hostname** field is sent as the DHCP client identifier and hostname when requesting a DHCP lease.

Alias IPv4 Address
This value used as a fixed IPv4 alias address by the DHCP client since a typical IP Alias VIP cannot be used with DHCP. This can be useful for accessing a piece of gear on a separate, statically numbered network outside the DHCP scope. One example would be for reaching a cable modem management IP address.

Reject Leases From
An IPv4 address for a DHCP server that should be ignored. For example, a cable modem that hands out private IP addresses when the cable sync has been lost. Enter the private IP address of the modem here, e.g. `192.168.100.1` and the firewall will never pick up or attempt to use an IP address supplied by the specified server.

DHCP VLAN Priority
Optionally sets a VLAN Priority tag (802.1p) on DHCP client traffic. Should only be enabled when required by an ISP and with the settings they provide.

Advanced Configuration
Enables options to control the protocol timing. In the vast majority of cases this must be left unchecked and the options inside unchanged.

Protocol Timing
The fields in this area give fine-grained control over the timing used by `dhclient` when managing an address on this interface. These options are almost always left at their default values. For more details on what each field controls, see the dhclient man page

Presets
Has several options for preset protocol timing values. These are useful as a starting point for custom adjustments or for use when the values need to be reset back to default values.

Configuration Override
Enables a field to use a custom `dhclient` configuration file. The full path must be given. Using a custom file is rarely needed, but some ISPs require DHCP fields or options that are not possible in the GUI.

11.3.4 PPP Types

The various PPP-based connection types such as PPP, PPPoE, PPTP, and L2TP are all covered in detail at *PPPs*. When one of these types is selected here on the interfaces screen, their basic options can be changed as described. To access the advanced options, follow the link on this page or navigate to **Interfaces > Assignments** on the **PPPs** tab, find the entry, and edit it there.

11.4 IPv6 Configuration Types

Similar to IPv4, the **IPv6 Configuration Type** controls if and how an IPv6 address is assigned to an interface. There are several different ways to configure IPv6 and the exact method depends on the network to which this firewall is connected and how the ISP has deployed IPv6.

> **Warning:** Every ISP is different and large providers can even vary by region.
>
> The ISP determines IPv6 settings for a circuit, and they are the only valid source for that information. As such, this documentation does not include examples for specific providers. Contact the ISP for information about their IPv6 client settings and requirements.
>
> The ISP should provide instructions and specific values for configuring IPv6 on their service. For example, on a circuit with a static IPv6 configuration the ISP should supply the subnet addresses and prefix values for the WAN itself, as well as for routed prefixes. Providers who require DHCPv6 should supply values for settings such as the prefix delegation size, along with any requirements they have for client behavior.

See also:

For more information on IPv6, including a basic introduction, see *IPv6*.

11.4.1 None

When **IPv6 Configuration Type** is set to *None*, IPv6 is disabled on the interface. This is useful if the interface has no IPv6 connectivity or if the IPv6 address on the interface is being managed in some other way, such as for a VPN or tunnel interface.

11.4.2 Static IPv6

The Static IPv6 controls work identically to the Static IPv4 settings. See *Static IPv4* for details.

With **Static IPv6**, the interface contains a manually configured IPv6 address. When chosen, three additional fields are available on the interface configuration screen: **IPv6 Address**, a prefix length selector, and the **IPv6 Upstream Gateway** field.

Note: Do not set a gateway for internal interfaces such as a LAN or DMZ. Only select a gateway on externally-connected interfaces such as a WAN or a private site-to-site link which the firewall should consider a WAN.

Gateways may still be used on internal interfaces for the purpose of static routes without selecting an **IPv6 Upstream Gateway** here.

See *WAN vs LAN Interfaces* for more information.

The default IPv4 and IPv6 gateways work independently of one another. The two need not be on the same interface. Changing the default IPv4 gateway has no effect on the IPv6 gateway, and vice versa.

11.4.3 DHCP6

DHCP6 configures automatic IPv6 configuration of this interface via DHCPv6. DHCPv6 will configure the interface with an IPv6 address, prefix length, DNS servers, etc. but not a gateway. The gateway is obtained via router advertisements, so this interface will be set to accept router advertisements. This is a design choice as part of the IPv6 specification, not a limitation of this implementation. For more information on router advertisements, see *Router Advertisements*.

Several additional fields are available for IPv6 DHCP that do not exist for IPv4 DHCP:

Use IPv4 Connectivity as Parent Interface
When set, the IPv6 DHCP request is sent using IPv4 on this interface, rather than using native IPv6. This is only required in special cases when the ISP requires this type of configuration.

Request only an IPv6 Prefix
When set, the DHCPv6 client does not request an address for the interface itself, it only requests a delegated prefix.

DHCPv6 Prefix Delegation Size
If the ISP supplies a routed IPv6 network via prefix delegation, they will publish the delegation size, which can be selected here. It is typically a value somewhere between *48* and *64*. For more information on how DHCPv6 prefix delegation works, see *DHCP6 Prefix Delegation*.

Note: To use this delegation, another internal interface must be set to an **IPv6 Configuration Type** of *Track Interface* (*Track Interface*) so that it can use the addresses delegated by the upstream DHCPv6 server.

Send IPv6 Prefix Hint
When set, the **DHCPv6 Prefix Delegation Size** is sent along with the request to inform the upstream server how large of a delegation is desired by this firewall. If an ISP allows the choice, and the chosen size is within their allowed range, the requested size will be given instead of the default size.

Debug
When set, the DHCPv6 client is started in debug mode.

Do not wait for a RA
Informs the operating system not to wait for a router advertisement when configuring the interface. This is required by some ISPs.

Do not allow PD/Address release
Prevents the operating system from sending a DHCPv6 release message on exit.

Some ISPs will release the allocated address or prefix when a client sends this message. With this option set, the client is more likely to receive the same allocation with subsequent requests.

DHCPv6 VLAN Priority
Optionally sets a VLAN Priority tag (802.1p) on DHCPv6 client traffic. Should only be enabled when required by an ISP and with the settings they provide.

Advanced Configuration
Enables a wide array of advanced tuning parameters for the DHCPv6 client. These options are rarely used, and when they are required, the values are dictated by the ISP or network administrator. See the dhcp6c.conf man page for details.

Configuration Override
Enables a field to use a custom configuration file. The full path must be given. Using a custom file is rarely needed, but some ISPs require DHCP fields or options that are not possible in the pfSense GUI.

11.4.4 SLAAC

Stateless address autoconfiguration (*SLAAC*) as the IPv6 type makes the operating system attempt to configure the IPv6 address for the interface from router advertisements (RA) that advertise the prefix and related information.

Note: DNS is not typically provided via RA, so the firewall will still attempt to get DNS servers via DHCPv6 when using SLAAC. The RDNSS extensions to the RA process may allow DNS servers to be obtained from RA in some cases. For more information on router advertisements, see *Router Advertisements*.

This selection has one additional option:

Use IPv4 connectivity as parent interface
When set, IPv6 requests are sent over the IPv4 connectivity layer used by this interface (e.g. PPPoE) rather than the parent interface directly. May be required by certain ISPs.

11.4.5 6RD Tunnel

6RD is an IPv6 tunneling technology employed by ISPs to quickly enable IPv6 connectivity for their networks, passing IPv6 traffic inside specially crafted IPv4 packets between and end user router and the ISP relay. It is related to 6to4 but is intended to be used within the ISP network, using the IPv6 addresses from the ISP for client traffic. To use 6RD, the ISP must supply three pieces of information: The 6RD prefix, the 6RD Border Relay, and the 6RD IPv4 Prefix length.

6RD Prefix
The 6RD IPv6 prefix assigned by the ISP, such as `2001:db8::/32`.

6RD Border Relay
The IPv4 address of the ISP 6RD relay.

6RD IPv4 Prefix Length
Controls how much of the end user IPv4 address is encoded inside the 6RD prefix. This is normally supplied by the ISP. A value of `0` means the entire IPv4 address will be embedded inside the 6RD prefix. This value allows ISPs to effectively route more IPv6 addresses to customers by removing redundant IPv4 information if an ISP allocation is entirely within the same larger subnet.

11.4.6 6to4 Tunnel

Similar to 6RD, *6to4* is another method of tunneling IPv6 traffic inside IPv4. Unlike 6RD, however, 6to4 uses constant prefixes and relays. As such there are no user-adjustable settings for using the *6to4* option. The 6to4 prefix is always `2002::/16`. Any address inside the `2002::/16` prefix is considered a 6to4 address rather than a native IPv6 address. Also unlike 6RD, a 6to4 tunnel can be terminated anywhere on the Internet, not only at the end user ISP, so the quality of the connection between the user and the 6to4 relay can vary widely.

6to4 tunnels are always terminated at the IPv4 address of `192.88.99.1`. This is an IPv4 anycast address, meaning that although the IPv4 address is the same everywhere, it can be routed regionally toward a node close to the user.

Another deficiency of 6to4 is that it relies upon other routers to relay traffic between the 6to4 network and the remainder of the IPv6 network. There is a possibility that some IPv6 peers may not have connectivity to the 6to4 network, and thus these would be unreachable by clients connecting to 6to4 relays, and this could also vary depending upon the 6to4 node to which the user is actually connected.

11.4.7 Track Interface

The *Track Interface* choice works in concert with another IPv6 interface using DHCPv6 Prefix Delegation. When a delegation is received from the ISP, this option designates which interface will be assigned the IPv6 addresses delegated by the ISP and in cases where a larger delegation is obtained, which prefix inside the delegation is used.

IPv6 Interface
: A list of all interfaces on the system currently set for dynamic IPv6 WAN types offering prefix delegation (DHCPv6, PPPoE, 6rd, etc.). Select the interface from the list which will receive the delegated subnet information from the ISP.

IPv6 Prefix ID
: If the ISP has delegated more than one prefix via DHCPv6, the IPv6 Prefix ID controls which of the delegated `/64` subnets will be used on this interface. This value is specified in hexadecimal.

  For example, If a `/60` delegation is supplied by the ISP that means 16 `/64` networks are available, so prefix IDs from `0` through `f` may be used.

For more information on how prefix delegation works, see *DHCP6 Prefix Delegation*.

11.5 Interface Groups

Unlike the other interfaces in this chapter, an **Interface Group** is not a type of interface that can be assigned. Interface groups are used to apply firewall or NAT rules to a set of interfaces on a common tab. If this concept is unfamiliar, consider how the firewall rules for OpenVPN, the PPPoE server, or L2TP server work. There are multiple interfaces in the underlying OS, but the rules for all of them are managed on a single tab for each type.

If many interfaces of a similar function are present on the firewall that need practically identical rules, an interface group may be created to add rules to all the interfaces at the same time. Interfaces can still have their own individual rules, which are processed after the group rules.

11.5.1 Interface Group Options

When creating or editing an Interface Group, the following options are available:

Group Name
: The name of the interface group. Has the same restrictions as the name of an interface. The name may only contain upper and lowercase letters, no numbers, spaces, or special characters.

Group Description
: An optional text description for reference.

Group Members
: A multi-select list of assigned interfaces on the firewall from which group members can be added. Add interfaces to the group by selecting them with Ctrl-click (PC) or cmd-click (MAC).

11.5.2 Creating an Interface Group

To create an interface group:

- Navigate to **Interfaces > Assignments**, **Interface Groups** tab
- Click  **Add** to create a new group
- Fill in the options as described in *Interface Group Options*
- Click **Save**

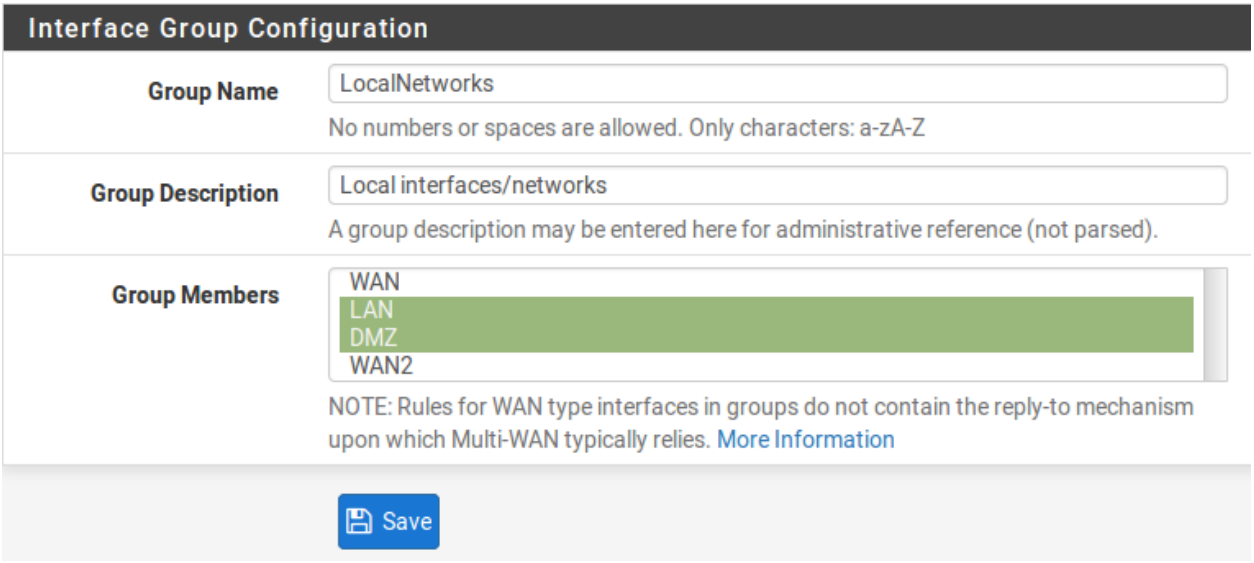


Fig. 1: Add Interface Group

11.5.3 Using an Interface Group

Interface groups each have an individual tab under **Firewall > Rules** to manage their rules. Figure *Interface Group Firewall Rules Tab* shows the firewall rule tab for the group defined in figure *Add Interface Group*

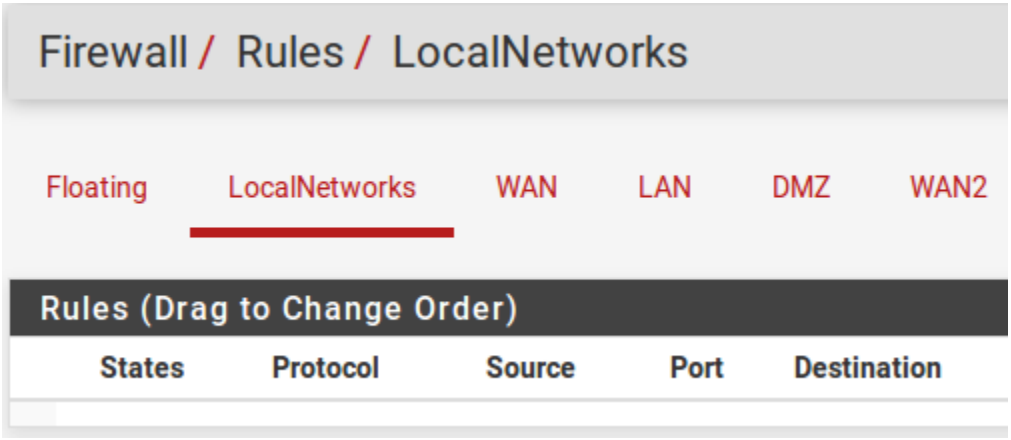


Fig. 2: Interface Group Firewall Rules Tab

See also:

Configuring Firewall Rules for information on managing firewall rules.

11.5.4 Group Rule Processing Order

The rule processing order for user rules is:

- Floating rules
- Interface group rules
- Rules on the interface directly

For example, if a rule on the group tab matches a connection, the interface tab rules will not be consulted. Similarly, if a floating rule with **Quick** set matched a connection, the interface group rules will not be consulted.

The processing order prevents some combination of rules that otherwise might be a good fit. For example, if a general blocking rule is present on the group, a rule on a specific interface cannot override it. Same with a pass rule, a specific interface rule cannot block traffic passed on a group tab rule.

11.5.5 Use with WAN Interfaces

The best practice is to not use interface groups with multiple WANs. Doing so may appear to be convenient, but the group rules do not receive the same treatment as actual WAN tab rules. For example, rules on a tab for a WAN-type interface will receive `reply-to` which allows PF to return traffic back via the interface from which it entered. Group tab rules do not receive `reply-to` which effectively means that the group rules only function as expected on the WAN with the default gateway.

11.6 PPPs

There are four types of PPP interfaces:

- **PPP** for cellular and modem devices
- **PPPoE** for DSL or similar connections
- **PPTP** and **L2TP** for ISPs that require them for authentication

In most cases these are managed by the interface settings directly, but the settings are also available under **Interfaces > Assignments** on the **PPPs** tab.

See also:

- *PPP Logs*

11.6.1 Multi-Link PPP (MLPPP)

Multi-Link PPP (MLPPP) is available for any type of PPP instance by selecting multiple **Link Interface** entries at the same time.

> **Warning:** MLPPP only works on multiple circuits from the same provider where the provider supports MLPPP.
>
> MLPPP is not compatible with the `if_pppoe` backend for PPPoE WANs (*Use if_pppoe Kernel Module*).

MLPPP bonds multiple PPP links into a single larger aggregate channel. Unlike other multi-WAN techniques MLPPP can utilize the full bandwidth of all links for a single connection. MLPPP also does not have the usual concerns about load balancing and failover. The MLPPP link is presented as one interface with one IP address. If one link fails the connection functions the same but with reduced capacity.

For more information on MLPPP, see *Multiple WAN Connections*.

11.6.2 PPP (Point-to-Point Protocol) Interface Types

Add or edit a PPP entry as follows:

- Navigate to **Interfaces > Assignments** on the **PPPs** tab
- Click  to edit an existing entry or  to add a new entry
- Set the **Link Type**

The **Link Type** determines the remaining options on the page. The available link types are explained throughout the remainder of this document.

PPP (Cellular Modem)

The **PPP** link type is used for talking to a modem over a serial device. This can be anything from a USB modem dongle for accessing a cellular network down to an old hardware modem for dial-up access.

Note: Some cellular modems appear as Ethernet devices and not serial devices. Those are configured as regular interfaces, not as PPP devices.

See also:

- *Cellular Wireless*

When configuring a PPP device, the following options are available:

Link Interface
A list of serial devices that the firewall can use to communicate with a modem. Click on a specific entry to select it for use by the firewall.

Note: The firewall does not automatically detect the serial device for a modem. Some modems present themselves as several devices and the sub-device for the PPP line may be any of the available choices. Start with the last device, then try the first, and then others in between until a functional device is found.

Description
A text description of this PPP instance, for reference (e.g. `VZW Modem`).

Country
The country in which this modem resides (e.g. *United States*).

The firewall populates the **Provider** list based on the value of this field.

Provider
The cellular service provider for this modem (e.g. *Verizon*).

The firewall populates the **Plan** list based on the value of this field.

Plan
The type of cellular service this modem uses from **Provider**.

This populates the remaining fields where possible with values specific to the **Plan**.

The remaining options can be configured manually if other values are needed, or when using an unlisted provider:

Username and Password
The credentials used for the PPP login, if any.

Phone Number
The number to dial at the ISP to gain access. For cellular providers this tends to be a number such as `*99#` or `#777`. For dial-up this is usually a traditional telephone phone number.

Access Point Name (APN)
Some ISPs require this value to identify the service to which the client connects. Some providers use this to distinguish between consumer and business plans or legacy networks.

APN Number
Optional setting. Defaults to `1` if the APN is set, and ignored when APN is unset.

SIM PIN
Security code on the SIM to prevent unauthorized use of the card.

> **Warning:** Do not enter anything here if the SIM does not have a PIN.

SIM PIN Wait
Number of seconds the firewall will wait for the SIM to discover network after the PIN is sent to the SIM. If the delay is not long enough the SIM may not have time to initialize properly after unlocking.

Init String
The modem initialization string, if necessary. Most modern modems do not require a custom initialization string.

Note: Do not include `AT` at the beginning of the command.

Connection Timeout
Time the firewall will wait for a connection attempt to succeed, in seconds. Default is `45` seconds.

Uptime Logging
When checked, the firewall tracks the uptime for the connection and displays it on **Status > Interfaces**.

PPPoE (Point-to-Point Protocol over Ethernet)

PPPoE is a popular method of authenticating and gaining access to an ISP network. It is commonly found on DSL and some fiber networks in certain regions, but may also be used on other link types.

Tip: For the best performance on PPPoE WANs, use the kernel-based `if_pppoe` backend. For details, see *Use if_pppoe Kernel Module*.

Due to limitations in the way PPPoE frames are processed by network cards when using MPD, incoming PPPoE traffic is limited to a single network interface queue. As such, performance may be limited or otherwise lower than expected. See *PPPoE with Multi-Queue NICs* for details.

To configure a PPPoE link, start by setting **Link Type** to *PPPoE* and complete the remainder of the settings as follows:

Link Interface(s)
A list of network interfaces the firewall can use for PPPoE. These are typically physical interfaces, but PPPoE can also work over some other interface types such as VLANs. Select one entry for normal PPPoE or multiple entries for MLPPP.

Description
An optional text description of the PPP entry.

Username and Password
The credentials for this PPPoE connection, provided by the ISP. Usernames are typically in the form of an e-mail address, such as `mycompany@ispexample.com`.

Service Name
Left blank for most ISPs, but some ISPs require this to be set to a specific value.

Contact the ISP to confirm the value if the connection does not function when left blank.

Configure NULL Service Name
Some ISPs require clients to send a `NULL` value instead of a blank service name. Check this option when the ISP requires this behavior.

Periodic Reset
Configures a pre-set time when the firewall will drop the connection and reconnect. This is rarely needed, but in certain cases it can better handle reconnections when an ISP has forced daily reconnections or similar quirky behavior.

PPTP (Point-to-Point Tunneling Protocol)

Not to be confused with a PPTP VPN, this type of PPTP interface is meant to connect to an ISP and authenticate, much the same as PPPoE. The options for a PPTP WAN are identical to the PPPoE options of the same name. Refer to the previous section for configuration information.

L2TP (Layer 2 Tunneling Protocol)

L2TP, as it is configured here, is used for connecting to an ISP that requires it for authentication as a type of WAN. L2TP works nearly identically to PPTP. Refer to the previous sections for configuration information.

L2TP has one additional option not found on other types:

Shared Secret
A shared secret the firewall will use to authenticate the tunnel connection and encrypt control L2TP control packets. May be left blank if the server does not support a shared secret.

> **Warning:** This **must** match the shared secret set on the L2TP server.

11.6.3 Advanced PPP Options

All PPP types have several advanced options in common. In most cases these settings can remain at their default values.

Click  **Display Advanced** to display these options.

Dial On Demand
The default behavior for a PPP link is to immediately connect and immediately attempt to reconnect when a link is lost. This behavior is described as **Always On**. **Dial-on-Demand** delays this connection attempt. When set, the firewall waits until a packet attempts to leave the via this interface to make a connection attempt. Once the firewall connects it will not automatically disconnect.

Idle Timeout
The firewall will hold a PPP connection open indefinitely by default. A value in **Idle Timeout**, specified in seconds, will cause the firewall to monitor the line for activity. If there is no traffic on the link for the given amount of time, the firewall will disconnect the link. If **Dial-on-Demand** has also been set, the firewall will return to dial-on-demand mode.

Note: The firewall performs gateway monitoring by default which generates two ICMP pings per second on the interface. **Idle Timeout** will not function in this case. This can be worked around by editing the gateway for this PPP link and checking **Disable Gateway Monitoring**.

Compression (vjcomp)
This option controls whether the firewall will use Van Jacobson TCP header compression for this connection. By default, the firewall will negotiate this with the peer during login and enable it if both sides are compatible. Checking **Disable vjcomp** will disable this feature. This feature is beneficial because it saves several bytes per TCP data packet when possible. The best practice is to keep the option enabled unless the remote requires it to be disabled.

Note: This compression is ineffective for TCP connections with enabled modern extensions like time stamping or SACK, which modify TCP options between sequential packets.

TCP MSS Fix
This option causes the PPP daemon to adjust incoming and outgoing TCP SYN segments so that the requested maximum segment size (MSS) is not greater than the amount allowed by the interface MTU.

This is necessary in most cases to avoid problems caused by routers which drop ICMP "Datagram Too Big" messages. Without these messages, peers cannot detect a when packets attempt to cross a link which cannot carry frames of the required size. Consider this scenario. The originating machine sends data which passes a rogue router then arrives at a host that has an MTU that is not big enough for the data. Because the IP "Don't Fragment" option is set, this machine sends an ICMP "Datagram Too Big" message back to the originator and drops the packet. The rogue router drops the ICMP message and the originator never gets to discover that it must reduce the fragment size or drop the IP "Don't Fragment" option from its outgoing data. If this behavior is undesirable, check **Disable tcpmssfix**.

Note: The MTU and MSS values for the interface may also be adjusted on the configuration page for the interface under the **Interfaces** menu, such as **Interfaces > WAN** (*Interface Configuration*).

Short Sequence (ShortSeq)

This option is only meaningful when the firewall is negotiating MLPPP with the provider. It proscribes shorter multi-link fragment headers, saving two bytes on every frame. It is not necessary to disable this for connections that are not multi-link. If MLPPP is active and this feature must be disabled, check **Disable shortseq**.

Address Control Field Compression (ACFComp)
This option only applies to asynchronous link types. It saves two bytes per frame. To disable this, check **Disable ACF Compression**.

Protocol Field Compression (ProtoComp)
This option saves one byte per frame for most frames. To disable this, check **Disable Protocol Compression**.

PPPoE has two additional advanced options:

Multilink over single link
When set, the firewall will use LCP multi-link extensions over a single link. This ignores the MTU/MRU settings. Only enable when required by the ISP.

Force MTU
When set, overrides the MTU negotiated with the ISP with a higher value known to work on the link.

> **Warning:** This option violates RFC 1661 and can break connectivity. While it may result in faster speed as larger packets can be transferred, there is no guarantee that it will function in the future if the provider makes changes.

11.7 GRE (Generic Routing Encapsulation)

Generic Routing Encapsulation (GRE) is a method of tunneling traffic between two endpoints without encryption. It can be used to route packets between two locations that are not directly connected, which do not require encryption. It can also be combined with a method of encryption that does not perform its own tunneling.

Note: The GRE protocol was originally designed by Cisco, and it is the default tunneling mode on many of their devices.

GRE tunnels can carry either IPv4, IPv6, or both types of traffic at the same time.

11.7.1 GRE Interface Settings

Parent interface
The interface upon which the GRE tunnel will terminate. Often this will be WAN or a WAN-type connection.

Remote Address
The address of the remote peer. This is the address where the GRE packets will be sent by this firewall; The routable external address at the other end of the tunnel.

Local IPv4/IPv6 Tunnel Address
The **internal** IPv4 and IPv6 address for the end of the tunnel on this firewall. The firewall will use this address for its own traffic in the tunnel, and tunneled remote traffic would be sent to this address by the remote peer.

Remote IPv4/IPv6 Tunnel Address
The IPv4 and IPv6 address used by the firewall **inside** the tunnel to reach the far side. Traffic destined for the other end of the tunnel must use this address as a gateway for routing purposes.

IPv4/IPv6 Tunnel Subnet
The subnet mask for the GRE interface address.

Add Static Route
When set, the firewall adds an explicit static route for the remote inner tunnel address/subnet via the local tunnel address. This can help with reaching the remote subnet in cases where other route table entries may select the wrong path to that destination.

Description
A short description of this GRE tunnel for documentation purposes.

11.7.2 GRE Interface Management

To create or manage a GRE interface:

- Navigate to **Interfaces > Assignments**, **GRE** tab

Note: The items in this list are managed in the usual way. See *Managing Lists in the GUI*.

- Click  **Add** to create a new GRE instance
- Complete the settings as described in *GRE Interface Settings*
- Click **Save**
- Navigate to **Interfaces > Assignments**
- Select the new GRE interface in the **Available network ports** list
- Click  **Add**
- Note the name given to the new interface (e.g. OPT1)
- Navigate to **Interfaces > <name>** where <name> corresponds to the name of the GRE interface (e.g. OPT1)
- Check **Enable interface**
- Enter a new name for the interface in **Description** (optional)
- Click **Save**

Then use the interface as any other WAN-type interface. The firewall automatically creates a dynamic gateway for routing purposes. Depending on the use case, the interface may need NAT or firewall rules, static routes, and so on.

11.8 GIF (Generic tunnel InterFace)

A Generic Tunneling Interface (GIF) is similar to *GRE*; Both protocols are a means to tunnel traffic between two hosts without encryption. In addition to tunneling IPv4 or IPv6 directly, GIF may be used to tunnel IPv6 over IPv4 networks and vice versa. GIF tunnels are commonly used to obtain IPv6 connectivity to a tunnel broker such as Hurricane Electric in locations where IPv6 connectivity is unavailable.

See also:

See *Configuring IPv6 Through A Tunnel Broker Service* for information about connecting to a tunnel broker service.

GIF interfaces carry more information across the tunnel than can be done with GRE, but GIF is not as widely compatible. For example, a GIF tunnel is capable of bridging layer 2 between two locations while GRE cannot.

GIF interfaces can carry IPv4 or IPv6 traffic, but not both at the same time.

Note: GIF compatibility varies by vendor, but is not as common as others like GRE.

See also:

VXLANs are another method for tunneling L2 data without encryption.

11.8.1 GIF Interface Settings

Parent interface
The interface upon which the GIF tunnel will terminate. Often this will be WAN or a WAN-type connection.

GIF Remote Address
The address of the remote peer. This is the address where the GIF packets will be sent by this firewall; The routable external address at the other end of the tunnel. For example, in a IPv6-in-IPv4 tunnel to Hurricane Electric, this would be the **IPv4 address** of the tunnel server, such as `209.51.181.2`.

GIF tunnel local address
The **internal address** for the end of the tunnel on this firewall. The firewall will use this address for its own traffic in the tunnel, and tunneled remote traffic would be sent to this address by the remote peer. For example, when tunneling IPv6-in-IPv4 via Hurricane Electric, they refer to this as the **Client IPv6 Address**.

GIF tunnel remote address
The address used by the firewall **inside** the tunnel to reach the far side. Traffic destined for the other end of the tunnel must use this address as a gateway for routing purposes. For example, when tunneling IPv6-in-IPv4 via Hurricane Electric, they refer to this as the **Server IPv6 Address**.

GIF Tunnel Subnet
The subnet mask or prefix length for the interface address. Typically, `64`. This option is ignored with IPv6 and a `128` prefix is enforced by the kernel instead.

ECN Friendly Behavior
The ECN friendly behavior option controls whether the Explicit Congestion Notification (ECN)-friendly practice of copying the TOS bit into/out of the tunnel traffic is performed by the firewall. By default, the firewall clears the TOS bit on the packets or sets it to `0`, depending on the direction of the traffic. With this option set, the bit is copied as needed between the inner and outer packets to be more friendly with intermediate routers that can perform traffic shaping. This behavior breaks RFC 2893 so it must only be used when both peers agree to enable the option.

Outer Source Filtering
When set, the firewall will not automatically filter based on the outer GIF source. This is normally desirable as it ensures a match with the configured remote peer, which is more secure. When disabled, Martian and inbound filtering is not performed which allows asymmetric routing of the outer traffic. This is less secure, but some GIF peers may source traffic in this manner.

Description
A short description of this GIF tunnel for documentation purposes.

11.8.2 GIF Interface Configuration

To create or manage a GIF interface:

- Navigate to **Interfaces > Assignments**, **GIF** tab

Note: The items in this list are managed in the usual way. See *Managing Lists in the GUI*.

- Click  **Add** to create a new GIF instance
- Complete the settings as described in *GIF Interface Settings*
- Click **Save**
- Navigate to **Interfaces > Assignments**
- Select the new GIF interface in the **Available network ports** list
- Click  **Add**
- Note the name given to the new interface (e.g. OPT1)
- Navigate to **Interfaces > <name>** where <name> corresponds to the name of the GIF interface (e.g. OPT1)
- Check **Enable interface**
- Enter a new name for the interface in **Description** (optional)
- Click **Save**

Then use the interface as any other WAN-type interface. The firewall automatically creates a dynamic gateway for routing purposes. Depending on the use case, the interface may need NAT or firewall rules, static routes, and so on.

11.9 LAGG (Link Aggregation)

Link aggregation is handled by `lagg(4)` type interfaces (LAGG) on pfSense® software. LAGG combines multiple physical interfaces together as one logical interface. There are several ways this can work, either for gaining extra bandwidth, redundancy, or some combination of the two.

Note: LACP will only work across multiple switches if the switches are Stackable.

11.9.1 LAGG Interface Settings

When creating or editing a LAGG interface, the following settings are available:

Parent Interfaces
This list contains all currently unassigned interfaces, plus members of the current LAGG interface when editing an existing instance.

To add interfaces to this LAGG, select one or more interfaces in this list.

Note: An interface may only be added to a LAGG group if it is not assigned. If an interface is not present in the list, it is likely already assigned as an interface.

LAGG Protocol
The operating modes for LAGG interfaces are: LACP, Failover, Load Balance, Round Robin, and None.

LACP
The most commonly used LAGG protocol. This mode can utilize IEEE 802.3ad Link Aggregation Control Protocol (LACP) and the Marker Protocol. In LACP mode, negotiation is performed with the switch – which must also be configured for LACP – to form a group of ports that are all active at the same time. This is known as a Link Aggregation Group, or LAG. The speed and MTU of each port in a LAG must be identical, and the ports must also run at full-duplex. If link is lost to a port on the LAG, the LAG continues to function but at reduced capacity. In this way, an LACP LAGG bundle can gain both redundancy and increased bandwidth.

Traffic is balanced between all ports using the selected **Hash Algorithm**.

In addition to configuring this option on the firewall, the switch must enable LACP on these ports or have the ports bundled into a LAG group. Both sides must agree on the configuration in order for it to work properly.

LACP Timeout Mode controls how often the firewall sends LACP PDUs. An LACP timeout occurs when three consecutive PDUs are missed.

Slow
Default. LACP PDUs are sent every `30` seconds. A timeout occurs after `90` seconds.

Fast
LACP PDUs are sent every second. A timeout occurs after `3` seconds.

Failover
When using the Failover LAGG protocol traffic will only be sent on the *primary* interface of the group. If the primary interface fails, then traffic will use the next available interface.

Note: By default, traffic may only be received by the active interface. Create a system tunable for `net.link.lagg.failover_rx_all` with a value of `1` to allow traffic to be received on every member interface.

Failover mode has one additional option:

Failover Primary Interface
This option sets the primary interface for failover mode, or **auto** to allow the

firewall to select the primary interface automatically. In auto mode, the first selected interface in the list is primary.

Each non-primary interface is eligible for use in failover if the primary fails.

Load Balance
Load Balance mode accepts inbound traffic on any port of the LAGG group and balances outgoing traffic on any active ports in the LAGG group. It is a static setup that does not monitor the link state nor does it negotiate with the switch. Outbound traffic is load balanced based on all active ports in the LAGG using the chosen **Hash Algorithm**.

Round Robin
This mode accepts inbound traffic on any port of the LAGG group and sends outbound traffic using a round-robin scheduling algorithm. Typically, this means that traffic will be sent out in sequence, using each interface in the group in turn.

None
This mode disables traffic on the LAGG interface without disabling the interface itself. The OS will still believe the interface is up and usable, but no traffic will be sent or received on the group.

Hash Algorithm
The *LACP* and *Load Balance* protocols choose how to balance traffic between their interfaces when transmitting data using details from packets checked against the chosen **Hash Algorithm**.

Note: This can only affect traffic transmitted from this system, it has no influence over how the peer sends data or how data received from the peer is processed by the LAGG. That behavior must be configured on the peer.

For a packet to egress via a different interface, it must differ in some way from other packets. This option allows the user to choose which packet properties are checked when differentiating between packets for determining the egress interface.

The options are based on combinations of properties on different *Layers of the OSI model*:

Layer 2
Source/Destination MAC Address and optional VLAN number.

Layer 3
Source/Destination IPv4/IPv6 Address.

Layer 4
Source/Destination port.

When the selected **Hash Algorithm** includes a given layer, packets are checked for any difference in values at that layer.

The default `Layer 2/3/4` selection means that any difference in the VLAN, MAC address, source or destination IP address, or source/destination port on a packet is eligible to flow out a different interface.

Note: In some equipment the default is Layer 2 only meaning that any traffic for a single MAC address would only utilize one interface, which primarily affects links between routers, essentially turning them into failover only LAGGs. If possible, choose additional layers for the hash algorithm to better balance the load.

Description
A short note about the purpose of this LAGG instance.

11.9.2 LAGG Interface Configuration

To create or manage LAGG interfaces:

- Navigate to **Interfaces > Assignments**, **LAGGs** tab
- Click  **Add** to create a new LAGG, or click  to edit an existing instance.
- Complete the settings as described in *LAGG Interface Settings*
- Click **Save**

After creating a LAGG interface, it works like any other physical interface. Assign the lagg interface under **Interfaces > Assignments** and give it an IP address, or build other things on top of it such as VLANs.

Note: If the only purpose of the LAGG interface is to carry VLANs, it does not need to be assigned.

11.9.3 LAGG and Traffic Shaping

Due to limitations in FreeBSD, `lagg(4)` is not compatible with `altq(4)` so it is not possible to use the traffic shaper on LAGG interfaces directly. `vlan(4)` interfaces are compatible with `altq(4)` and VLANs can be used on top of LAGG interfaces, so using VLANs can work around the problem. As an alternate workaround, Limiters can control bandwidth usage on LAGG interfaces.

11.9.4 LAGG Throughput

Using a LAGG does not necessarily guarantee full throughput equal to the sum of all interfaces. In particular, a single flow will not exceed the throughput of a LAGG member interface. Traffic on a LAGG is hashed in such a way that flows between two hosts, such as this firewall and an upstream gateway, would only use a single link since the flow is between a single MAC address on each side.

In networks where many hosts communicate with different MAC addresses, the usage can approach the sum of all interfaces in the LAGG.

11.10 QinQ Configuration

QinQ, also known as IEEE 802.1ad or stacked VLANs, is a means of nesting VLAN tagged traffic inside of packets that are already VLAN tagged, or "double tagging" the traffic.

See also:

- *Virtual LANs (VLANs)*

QinQ is used to move groups of VLANs over a single link containing one outer tag, as can be found on some links between locations from ISPs or datacenters. QinQ can be a quick and easy way of trunking VLANs across locations without having a trunking-capable connection between the sites, provided the infrastructure between the locations does not strip tags from the packets.

11.10.1 QinQ Interface Settings

When creating or editing a QinQ interface entry, the following options are available:

Parent Interface
The interface that will carry the QinQ traffic.

VLAN Tag Type
The VLAN tag protocol identifier (TPID) which specifies the type of VLAN tag in a packet. This can be one of two values:

C-Tag
"Customer" tag, value `0x8100`. When a packet has multiple tags, C-Tags are the inner tag.

S-Tag
"Service" tag, value `0x88a8`. This is the default and most common value for typical 802.1ad QinQ interfaces. In the majority of cases, S-tags are used for the outer tag when double tagging with QinQ.

First level tag
The outer VLAN ID on the QinQ interface, or the VLAN ID given by the provider for the site-to-site link.

Adds interface to QinQ interface groups
When checked, a new interface group will be created called **QinQ** that can be used to filter all the QinQ sub-interfaces at once.

When hundreds or potentially thousands of QinQ tags are present, this greatly reduces the amount of work needed to use the QinQ interfaces

Description
Optional text for reference, used to identify the entry

Member(s)
Member VLAN IDs for QinQ tagging. These can be entered one per row or in ranges such as `100-150`.

Click **Add Tag** to add another line for more tags or ranges.

11.10.2 QinQ Interface Configuration

Setting up QinQ interfaces is fairly simple:

- Navigate to **Interfaces > Assignments**
- Click the **QinQ** tab
- Click **Add** to add a new QinQ entry
- Configure the QinQ entry as described in *QinQ Interface Settings*
- Click **Save** to complete the interface

11.10.3 QinQ Example

In the following example (Figure *QinQ Basic Example*), a QinQ interface is configured to carry tagged traffic for VLANs `10` and `20` across the link on `igc0` with a first level tag of `2000`.

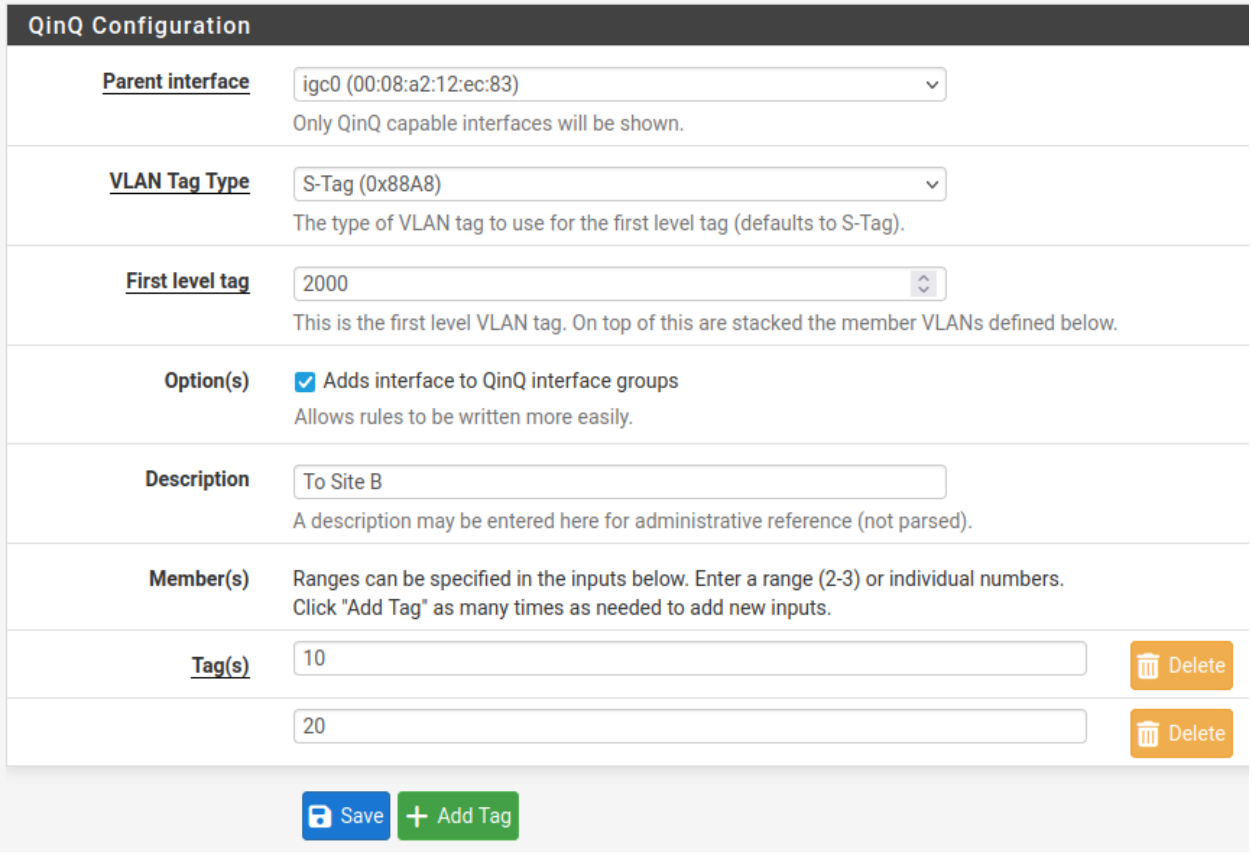


Fig. 3: QinQ Basic Example

In Figure *QinQ List*, this entry is shown on the QinQ tab summary list.

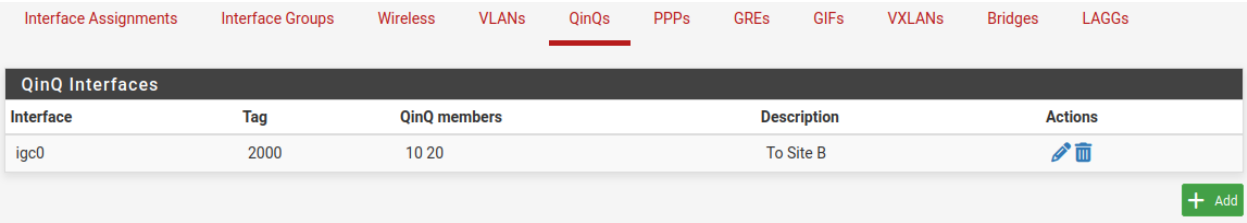


Fig. 4: QinQ List

The automatic interface group, shown in Figure *QinQ Interface Group*, must not be manually edited. Because these interfaces are not assigned, it is not possible to make alterations to the group without breaking it. To re-create the group, delete it from this list and then edit and save the QinQ instance again to add it back.

Rules may be added to the **QinQ** tab under **Firewall > Rules** to pass traffic in both directions across the QinQ links.

From here, how the QinQ interfaces are used is mostly up to the needs of the network. Most likely, the resulting interfaces may be assigned and then configured in some way, or bridged to their local equivalent VLANs (e.g. bridge an assigned `igc1.10` to `igc0.2000.10` and so on).

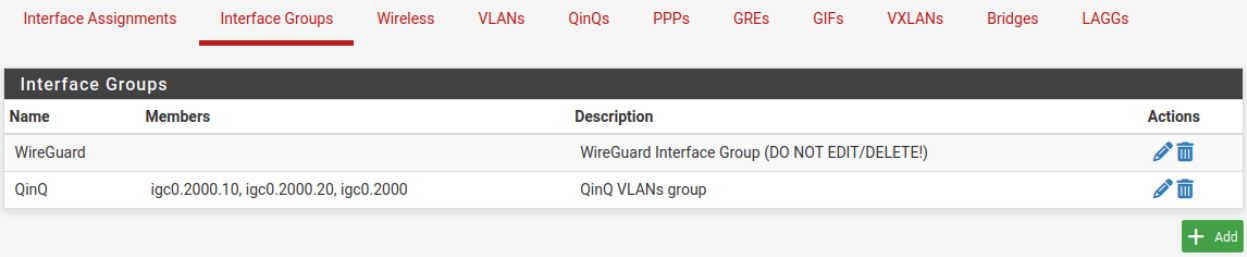


Fig. 5: QinQ Interface Group

The QinQ configuration will be roughly the same on both ends of the setup. For example, if both sides use identical interface configurations, then traffic that leaves Site A out on `igc0.2000.10` will go through VLAN `2000` on `igc0`, come out the other side on VLAN `2000` on `igc0` at Site B, and then in `igc0.2000.10` at Site B.

11.11 Virtual Extensible LAN (VXLAN) Configuration

Virtual Extensible LAN, or VXLAN (RFC 7348), is a tunneling protocol designed to carry L2 information similarly to VLANs, but using UDP as its transport. A VXLAN tunnel is terminated on each end at a VXLAN tunnel endpoint (VTEP), such as this device or a switch compatible with VXLANs.

See also:

GIF (Generic tunnel InterFace) is another tunneling protocol designed to carry L2 information, but it is not as widely implemented as VXLANs.

Warning: VXLANs do not encrypt data they carry. Do not allow VXLAN traffic to pass over untrusted networks.

See also:

- *Virtual LANs (VLANs)*
- *Bridging*

11.11.1 Compatibility

Individual VXLAN instances with unique VXLAN IDs can be bridged with individual VLANs to transport VLANs between two endpoints, provided that the VLAN parent interface itself is not a bridge.

It is not currently possible to bridge an interface containing multiple VLANs onto a single VXLAN instance.

11.11.2 VXLAN Interface Settings

When creating or editing a VXLAN interface, the following options are available:

Parent Interface
: The interface on the firewall upon which the VXLAN tunnel will terminate.

VXLAN Remote Address
: The address of the remote VTEP where this firewall will send VXLAN packets.

VXLAN ID
The 24-bit VXLAN Network Identifier (VNI) that identifies the virtual network segment membership of the interface.

Note: The VXLAN ID must be unique for VXLAN instances between the same endpoints.

VXLAN Local Port
The port number the VXLAN interface will listen on. The default port number is `4789`.

VXLAN Remote Port
The destination port number used in the encapsulating IPv4/IPv6 header. The remote host must accept VXLAN traffic on this port. The default port number is `4789`.

Note: Some implementations listen on port `8472` instead, such as Linux.

VXLAN TTL
The Time To Live (TTL) used in the encapsulating IPv4/IPv6 header. Default value is `64`.

VXLAN Learn
Populate the forwarding table with the source IP address and inner source Ethernet MAC address of received packets.

The source IP address and inner source Ethernet MAC address of received packets are used to dynamically populate the forwarding table. When in multicast mode, an entry in the forwarding table allows the interface to send the frame directly to the remote host instead of broadcasting the frame to the multicast group. This is the default.

Description
A brief text description of this VXLAN tunnel.

11.11.3 VXLAN Interface Configuration

To configure a VXLAN interface:

- Navigate to **Interfaces > Assignments**
- Click the **VXLAN** tab
- Click  **Add** to add a new VXLAN entry
- Configure the VXLAN entry as described in *VXLAN Interface Settings*
- Click **Save** to complete the interface

11.12 Integrated Switches

Certain models of hardware sold by Netgate have integrated switches. These switches can be configured in a variety of ways, with multiple ports on the same network or with each port on a separate network. The default configuration of the switch and the procedure to change that configuration varies by model.

Models with integrated switches include:

- Netgate 7100
- Netgate 3100
- Netgate 2100
- Netgate 1100

See also:

- *Virtual LANs (VLANs)*
- *Bridging*
- *Wireless*

pfSense® software is compatible with numerous types of network interfaces, either using physical interfaces directly or by employing other protocols such as PPP or VLANs.

Interface assignments and the creation of new virtual interfaces are all handled under **Interfaces > Assignments**.

11.13 Physical and Virtual Interfaces

Most interfaces discussed in this chapter can be assigned as WAN, LAN, or an OPT interface under **Interfaces > Assignments**. All currently-defined and detected interfaces are listed directly on **Interfaces > Assignments** or in the list of interfaces available for assignment. By default, this list includes only the physical interfaces, but the other tabs under **Interfaces > Assignments** can create virtual interfaces which can then be assigned.

Interfaces can be configured with various combinations of options. They can also utilize multiple networks and protocols on a single interface, or multiple interfaces can be bound together into a larger capacity or redundant virtual interface.

All interfaces are treated equally; Every interface can be configured for any type of connectivity or role. The default WAN and LAN interfaces can be renamed and used in other ways.

Physical interfaces and virtual interfaces are treated the same once assigned, and have the same capabilities. For example, a VLAN interface can have the same type of configuration that a physical interface can have. Some interface types receive special handling once assigned, which are covered in their respective sections of this chapter.

This section covers the various types of interfaces that can be created, assigned, and managed.

11.14 Switches

Some Netgate Appliances sold in the Netgate Store contain built-in switches which can be configured in the GUI under **Interfaces > Switches**. Documentation for the switch configuration can vary by model, and may be found in the Netgate Product Manuals which match a given product.

11.15 Limitations

While the firewall does not impose any limits on the number of interfaces, large numbers of interfaces may function in suboptimal ways. For example, the firewall may take much longer to configure interfaces and the GUI may have rendering issues with large numbers of tabs or menu entries.

Most hardware will accommodate as many physical interfaces as can fit into the case. Issues may vary from driver to driver but generally are hardware-related and not the result of the operating system or pfSense software.

Note: With numerous physical interfaces, the number of mbufs will likely need to be increased. See *Hardware Tuning and Troubleshooting*.

Physical limitations aside, significant numbers of virtual interfaces such as VLANs, LAGGs, VPNs, and more may be added to the firewall. These types interfaces tend to outnumber physical interfaces, especially VLANs.

Issues reported by users with large numbers of interfaces (physical and virtual) vary by hardware, configuration, and browser. These issues tend to increase as the number of interfaces approaches 200. Should a particular environment require more than 128 interfaces, consider alternate designs that do not involve using all the interfaces on the firewall directly. If the firewall must handle numerous interfaces, be wary of potential performance and GUI concerns.

CHAPTER

TWELVE

USER MANAGEMENT AND AUTHENTICATION

12.1 Default Username and Password

The factory default credentials for a pfSense® software installation are:

Username
`admin`

Password
`pfsense`

> **Warning:** Change the password to a secure value as soon as possible. Do not leave the password at the default value, even in a lab or test environment.

Note: On pfSense Plus software version 24.03 and later, administrators are required to change this password to a custom non-default value the first time they log in.

12.2 Privileges

Managing privileges for users and groups is done similarly, so both will be covered here rather than duplicating the effort. Whether a user or group is managed, the entry must be created and saved first before privileges can be added to the account or group.

To add privileges, edit an existing user or group and click  **Add** in the **Assigned Privileges** or **Effective Privileges** section.

The GUI presents a list of all available privileges. Privileges may be added one at a time by selecting a single entry, or by multi-select using Ctrl-click or cmd-click. If other privileges are already present on the user or group, they are hidden from this list so they cannot be added twice. To search for a specific privilege by name, enter the search term in the **Filter** box and click  **Filter**.

Selecting a privilege will show a short description of its purpose in the information block area under the permission list and action buttons. Most of the privileges are self-explanatory based on their names, but a few notable permissions are:

WebCfg - All Pages
Grants the user access to any page in the GUI

WebCfg - Dashboard (all)
Grants the user access to the dashboard page and all of its associated functions (widgets, graphs, etc.)

WebCfg - System: User Password Manager Page
If the user has access to only this page, they can log in to the GUI to set their own password but do nothing else.

User - VPN - IPsec xauth Dialin
Allows the user to connect and authenticate for IPsec xauth

User - Config - Deny Config Write
Prevents the user from making changes to the firewall configuration (`config.xml`).

> **Warning:** This does not prevent the user from taking other actions that do not involve writing to the configuration.

User - System - Shell account access
Grants the user the ability to login over SSH, though the user will not have root-level access, so functionality is limited. A package for *sudo* is available to enhance this feature.

After login, the firewall will attempt to display the dashboard. If the user does not have access to the dashboard, the GUI will forward the user to the first page in their privilege list to which they have access.

Menus on the firewall only contain entries for which privileges exist on a user account. For example, if the only Diagnostics page that a user has access to is **Diagnostics > Ping** then no other items will be displayed in the **Diagnostics** menu.

12.3 Manage Local Users

The **Users** tab under **System > User Manager** is where individual users are managed.

Note: The admin user cannot be deleted, and its username may not be changed.

12.3.1 Creating and Editing Users

The first step is always to add the user and save. Privileges can only be added to existing users, they cannot be added when creating a new user.

Tip: If multiple users need the same privileges, the most efficient method is to add a *group* and then add users to the group.

To add a new user:

- Navigate to **System > User Manager**
- Click  **Add**

To edit an existing user:

- Navigate to **System > User Manager**
- Click  on the row containing the user

12.3.2 User Settings

When creating or editing a user, the following options are available:

Disabled
This checkbox controls whether this user will be active. To deactivate this account, check the option.

Username
Sets the login name for the user. This field is required, must be 16 characters or fewer and may only contain letters, numbers, and a period, hyphen, or underscore.

Password / Confirm Password
The password for this user. Ensure the two fields match to confirm the password.

This password cannot be set to the same value as the username. Additionally, on pfSense Plus software version 24.03 and later, the password cannot be set to the default value (*Default Username and Password*).

Note: Passwords are stored in the configuration as salted hashes, not plain text.

Tip: GUI users can also change their own password using the *User Password Manager* page.

Full Name
Optional field which can be used to enter a longer name or a description for this user account.

Expiration Date
Optional date at which the firewall will automatically deactivate this user account. The date must be entered in `MM/DD/YYYY` format.

Custom Settings
Enables options for per-user custom GUI settings. See *Per-user GUI Options and Dashboard Layout* for details.

Group Memberships
If one or more groups exist on the firewall (*Manage Local Groups*), this control can add the user as a member.

To add a group for this user:

- Click the group name in the **Not Member Of** column
- Click  to move it to the **Member Of** column

To remove a group from the user:

- Click the group name in the **Member Of** column
- Click  to move it to the **Not Member Of** column

Effective Privileges

A list of privileges this user has, either directly assigned or inherited by group membership.

Appears only when editing an existing user, not when creating a user.

Privileges assigned to the user may be edited by these controls, but group privileges cannot. Group privileges must be managed on the *group*.

See also:

See *Privileges* for information on managing privileges.

Certificate

Certificates associated with this user account.

The behavior of this section changes depending on whether the page is creating a new user or editing an existing user. This section is disabled if there are no internal certificate authorities defined on the firewall capable of signing a certificate.

To create a certificate while adding a user:

- Check **Click to create a user certificate**
- Fill in the **Descriptive name**
- Choose a **Certificate Authority**
- Select a **Key Type** and **Key Length**
- Select a **Digest Algorithm**
- Enter a **Lifetime**

See also:

For more information on these parameters, see *Create an Internal Certificate*.

When editing a user, this section of the page instead becomes a list of certificates associated with this user account.

To create a certificate for an existing user:

- Click  **Add**
- Fill in the settings on the page as described in *Create an Internal Certificate* (some data is pre-filled)

To associate an existing certificate with this user:

- Set **Method** to *Choose an Existing Certificate*
- Select an entry from the **Existing Certificate** list
- Click **Save**

Authorized SSH keys

Public keys for SSH and SCP authentication.

To add a key, paste or enter in the key data. Multiple keys are allowed, one per line.

> **Warning:** Only enter authorized keys into this field. Do not add them to files in user home directories. Those files will be overwritten by the GUI the next time account information is synchronized to disk (e.g. at boot time).

IPsec Pre-Shared Key
Pre-Shared Key (PSK) for this user to connect to a non-xauth Pre-Shared Key mobile IPsec setup.

If a PSK is entered here, the username is used as the identifier. The PSK is also displayed under **VPN > IPsec** on the **Pre-Shared Keys** tab.

Note: This field has no effect for IKEv2 or xauth mobile IPsec.

Keep Command History
If this user has shell access, this option preserves the last 1000 unique commands entered at a shell prompt between login sessions. The user can access history using the up and down arrows at an SSH or console shell prompt and search the history by typing a partial command and then using the up or down arrows.

Per-user GUI Options and Dashboard Layout

Each user can have their own settings for various GUI options and their dashboard layout. To enable this for a user, check the **Custom Settings** box when adding or editing the user.

When that option is active, additional GUI options for the user are present on the user account page. Additionally, the user can have their own personal dashboard layout, starting from the system-wide layout.

Choose the other GUI options desired for the user such as theme, top navigation, host name in menu, dashboard columns, show/hide associated panels, left column labels and browser tab text.

Tip: Users with the **WebCfg - System: User Settings** privilege may adjust their own GUI options.

Users in the `admin` group already have this privilege.

A user with **Custom Settings** enabled and the **User Settings** privilege will have menu option **System > User Settings**. The user can select this to change the GUI options for their account.

When a user with **Custom Settings** adds, moves or removes dashboard widgets, the custom dashboard layout is saved in the preferences for only that user.

12.4 Manage Local Groups

Groups manage sets of user privileges so they do not need to be maintained individually on every user account. For example, a group can be used for IPsec xauth users, or a group that can access the firewall dashboard, a group of firewall administrators, or many other possible scenarios using any combination of privileges.

Groups are managed under **System > User Manager** on the **Groups** tab.

Note: The `all` and `admins` groups cannot be deleted.

12.4.1 Groups and Remote Authentication

When working with group privileges while authenticating against LDAP and RADIUS (*Authentication Servers*), local groups must exist with names that exactly match groups from the server. For example, if an LDAP group named `firewall_admins` exists then the firewall must also contain an identically named group, `firewall_admins`, with the desired privileges.

If a user attempts to authenticate against a remote authentication server and there are no matching groups, the user will not have any privileges from groups, and cannot access resources which require privileges.

12.4.2 Creating and Editing Groups

As with users, the first step is to add the group and save. Privileges can only be added to existing groups, they cannot be added when creating a new group.

To add a new group:

- Navigate to **System > User Manager**, **Groups** tab
- Click  **Add**

To edit an existing group:

- Navigate to **System > User Manager**, **Groups** tab
- Click  on the row containing the group

12.4.3 Group Settings

Group name
The name of the group.

For groups in the *Local* scope, this setting has the same restrictions as a username: It must be 16 characters or fewer and may only contain letters, numbers, and a period, hyphen, or underscore.

Groups in the *Remote* scope do not have strict name restrictions, for example they may have longer names.

Scope
The scope in which this group is available for use.

Note: LDAP and RADIUS groups can match names in both local and remote scopes.

Local
Groups on the firewall itself, such as those for use in the shell, filesystem, and other local uses. These groups are added to the operating system, so they are subject to naming restrictions imposed there.

Remote
Groups from remote sources, such as authentication servers (RADIUS or LDAP). These groups are not exposed to the operating system, and thus are only available for use in the GUI and other similar uses not involving the operating system layer. This scope has relaxed name restrictions, for example, group names may be longer and may contain spaces.

Description
Optional free-form text for reference and to better identify the purpose of the group in case the **Group name** is not sufficient.

Group Memberships
This set of controls defines which existing users will be members of the new group. Firewall users are listed in the **Not Members** column by default.

To add a user to this group:

- Click the username in the **Not Members** column
- Click  to move it to the **Members** column

To remove a user from this group:

- Click the username in the **Members** column
- Click  to move it to the **Not Members** column

Assigned Privileges
A list of privileges assigned to this group. Appears only when editing an existing group.

See also:

See *Privileges* earlier in this for information on managing privileges.

12.5 Settings

The **Settings** tab in the User Manager controls how the firewall authenticates users for the GUI and SSH.

Session Timeout
This field specifies how long a GUI login session will last when *idle*. This value is specified in **minutes**, and the default is four hours (`240` minutes). A value of `0` may be entered to disable session expiration, making the login sessions valid forever. A shorter timeout is better, though it should be long enough that an active administrator would not be logged out unintentionally while making changes.

> **Warning:** Allowing a session to stay valid when idle for long periods of time is insecure. If an administrator leaves a terminal unattended with a browser window open and logged in, someone or something else could take advantage of the open session.

Authentication Server
This selector chooses the primary authentication source for users logging into the GUI. This can be a RADIUS or LDAP server, or the default *Local Database*.

Note: Authentication falls back to *Local Database* if the RADIUS or LDAP server is unreachable, returns an authentication failure, or otherwise results in an error, even if another method is chosen.

This ensures that an administrator can always access the device, even if the authentication server is broken.

Password Hash Algorithm
Selects which algorithm the firewall will use when creating hashes for passwords in user manager accounts.

May be one of the following choices:

bcrypt - Blowfish-based crypt
Secure password hashing with a crypt algorithm based on Blowfish. The most secure option currently available.

Note: This hashing algorithm is restricted to a maximum password length of 72 characters.

SHA-512 - SHA-512-based crypt
Secure password hashing with a crypt algorithm based on SHA-512. Weaker than bcrypt but still has an acceptable level of security in many environments.

Some users may prefer SHA-512-based crypt hashes for compatibility or compliance purposes.

Shell Authentication
When set, the selected **Authentication Server** will also be configured as the authentication source for SSH access to the firewall. By default, only accounts in the *User Manager* with shell privileges can log in over SSH.

This works with both RADIUS and LDAP servers, with some caveats:

RADIUS Servers
When used with a RADIUS server, accounts must exist on the firewall with the same names and the expected privileges. They will authenticate against RADIUS but use the local accounts settings otherwise.

LDAP Servers
When used with an LDAP server, the **Shell Authentication Group DN** must be set on the *LDAP Authentication Server* entry. Users must be a member of that group and have valid `posixAccount` attributes in their LDAP account.

Auth Refresh Time
Time in seconds for which the firewall cache authentication results. The default is `30` seconds, maximum `3600` (one hour). Shorter times result in more frequent queries to authentication servers.

The firewall periodically re-authenticates users against the remote server to ensure the account is still valid and has the expected privileges. Checking frequently is more secure, but puts a larger burden on the authentication server and can increase page load times on the firewall.

12.5.1 Remote Authentication Servers and Privileges

When using a RADIUS or LDAP server to authentication for the GUI, the *users* and/or *group memberships* must be defined in the firewall in order to properly allocate permissions, as there is no method to obtain permissions dynamically from an authentication server.

For group membership to work properly, the firewall must be able to recognize the groups as presented by the authentication server. This requires two things:

- The local groups must exist with identical names (*Manage Local Groups*).
- The firewall must be able to locate or receive a list of groups from the authentication server.

See *Authentication Servers* for details specific to each type of authentication server.

12.6 User Password Manager

The self-service user password manager page (**System > User Password Manager** or **System > User Manager**, **Change Password** tab) allows a user to change their own password in the User Manager Local Database.

This page is separate from the User Manager so that privileges can be granted to access this page without giving a user access to the user manager as a whole. The privilege governing this page is labeled "WebCfg - System: User Password Manager".

The page prompts the user to enter and confirm a new password. Saving the settings changes the password.

Warning: This password cannot be set to the same value as the username. Additionally, on pfSense Plus software version 24.03 and later, the password cannot be set to the default value (*Default Username and Password*).

Note: Passwords are stored in the configuration as salted hashes, not plain text.

12.7 Authentication Servers

The firewall can use RADIUS and LDAP servers to authenticate users from remote sources.

User Manager Integration contains information on which areas of the firewall can utilize these servers.

To add a new server:

- Navigate to **System > User Manager**, **Authentication Servers** tab
- Click  **Add**

To edit an existing server, click  next to its entry on the same page.

Each type of authentication server is covered in the following documents

12.7.1 RADIUS Authentication Servers

Remote Authentication Dial-In User Service (RADIUS) is a protocol compatible with a wide variety of networking equipment for user authentication, authorization, and accounting (AAA).

Servers are commonly available as well, including *FreeRADIUS* and *Active Directory via NPS*.

Though most areas on pfSense® software which are capable of RADIUS authentication now integrate their RADIUS settings via the user manager, a few remain which use separate settings, such as the PPPoE and L2TP servers.

See also:

- *Controlling Client Parameters via RADIUS*

Warning: Secure the link between the firewall and the RADIUS server. If the server is local, use a trusted management network. If the server is remote, communicate only over VPN tunnels.

Some RADIUS protocols transmit passwords in plain text, and though others attempt to protect the password in other ways, other aspects of the protocol are not encrypted and may contain sensitive information.

RADIUS Configuration

Descriptive name
The name for this RADIUS server. This name will be used to identify the server throughout the GUI.

Protocol
The protocol used by the firewall when performing RADIUS requests. May be one of:

PAP
Password Authentication Protocol. Sends passwords unencrypted, and is considered weak. It is more widely compatible than other methods, and may be required by specific features (e.g. mOTP).

Warning: Due to its security deficiencies, avoid using PAP where possible.

MD5-CHAP
Challenge-Handshake Authentication Protocol using MD5 hashing. The RADIUS server sends a challenge value and the client responds with a hash of the challenge value and the password together. More secure than PAP as it does not transmit passwords in the clear, but both parties must know the plain text of the password.

MS-CHAPv1
A Microsoft variation of CHAP where neither side needs to know the plain text of the password. Though it is generally more secure, it has other known weaknesses which make it vulnerable to attack.

MS-CHAPv2
An updated variation of MS-CHAPv1. It is used in EAP as well as 802.1x/WPA Enterprise for wireless. However, it also has known weaknesses.

Note: Certain RADIUS features may require specific modes. For example, mOTP typically requires PAP since it reads the password in the clear to separate the PIN and OTP code. Services utilizing EAP typically use MS-CHAPv2.

Hostname or IP address
The address of the RADIUS server. This can be a fully qualified domain name or an IPv4 IP address.

Shared Secret
The password established for this firewall *on the RADIUS server* software.

Services offered
This selector sets which services are offered by this RADIUS server.

Authentication
The firewall will use this RADIUS server to authenticate users.

Accounting
The firewall will send RADIUS start/stop accounting packet data for login sessions when possible.

Authentication and Accounting
The server will be used for both types of actions.

Authentication port
Only appears if an Authentication mode is chosen. Sets the UDP port where RADIUS authentication will occur. The default RADIUS authentication port is `1812`.

Accounting port
Only appears if an Accounting mode is chosen. Sets the UDP port where RADIUS accounting will occur. The default RADIUS accounting port is `1813`.

Authentication Timeout
Controls how long, in seconds, that the RADIUS server may take to respond to an authentication request. If left blank, the default value is `5` seconds. If an interactive two-factor authentication system is in use, increase this timeout to account for how long it will take the user to receive and enter a token, which can be 60-120 seconds or more if it must wait for an external action such as a phone call, SMS message, etc.

Note: The system will retry authentication **three** times before giving up, and the timeout applies to each attempt individually. Thus, authentication may take up to 3x this value to terminate if the server is unreachable.

RADIUS NAS IP Attribute
Sets the value the firewall will send in the RADIUS request `NAS-IP-Address` attribute. This value is used by the RADIUS server to identify this firewall. The server can use this value to make authentication decisions, or to denote which node users were authenticated by in accounting data.

In most cases, the `NAS-IP-Address` value does not matter so long as it is unique to this firewall. However, more complicated RADIUS environments may use this attribute to let the server make more informed decisions about users logging into different services. For example, if there are multiple Captive Portal instances on the firewall, multiple RADIUS server entries can be created, each using the specific interface address for a given portal. The RADIUS server could then choose to only let certain sets of users login to each portal.

Adding a RADIUS Server

To add a new RADIUS server:

- Add the firewall as a client on the RADIUS server
- Navigate to **System > User Manager**, **Authentication Servers** tab
- Click  Add
- Set the **Type** selector to *RADIUS*

  The GUI will change the form to display RADIUS Server Settings
- Fill in the fields as described in *RADIUS Configuration*
- Click **Save** to create the server
- Navigate to **Diagnostics > Authentication** to test the RADIUS server using a valid account.

RADIUS Groups

There are two requirements for RADIUS groups to function properly:

- The RADIUS server must return a list of groups in the `Class` RADIUS reply attribute as a string.
- The same groups must exist locally (*Manage Local Groups*)

Multiple groups returned by the RADIUS server in the `Class` attribute must be separated by a semicolon. For example, in FreeRADIUS, to return the `admins` and `VPNUsers` groups, use the following Reply-Item RADIUS Attribute:

```
Class := "admins;VPNUsers"
```

If the RADIUS server returns the group list properly for a user, and the groups exist locally, then the groups will be listed on the results when using the **Diagnostics > Authentication** page to test an account.

If the groups do not show up when testing, ensure the groups exist in the *Group Manager* with matching names and that the server is returning the `Class` attribute as a string, not binary.

12.7.2 LDAP Authentication Servers

Though Lightweight Directory Access Protocol (LDAP) is technically a repository for user information, it also offers mechanisms for user authentication via bind operations.

There are many popular user directory implementations which use LDAP, including Active Directory, OpenLDAP, FreeIPA, and more.

Note: LDAP server implementations and schemas vary widely. As such, there are no complete and specific examples in this document.

LDAP Configuration

Hostname or IP address
The address of the LDAP server. This can be a fully qualified domain name, an IPv4 IP address, or an IPv6 IP address.

Note: If this LDAP server uses SSL, the value of this field **must** match the certificate presented by the LDAP server. Typically, this means it must be a hostname which resolves to the IP address of the LDAP server, but the specific requirements depend on the contents of the server certificate.

For example, with a value of `ldap.example.com` in this field, the server certificate must include an FQDN value of `ldap.example.com`, and `ldap.example.com` must resolve to `192.168.1.5`. One exception to this is if the IP address of the server also happens to be the listed in the server certificate.

This can be worked around in some cases by creating a DNS host override to make the server certificate hostname resolve to the correct IP address if they do not match in this network infrastructure, and they cannot be easily fixed.

Port value
This setting specifies the port on which the LDAP server is listening for LDAP queries. The default port is `389` for Standard TCP and STARTTLS, and `636` for SSL. This field is updated automatically with the proper default value based on the selected **Transport**.

Note: When using port `636` for SSL, the firewall uses an `ldaps://` URL, not STARTTLS. Ensure that the LDAP server is listening on the correct port with the correct mode.

Transport
This setting controls which transport method will be used by the firewall to communicate with the LDAP server.

Warning: LDAP queries will contain sensitive data, such as usernames, passwords, and other information about the user. The best practice is for the firewall to use encryption when communicating with the LDAP server, if the LDAP server is configured appropriately. Both SSL/TLS and STARTTLS will encrypt traffic between the firewall and the LDAP server.

Standard TCP
(Default) Plain unencrypted TCP connections on port `389`. This is not secure, but is the most compatible method and also useful for debugging with packet captures. Do not use this method across untrusted networks.

STARTTLS Encrypted
Connects using TCP port `389` but negotiates encryption with the server using STARTTLS.

Note: Not all LDAP servers accept STARTTLS connections, check the LDAP server documentation and configuration.

SSL/TLS Encrypted
Connects using SSL/TLS on TCP port `636` to encrypt LDAP queries.

Note: Not all LDAP servers offer SSL/TLS encryption for connections, check the LDAP server documentation and configuration.

Peer Certificate Authority
The CA chosen with this selector is used by the firewall to validate the LDAP server certificate when **Transport** is set to **SSL/TLS Encrypted** or **STARTTLS Encrypted** mode.

The selected CA must match the CA which signed the LDAP server certificate, otherwise validation will fail. If the LDAP server is using a globally trusted certificate (e.g. Let's Encrypt or another public CA), choose *Global Root CA List*.

See *Certificate Authority Management* for more information on creating or importing CAs.

Client Certificate
(Plus only) This certificate is sent to the LDAP server to identify this client when using an encrypted transport mode. If the LDAP server requires a client certificate, the server will use this certificate to ensure that the firewall is authorized to make LDAP queries.

This certificate must be issued by the CA used by the LDAP server to validate connecting clients.

Protocol version
Chooses which version of the LDAP protocol is employed by the LDAP server, either *2* or *3*, typically *3*.

Server Timeout
The time, in seconds, after which LDAP operations are considered as failed. Using a lower value will

allow the GUI to try other authentication sources faster when the server fails. If the LDAP server is slow or overloaded, a larger value can help the firewall accept delayed responses.

Search scope
Determines where, and how deep, an LDAP search will be performed to locate a match.

Level
Controls the depth of the LDAP search.

One Level
Search only one level, defined by the **Authentication Containers**.

Entire Subtree
Search the entire subtree of the directory, starting with the **Authentication Containers**.

Tip: This is typically the best choice, and is nearly always required for Active Directory configurations.

Base DN
Controls where the search will start. Typically set to the root of the LDAP structure, e.g. `DC=example,DC=com`

Authentication containers
A list of potential account locations or containers, separated by semicolons. These containers will be prepended to the **Base DN** above when the firewall crafts LDAP queries. Alternately, specify a full container path here and leave the **Base DN** blank.

Tip: If the LDAP server allows it, and the bind settings are correct, click  **Select a container** to browse the LDAP server and select containers from a list.

Some examples of containers are:

- `CN=Users;DC=example;DC=com` This searches for users inside the domain component `example.com`, a common syntax for Active Directory
- `CN=Users,DC=example,DC=com;OU=OtherUsers,DC=example,DC=com` This searches in two different locations, the second of which is restricted to the `OtherUsers` organizational unit.

Extended Query
Specifies an extra restriction to query after the username, which allows group membership to be used as a filter. This must include both the item to search and the method of searching. For example, a restriction based on group membership would use `memberOf`. Check the LDAP server documentation for information on forming such queries.

To set an extended query, check the box and fill in the **Query** value with a filter such as:

```
memberOf=CN=VPNUsers,CN=Users,DC=example,DC=com
```

For users of RFC2307 groups, such as with OpenLDAP, an extended filter might look more like the following:

```
&(objectClass=posixGroup)(cn=VPNUsers)(memberUid=*)
```

Bind credentials
Controls how this LDAP client will attempt to bind to the server.

Note: Active Directory typically requires the use of bind credentials and may need a service account or administrator-equivalent depending on the server configuration. Consult Windows documentation to determine which is necessary in a specific environment.

Bind Anonymous
(Default) When checked the firewall will use anonymous binds. When unchecked the GUI presents the **Bind Credentials** fields.

Bind Credentials (User DN/Password)
When **Bind Anonymous** is unchecked, the credentials in these fields are used by the firewall to make authenticated binds when performing a query.

The **User DN** may be a username or a full DN, depending on what the LDAP server requires.

Attributes

Initial Template
This option only appears when initially creating an LDAP server entry. It pre-fills the remaining options on the page with common defaults for a given type of LDAP server. The choices include *OpenLDAP*, *Microsoft AD*, and *Novell eDirectory*.

User naming attribute
The attribute used to identify the name of a user, most commonly `cn` or `samAccountName`.

Group naming attribute
The attribute used to identify a group, such as `cn`.

Group member attribute
The attribute of a user that signifies it is the member of a group, such as `member`, `memberUid`, `memberOf`, or `uniqueMember`.

RFC2307 Groups
Specifies how group membership is organized on the LDAP server. When unset (default), the queries assume the server uses Active Directory style group membership (RFC 2307bis) where groups are listed as an attribute of the user object. When checked, queries use RFC 2307 style group membership where the users are listed as members on the group object.

Note: In this mode the `Group member attribute` will typically be set to `memberUid`, but may vary by LDAP schema.

RFC2307 User DN
When set, queries include the user DN when searching for groups.

Group Object Class
Specifies the object class of RFC 2307 style groups. Typically, `posixGroup` but it may vary by LDAP schema. Not necessary for Active Directory style groups.

Shell Authentication Group DN
The LDAP group DN for users allowed to log in via SSH. This is used with the **Shell Authentication** option on the *Settings* tab to allow LDAP users to login via SSH.

To login via SSH, users must be a member of this group and have valid `posixAccount` attributes in their LDAP account.

UTF8 Encode
When checked, queries to the LDAP server are encoded for UTF-8 and the responses are decoded from UTF-8. Capabilities vary depending on the LDAP server. Generally only necessary if usernames, groups, passwords, and other attributes contain UTF-8 or international style accented characters.

Username Alterations
When unchecked, a username given as `user@hostname` will have the `@hostname` portion stripped, so only the username is sent in the LDAP bind request. When checked, the username is sent in full.

Allow Unauthenticated Bind
When set, bind requests with empty passwords will be rejected locally. Some LDAP servers, specifically Microsoft Active Directory, will accept unauthenticated bind requests and treat them as successful.

> **Warning:** This behavior must be disabled on the LDAP server where possible. Allowing requests to succeed with an empty password is a significant security risk, and it affects any device or service authenticating against an LDAP server.
>
> Though this option allows the firewall to reject such authentication attempts, other LDAP clients may not offer the same choice. Disabling the feature on the server is the most secure means of correcting the problem. Consult the LDAP server documentation for information on disabling this behavior.

Adding an LDAP Server

To add a new LDAP server:

- Make sure that the LDAP server can be reached by the firewall
- Import the Certificate Authority used by the LDAP server before proceeding if using SSL/TLS or STARTTLS encryption

  See *Certificate Authority Management* for more information on creating or importing CAs.
- Navigate to **System > User Manager**, **Authentication Servers** tab
- Click  **Add**
- Set the **Type** selector to *LDAP*

  The GUI will change the form to display LDAP server settings
- Fill in the fields as described previously in *LDAP Configuration*
- Click **Save** to create the server
- Visit **Diagnostics > Authentication** to test the LDAP server using a valid account

Tip: The debug option on **Diagnostics > Authentication** writes messages to the system log containing a lot of information about LDAP queries and results, which can be of great assistance when testing LDAP server entries.

Multiple Entries to the Same Server

It's not only possible, but useful, to have multiple entries defined for the same LDAP server with different filtering for separate purposes. For example:

- An entry without any filtering for general authentication with areas that recognize groups and the privilege system or where groups do not matter (e.g. GUI administrative access, Captive Portal)
- An entry with an extended query limiting to a single group for VPN access (e.g. OpenVPN for employees) – This way the VPN does not need to check group membership, the authentication will only succeed for members of the LDAP group.
- An entry with an extended query limiting to a different single group for a special purpose VPN (e.g. Remote vendor access, Management VPN)

Name the entries appropriately and choose them for their respective intended purposes.

LDAP Groups

There are two requirements for LDAP groups to function properly:

- The LDAP authentication settings must match the group membership style used by the LDAP server
- The same groups must exist locally (*Manage Local Groups*)

If the LDAP query returns the group list properly for a user, and the groups exist locally, then the groups will be listed on the results when using the **Diagnostics > Authentication** page to test an account.

If the groups do not show up, ensure they exist in the *Group Manager* with matching names and that the proper group structure is present on the LDAP authentication server entry (e.g. RFC 2703 options.)

See also:

- Hangouts Archive to view the August 2015 Hangout on RADIUS and LDAP.
- *External User Authentication Examples*

12.8 User Manager Shell Commands

There are two commands available in a shell (console or SSH) to interact with the User Manager:

usermgrwhoami
: Prints information about the current user from the User Manager database.

usermgrpasswd
: Allows `admin` or `root` to change the password for accounts in the User Manager database.

  When run without any parameters, the script changes the password for the current user (`admin`).

  Other parameters include:

  -c, --check
  : Checks the password for the user to see if it matches known problematic values (e.g. the default value or the username).

  -u <name>, --username <name>
  : Passes a specific username to use when checking or changing the password.

Note: Users other than `admin` and `root` cannot use this shell command because they do not have direct write access to the firewall configuration. Those users can visit the *User Password Manager* page to change their own password.

12.9 Logging Out of the GUI

To end a GUI login session navigate to **System > Logout** or close the browser window.

Sessions will automatically expire if they are idle for longer than the **Session Timeout** defined on **System > User Manager**, **Settings** tab. The default session timeout is 4 hours (240 minutes) of idle time.

See also:

- *Sudo Package*
- *External User Authentication Examples*
- *Granting Users Access to SSH*
- *Accessing Files on the Firewall with SCP*
- *Authenticating Users with Google Cloud Identity*
- *Troubleshooting Authentication*
- *Troubleshooting Access when Locked Out of the Firewall*

The User Manager in pfSense® software provides the ability to create and manage multiple user accounts. These accounts can be used to access the GUI, use VPN services like IPsec and OpenVPN, and use the Captive Portal.

The User Manager is located at **System > User Manager**. From there users, groups, servers may be managed, and settings that govern the behavior of the User Manager may be changed.

The User Manager can also be used to define external authentication sources such as RADIUS and LDAP.

See also:

Hangouts Archive to view the February 2015 Hangout on User Management and Privileges, and the August 2015 Hangout on RADIUS and LDAP.

12.10 User Manager Integration

As of this writing, not all areas of the firewall hook back into the User Manager.

GUI
: Can authenticate users from the User Manager, and via RADIUS or LDAP. Groups or Users from RADIUS or LDAP require definitions in the local User Manager to manage their access permissions.

XMLRPC Configuration Synchronization
: Can authenticate users from the User Manager, and via RADIUS or LDAP. Requires special privilege granted to users or groups.

SSH/SCP
: Can authenticate users from the User Manager, and via RADIUS or LDAP. Requires special privilege granted to users or groups.

IPsec
Can authenticate users in the User Manager, RADIUS or LDAP via User Manager for Xauth, and RADIUS for IKEv2 with EAP-RADIUS.

OpenVPN
Can authenticate users in the User Manager, RADIUS or LDAP via User Manager.

Captive Portal
Can authenticate local users, RADIUS, or LDAP via User Manager.

L2TP
Can authenticate users in the L2TP settings, and via RADIUS in the L2TP settings.

PPPoE Server
Can authenticate users in the PPPoE settings, and via RADIUS in the PPPoE settings.

CHAPTER

THIRTEEN

CERTIFICATE MANAGEMENT

13.1 Certificate Properties

Certificate authority and certificate entries have several properties in common. The common properties of both types are covered here.

13.1.1 Keys

The public and private keys of the certificate are used for cryptographic operations.

Key Type
Certificate key type can be either **RSA** or **ECDSA** (Elliptic Curve Digital Signature Algorithm).

RSA
RSA keys are more widely-compatible than ECDSA, as well as having some performance benefits.

Key Length
When using RSA keys, the security is proportional to the key size. Larger keys are more secure, but they also take longer to generate and are slower to use. RSA performance decreases rapidly as the key size increases.

The best practice is to not use keys smaller than `2048` bits where possible. Larger keys may not be compatible with legacy and embedded systems.

ECDSA
ECDSA is a newer method, and is not as widely adopted. Its main advantage is that is can use smaller keys to provide equivalent levels of security to RSA. ECDSA is slower at verifying signatures than RSA, but scales better.

Curve Name
There are a variety of ECDSA curves available, but only a few have been confirmed to work with various services on the firewall. The services compatible with each curve are noted in the list. Pick the curve based on which services will use this certificate authority or certificate.

13.1.2 Digest Algorithm

Digest Algorithms, also known as Message Digest Algorithms and Hash Algorithms, are used to create a fixed-length hash of content for signing.

The larger the hash, the stronger it is and the less likely it is to be susceptible to collisions which compromise the integrity of the hash. The current best practice is to use a minimum of SHA-256.

Warning: Though the GUI still contains options for SHA-1, it is considered weak and should not be used. Rare exceptions can be made for legacy systems which are incompatible with stronger hashes.

13.1.3 Lifetime

The Lifetime of a certificate authority or certificate determines the length, in days, for which the certificate is valid. Shorter lifetimes are more secure, but require more work as the certificates must be renewed or replaced more frequently.

See also:

Renew or Reissue a CA or Certificate.

For certificate authorities, a longer lifetime such as `3650` days (10 years) is acceptable.

Certificates for users typically also have a long lifetime, but specific values depend largely on the needs of an organization. The GUI defaults to `3650` days for User Certificates, but a better practice is to use a lower value when practical.

Server certificates have stricter requirements for their lifetime. The current accepted maximum lifetime for server certificates is `398` days. Most browsers and other software will no longer accept new server certificates with longer lifetimes.

Note: Another special case is server certificates obtained using *ACME from Let's Encrypt*. These only have a lifetime of `90` days, but since they are automatically replaced well before they expire, there is little extra administrative overhead once the initial setup is complete.

13.1.4 Distinguished Name

The entity to which a certificate authority or certificate belongs, also known as the Subject, is identified by the unique components of the certificate. The primary component for this purpose is the Distinguished Name (DN). These are typically filled in with an organization's information, or in the case of an individual, personal information. This information is mostly cosmetic, and used to verify the accuracy of the CA, and to distinguish one CA from another.

A DN is composed of several fields which contain information about the subject.

Only the **Common Name** is required, the other fields may be left blank.

Warning: A DN with less unique information has the potential to be misidentified later when comparing certificate subjects. Always fill in enough information to uniquely identify the subject.

Common Name
: A short name, such as a username or hostname. Do not use spaces or punctuation, other than that which is typically found in a hostname.

Note: This name is not used directly for certificate validation on modern systems, which look at *Subject Alternative Name* values instead.

Country Code
The two-letter ISO country code for the certificate subject location.

Note: The ISO country code is not the same as the hostname TLD code for a country.

State or Province
The geographical state or province name for the certificate subject location. This value should be spelled out, not using an abbreviation or code.

City
The city for the certificate subject location.

Organization
The name of the organization to which the subject belongs. For example, a company name, government agency name, or similar.

Organizational Unit
A division or department inside the organization, if any. For example, "IT Department" or "Accounting".

Note: When creating a certificate, the GUI populates most of these fields with the values from the certificate authority chosen for signing. The contents of the fields may be changed before performing the signing operation.

13.1.5 Subject Alternative Name

The Subject Alternative Name (SAN) list is only present on certificates. It contains information used to validate the identity of the certificate. For example, when connecting to a device on the network, a system may compare the hostname or IP address to which it connected with values in the certificate SAN list. This way, it can be sure it is communicating with the intended host and not an impostor.

Note: The **Common Name** value from a certificate is automatically added to the SAN list internally, as its inclusion is a requirement of current standards.

The following types of SAN entries can be added to a certificate:

FQDN or Hostname
A fully qualified domain name (e.g. `host.domain.tld`) or a hostname (`host`). In most cases this hostname would also exist in DNS. In the case of user certificates, this could also be a username.

IP Address
An IP address (e.g. `x.x.x.x`), typically an address found on a network device using this certificate. Necessary for clients to properly validate the certificate when connecting by IP address instead of by hostname.

URI
A Uniform Resource Identifier for the certificate subject. In practice, only used as an alternate way to determine the hostname when communicating with servers. It does not restrict certificate validity to specific URIs on a server.

E-mail Address
An e-mail address for the certificate subject.

13.1.6 Certificate Properties in Lists

When viewing the lists of CA and certificate entries, the properties of the entry are available in the **Distinguished Name** column. The DN is printed there, and additional detailed information is available from the  icon.

Underneath that information, the GUI prints the start date and end date for the validity of the entry. The difference between the start and end date is the **Lifetime**. When an entry is nearing expiration, the GUI highlights the end date in yellow. When an entry is expired, it is red. The system also generates notifications for expiring certificates.

See also:

The certificate expiration warning threshold is `27` days by default, but can be customized. See *Notifications* for details.

13.2 Certificate Authority Management

Certificate Authority (CA) entries are managed from **System > Certificates**, on the **CAs** tab.

See also:

Renew or Reissue a CA or Certificate

13.2.1 Certificate Authority Settings

When creating or editing a CA entry, the following options are available:

Trust Store
Controls whether this CA is added to the certificate trust store on the firewall. When added to the trust store, a CA will be considered valid for all certificate operations performed by the operating system. If the firewall must contact a server using a certificate issued by a private CA, this allows such certificates to be trusted by client programs such as LDAP authentication, SMTP notifications, URL table connections, and many others.

Randomize Serial
Controls whether the CA will randomize serial numbers when it signs certificates or if it will use a sequential serial number.

The current best practice is to randomize serial numbers so they are unpredictable. This also reduces the chances of generating two certificates with the same serial number in circumstances where the CA is moved between different hosts or signs certificates in multiple places.

Common Properties
See *Certificate Properties* which covers the remaining fields on the page.

When importing or editing an existing CA entry, the following options are available:

Certificate Data
The PEM-encoded certificate data for the CA.

Certificate data is typically contained in a file ending with `.crt` or `.pem`. It would be plain text, and enclosed in a block such as:

```
-----BEGIN CERTIFICATE-----
[A bunch of random-looking base64-encoded data]
-----END CERTIFICATE-----
```

The format varies slightly for ECDSA certificates.

Certificate Private Key
The PEM-encoded private key for the CA. If this is omitted, the CA cannot sign certificates or CRLs, but it can be used for other purposes. When empty, the CA is marked as "External". They key can be filled in later to enable signing and to have the CA treated as "Internal".

The key data is typically in a file ending in `.key`. It would be plain text data enclosed in a block such as:

```
-----BEGIN RSA PRIVATE KEY-----
[A bunch of random-looking base64-encoded data]
-----END RSA PRIVATE KEY-----
```

The format varies slightly for ECDSA keys.

Next Certificate Serial
The serial number of the next certificate, used when the CA is not set to randomize serial numbers.

It is essential that each certificate have a unique serial, or there will be problems later with certificate revocation. If the next serial is unknown, attempt to estimate how many certificates have been made from the CA, and then set the number high enough a collision would be unlikely.

13.2.2 Create a new Certificate Authority Entry

To create a new CA entry, start the process as follows:

- Navigate to **System > Certificates**, **CAs** tab
- Click **Add** to create a new a CA
- Enter a **Descriptive name** for the CA

  This is used as a label for this CA throughout the GUI.
- Select the **Method** that best suits how the CA will be generated

  Create an Internal Certificate Authority
  Creates a new root CA. Fill in the settings as described in *Certificate Authority Settings*.

  Import an Existing Certificate Authority
  Exports a CA certificate created on another host, with or without a private key. This can be useful in two ways: One, for CAs made using another system, and two, for CAs made by others that must be trusted.

  Fill in the settings as described in *Certificate Authority Settings*.

  Note: If the CA has been signed by an intermediary and not directly by a root CA, then import each entry in the chain separately, starting with the root CA.

  Create an Intermediate Certificate Authority
  Creates a new intermediate CA, to be signed by another internal CA on this firewall.

Pick an existing internal CA for the **Signing Certificate Authority** and fill in the remaining settings as described in *Certificate Authority Settings*.

If errors are reported, such as invalid characters or other input problems, they will be described on the screen. Correct the errors, and attempt to Save again.

13.2.3 Edit a Certificate Authority

To edit an existing CA:

- Navigate to **System > Certificates**, **CAs** tab
- Locate the CA entry in the list
- Click the  icon at the end of its row

The edit screen presented by the GUI allows editing the fields as if the CA were being imported.

For information on the fields on this screen, see *Certificate Authority Settings*. In most cases the purpose of this screen would be to add the CA to the trust store, correct the **Serial** of the CA if needed, or to add a key to an imported CA so it can be used to create and sign certificates and CRLs.

13.2.4 Export a Certificate Authority

To export a CA:

- Navigate to **System > Certificates**, **CAs** tab
- Locate the CA entry in the list
- Click the  icon at the end of its row to export the CA certificate.

  The file will download with the descriptive name of the CA as the file name, with the extension `.crt`.
- Click the  icon to export the private key for the CA if necessary

  The file will download with the descriptive name of the CA as the file name, with the extension `.key`.

  In most cases the private key for a CA would not be exported unless the CA is being moved to a new location or a backup is being made. When using the CA for a VPN or most other purposes, only export the certificate for the CA and do not export the key.

  Warning: If the private key for a CA gets into the wrong hands, the other party could generate new certificates that would be considered valid against the CA.

13.2.5 Remove a Certificate Authority

To remove a CA, first it must be removed from active use.

- Check areas that can use a CA, such as OpenVPN, IPsec, and packages.

Note: In most cases, the areas using a CA are noted in the **In Use** column of the CA list. This does not necessarily include all areas, especially if the CA is used by a package.

- Remove entries utilizing the CA or select a different CA
- Navigate to **System > Certificates**, **CAs** tab
- Locate the CA entry in the list
- Click  at the end of the row for the CA

Note: This icon will only be present if the CA is not in use.

- Click OK on the confirmation dialog

13.2.6 Renew a Certificate Authority

To renew a CA entry:

- Navigate to **System > Certificates**, **CAs** tab
- Locate the CA entry in the list
- Click  at the end of the row for the CA
- Follow the rest of the renewal procedure as described in *Renew or Reissue a CA or Certificate*

13.3 Certificate Management

Certificates are managed from **System > Certificates**, on the **Certificates** tab.

When creating a certificate on any platform the process generally follows this flow:

- User creates a certificate signing request (CSR) and set of keys. The public key is a part of the CSR, but the private key is separate.
- The user transmits only the CSR to the CA, not the private key which remains private to the user.
- The CA signs the CSR, which results in a certificate.
- The CA transmits the certificate to the user.

The user now has a certificate trusted by the CA, and the private key for the certificate.

The GUI handles most of this process automatically, but it also allows performing individual steps separately. For example, when creating an internal certificate, there is no need to create and sign a CSR in separate steps, the GUI automates that process. Aside from that, the GUI can create a CSR which can be sent to a separate CA, and it can also sign CSRs.

13.3.1 Certificate Settings

When creating a certificate entry or working with a CSR, the following common options are available:

Common Properties
See *Certificate Properties* which covers properties of most certificate entries.

Certificate Type
Sets the intended purpose of this certificate. This influences which key usage properties are set in the certificate and thus limits the ways in which the certificate can operate.

> **Warning:** The certificate can only be used for purposes which match the selected type. Attempting to use it in other ways will produce errors and fail, or prevent the certificate from being shown for selection.

User Certificate
Certificates for end users and clients. For example, IPsec and OpenVPN client certificates.

> **Note:** User type certificates include Extended Key Usage attributes indicating they may be used for client authentication. They also are marked with a constraint indicating that they are not a CA.

Server Certificate
Certificates for servers, services, daemons, etc. For example, HTTPS servers (GUI, Captive Portal, HAProxy, etc.), IPsec IKEv2 mobile server, OpenVPN servers, and for packages such as FreeRADIUS.

> **Note:** Server type certificates include Extended Key Usage attributes indicating they may be used for server authentication as well as the OID `1.3.6.1.5.5.8.2.2` which is used by Microsoft to signify that a certificate may be used as an IKE intermediate. These are required for Windows 7 and later to trust the server certificate for use with certain types of VPNs. They also are marked with constraints indicating that they are not a CA, and they have `nsCertType` set to `server`.

Alternative Names
Identifiers for this certificate, such as a hostname. See *Subject Alternative Name* for details.

When importing an existing certificate entry, the following options are available:

Certificate Data
The PEM-encoded certificate data for the certificate.

Certificate data is typically contained in a file ending with `.crt` or `.pem`. It would be plain text, and enclosed in a block such as:

```
-----BEGIN CERTIFICATE-----
[A bunch of random-looking base64-encoded data]
-----END CERTIFICATE-----
```

The format varies slightly for ECDSA certificates.

Private Key Data
The PEM-encoded private key for the certificate.

The key data is typically in a file ending in `.key`. It would be plain text data enclosed in a block such as:

```
-----BEGIN RSA PRIVATE KEY-----
[A bunch of random-looking base64-encoded data]
-----END RSA PRIVATE KEY-----
```

The format varies slightly for ECDSA keys.

13.3.2 Create a new Certificate

To create a new certificate, start the process as follows:

- Navigate to **System > Certificates**, **Certificates** tab
- Click **Add** to create a new certificate
- Enter a **Descriptive name** for the certificate

  This is used as a label for this certificate throughout the GUI.
- Select the **Method** that best suits how the certificate will be generated

  These options and further instructions are in the corresponding sections below:

  - Create an Internal Certificate
  - Import an Existing Certificate
  - Create a Certificate Signing Request
  - Sign a Certificate Signing Request
- Complete the steps for the chosen method
- Click Save to finish the import process

Create an Internal Certificate

The most common **Method** is *Create an Internal Certificate*. This will make a new certificate using one of the existing certificate authorities.

- Select the **Certificate Authority** which will sign this certificate. Only a CA that has a private key present can be in this list, as the private key is required in order for the CA to sign a certificate.
- Set the properties of the certificate as described in *Certificate Settings*.
- Click **Save**.

Import an Existing Certificate

To import an existing certificate from an external source, set **Method** to *Import an Existing Certificate*. This can be useful for certificates made using another system or for certificates provided by a third party.

There are two ways to import a certificate, indicated by the **Certificate Type** option:

X.509 (PEM)
: Enter the **Certificate data** and **Private key data**, which are both required. See *Certificate Settings* for details on populating the contents of the fields.

The most common error is not pasting in the right portion of the certificate or private key. Make sure to include the entire block, including the beginning header and ending footer around the encoded data.

PKCS #12 (PFX)
This method reads the certificate data from a PKCS #12 file, commonly found with a `.p12` extension. If the `.p12` file contains a CA, it is also imported along with the certificate, provided it does not already exist locally.

PKCS #12 Certificate
Click **Browse** to locate the `.p12` file on the local client, it will be uploaded and read when saving.

PKCS #12 Certificate Password
Enter the password used to protect the contents of the `.p12` file

Intermediates
When set, if the PKCS #12 file contains multiple CA entries in a chain, this option will import all of them instead of only one.

Create a Certificate Signing Request

Choosing a **Method** of *Certificate Signing Request* creates a new request file that can signed by a CA at a later time, including by a third party CA not present on the firewall. This is commonly used to obtain a certificate from a trusted root certificate authority.

The parameters for creating this certificate are identical to those for creating a certificate and are covered in *Certificate Settings*.

Note: Though the GUI shows fields for **Certificate Type** and **Alternative Names** as described in *Certificate Settings*, they are only suggestions for the CA. The signing CA may ignore these options and replace them with values of its own.

Sign a Certificate Signing Request

Signing a certificate signing request (CSR) is a special process which uses an internal CA on the firewall to sign a CSR and turn it into a full-fledged certificate.

The following options are available when signing a CSR:

CA to sign with
The CA on the firewall which will sign this CSR. This must be an internal CA (private key present).

CSR to sign
This option chooses whether to sign a new CSR not present on the firewall or an existing CSR on the firewall.

New CSR
When chosen, the GUI presents fields in which the CSR data can be pasted.

CSR Data
The PEM-encoded CSR data. CSR data is typically contained in a file ending with `.req` or `.pem`. It would be plain text, and enclosed in a block such as:

```
-----BEGIN CERTIFICATE REQUEST-----
[A bunch of random-looking base64-encoded data]
-----END CERTIFICATE REQUEST-----
```

Key Data
The optional PEM-encoded private key for the certificate. This is not required to sign a CSR, but may be useful, or even necessary, if the resulting certificate will be used on the firewall. For example, a private key would be required for a local service or as a user certificate used with a VPN export package.

The key data is typically in a file ending in `.key`. It would be plain text data enclosed in a block such as:

```
-----BEGIN RSA PRIVATE KEY-----
[A bunch of random-looking base64-encoded data]
-----END RSA PRIVATE KEY-----
```

Existing CSR
The remaining items in the drop-down list are CSR entries which already exist on the firewall. Choose one to sign.

Certificate Lifetime
The lifetime of the new certificate. See *Lifetime* for details.

Digest Algorithm
The digest algorithm for the new certificate. See *Digest Algorithm* for details.

When signing a CSR, the signing CA may also give new values for **Certificate Type** and **Alternative Names** as described in *Certificate Settings*. The signing process in the GUI cannot automatically read those values from a CSR, so set them again here.

When complete, the result is a certificate entry in the list, which can then be used or exported.

13.3.3 Edit a Certificate

To edit an existing certificate:

- Navigate to **System > Certificates**, **Certificates** tab
- Locate the Certificate entry in the list
- Click the  icon at the end of its row to reach the Edit page for the certificate.

The Edit page can modify some aspects of the certificate, such as:

- The **Descriptive Name** of the certificate.
- The **Certificate Data**, which may need to be replaced if the certificate was renewed by a CA off the firewall.
- The **Private key data**, which may need updated if the private key is regenerated (e.g. with a stronger key, or a different key type)

The Edit page also contains options for exporting entries with a password. See *Export Password-Protected Files or Use Different Encryption Options* for details.

13.3.4 Export a Certificate

There are multiple methods to export certificates. The primary differences are whether the files will have password protection and which type of encryption is used to protect the PKCS #12 archive. The certificate itself does not contain private information and thus does not require protection. The private key and PKCS #12 format files do contain private information and thus can be exported in a protected manner.

Export Unprotected Files

- Navigate to **System > Certificates**, **Certificates** tab
- Locate the Certificate entry in the list
- Click the  icon at the end of its row to export the certificate.

  The file will download with the descriptive name of the certificate as the file name, with the extension `.crt`.
- Click the  icon to export the private key for the certificate.

  The file will download with the descriptive name of the certificate as the file name, with the extension `.key`.
- Click the  icon to export a PCKS #12 file containing the CA, certificate, and private key together.

  The file will download with the descriptive name of the certificate as the file name, with the extension `.p12`.

Export Password-Protected Files or Use Different Encryption Options

The GUI can also export password-protected versions of the private key and PKCS #12 archives. This is more secure, but some systems may be compatible with password-protected keys. There is also an option to control the type of encryption used to protect the PKCS #12 archive because some algorithms are not compatible with certain platforms.

- Navigate to **System > Certificates**, **Certificates** tab
- Locate the Certificate entry in the list
- Click the  icon at the end of its row to reach the Edit page for the certificate.
- Fill in the desired **Export Password** (or leave it blank to export without a password)
- Choose an appropriate **PKCS #12 Encryption** option:

  High
  Uses AES-256 and SHA256 to encrypt the archive (default). This is the current strongest option, and it is compatible with pfSense software, FreeBSD, Linux, and Windows 10/11. In most cases this is the most desirable option.

  Note: Use this level when exporting for platforms with OpenSSL 3.0.

  Low
  Uses 3DES and SHA1 to encrypt the archive. This algorithm is considered weak and deprecated by most modern operating systems, but it is required by the key management built into macOS (current versions, including Ventura 13.2) and older versions of Windows.

Legacy
Uses RC2-40 and SHA1 to encrypt the archive. Avoid using this level if at all possible as it is extremely weak by modern standards. This was the previous default on older versions of pfSense software (Plus 22.05, CE 2.6.0).

> **Warning:** OpenSSL 3.0 will not read PKCS #12 archives encrypted with this method as it has deprecated this type of weak encryption.

- Click the  **Export Private Key** button to export the private key for the certificate.

  The password-protected file will download with the descriptive name of the certificate as the file name, with the extension `.key`.

- Click the  **PKCS #12** button to export a PCKS #12 file containing the CA, certificate, and private key together.

  The password-protected file will download with the descriptive name of the certificate as the file name, with the extension `.p12`.

13.3.5 Export a Certificate Signing Request

- Navigate to **System > Certificates**, **Certificates** tab
- Locate the CSR entry in the list
- Click the  icon at the end of its row to export the CSR.

  The file will download with the descriptive name of the CSR as the file name, with the extension `.req`.

13.3.6 Remove a Certificate

To remove a certificate, first it must be removed from active use.

- Check areas that can use a certificate, such as the GUI options, OpenVPN, IPsec, and packages

> **Note:** In most cases, the areas using a certificate are noted in the **In Use** column of the certificate list. This does not necessarily include all areas, especially if the certificate is used by a package.

- Remove entries using the certificate, or choose another certificate
- Navigate to **System > Certificates**, **Certificates** tab
- Locate the certificate to delete in the list
- Click  at the end of the row for the certificate

> **Note:** This icon will only be present if the certificate is not in use.

- Click OK on the confirmation dialog

13.3.7 Renew a Certificate

To renew a certificate entry:

- Navigate to **System > Certificates**, **Certificates** tab
- Locate the certificate entry in the list
- Click  at the end of the row for the certificate
- Follow the rest of the renewal procedure as described in *Renew or Reissue a CA or Certificate*

13.3.8 User Certificates

If a VPN is being used that requires user certificates, they may be created in one of several ways. The exact method depends on where the authentication for the VPN is being performed and whether the certificate already exists.

No Authentication or External Authentication

If there is no user authentication, or if the user authentication is being performed on an external server (RADIUS, LDAP, etc.) then make a user certificate like any other certificate described earlier. Ensure that *User Certificate* is selected for the **Certificate Type** and set the **Common Name** to match the username.

Local Authentication

If user authentication is being performed by this firewall, the user certificate can be made inside the User Manager. The User Manager can create a certificate while creating a user, or it can add certificates to existing users. These processes are documented at *Manage Local Users*.

13.4 Renew or Reissue a CA or Certificate

When a CA or certificate expires it must be replaced, renewed, or reissued. The GUI can Renew or Reissue a certificate using a semi-automatic process. This process can retain the existing properties of the CA or certificate, but results in a freshly signed copy. This process can also make changes to the lifetime, keys, and digest so they meet current security best practices.

The new copy of this certificate must be distributed to the intended target as it was originally.

13.4.1 Certificate Properties

The **Renew or Reissue** page displays information about the entry, including:

Subject
: The subject of the certificate, containing its Distinguished Name (DN)

Serial
: The serial number of the certificate.

Subject Key ID
: Fingerprint of the certificate key.

Certificate Type
Either **User** or **Server**, if known.

Issued By
The CA which signed the certificate (Name and DN)

13.4.2 Renew or Reissue Options

There are two options available which control what happens when the certificate is renewed:

Reuse Key
When set (default), the existing key on the certificate is retained. When unset, a fresh key will be created when the certificate is reissued.

Reuse Serial
Set this option to retain the existing serial number when reissuing. Uncheck to generate a new serial.

Retaining the serial when renewing a CA allows existing certificates to remain valid, though some clients may not respect the new CA if the serial does not change.

Similarly, certificates should have a new serial every time they are renewed, or some peers will reject them.

The exact behavior depends on the service and clients, but generally speaking it is safe to reuse the serial on a CA but not safe to reuse the serial on a server or user certificate. For example, OpenVPN is OK with reusing the serial number on a CA when renewing, while web browsers will reject changing a server certificate, even self-signed, if the serial does not change when the contents of the certificate change.

Strict Security
When set, upgrades the security of the certificate to meet current standards.

The **Renew or Reissue** page performs a security analysis on the certificate, comparing its current values for Lifetime, Digest, and RSA Key size with current best security practices. This analysis is printed at the bottom of the page. If any of the values are weak, the **Would Change** column in the analysis indicates **Yes**.

13.4.3 Renew or Reissue Example

To start the renewal process, first locate the CA or certificate to renew:

- Navigate to **System > Certificates**
- Navigate to the **CAs** tab for CA entries, or the **Certificates** tab for certificates
- Locate the entry to renew in the list
- Click  at the end of the row for the certificate to load the Renew or Reissue page for the certificate

---

Note: The  icon only appears for entries which have been signed by an internal CA on the firewall.

- Review the contents of the page
- Set the *Renew or Reissue Options* as desired

- Click  **Renew/Reissue**
- Click **OK** to confirm the action

When the process completes, the certificate entry is updated in the configuration.

Note: If the certificate is in use by a service on the firewall, the associated service(s) are restarted automatically.

For user certificates, the updated certificate must be exported and transmitted to the user. If a new key was generated by the renewal process, it must also be transmitted to the user.

13.5 Certificate Revocation List Management

Certificate Revocation Lists (CRLs) are a part of the X.509 system that publish lists of certificates that must no longer be trusted. These certificates may have been compromised or otherwise need to be invalidated. An application using a CA, such as OpenVPN may optionally use a CRL so it can verify connecting client certificates. A CRL is generated and signed against a CA using its private key, so in order to create or add certificates to a CRL in the GUI, the private key of the CA must be present. If the CA is managed externally and the private key for the CA is not on the firewall, a CRL may still be generated outside the firewall and imported.

The traditional way to use a CRL is to only have one CRL per CA and only add invalid certificates to that CRL. The GUI, however, allows creating multiple CRLs for a single CA. In OpenVPN, different CRLs may be chosen for separate VPN instances. This could be used, for example, to prevent a specific certificate from connecting to one instance while allowing it to connect to another. For IPsec, all CRLs are consulted and there is no selection as currently exists with OpenVPN.

Certificate Revocation Lists are managed from **System > Certificates**, on the **Certificate Revocation** tab.

From this screen CRL entries can be added, edited, exported, or deleted. The list shows all existing CRLs and an option to add a new CRL from a given CA. The screen also indicates whether the CRL is internal or external (imported), and it shows a count of how many certificates have been revoked on each CRL, and indicates if the CRL is in use.

13.5.1 Create a new Certificate Revocation List

To create a new CRL:

- Navigate to **System > Certificates**, **Certificate Revocation** tab
- Select a CA from the drop-down menu under the **Create or Import a New Certificate Revocation List**
- Click  **Add** at the end of the row to create a new CRL
- Set the **Method** to *Create an Internal Certificate Revocation List*
- Enter a **Descriptive Name** for the CRL

  This is used to identify this CRL in lists around the GUI. It's usually best to include a reference to the name of the CA and/or the purpose of the CRL.
- Enter the **Lifetime** value as a number of days for which the CRL should be valid

  The default value is `730` days (2 years).

Note: In practice, this limit would almost never be reached as the CRL is regenerated any time the CRL is edited or when a service which utilizes a CRL is reconfigured.

Note: The system attempts to prevent using too large a value for the lifetime to ensure the date doesn't overflow. On 32-bit platforms, the limit is before the UNIX time rollover in 2038. On other platforms, the limit is before UTC time 2-digit dates roll over in 2050. See Redmine #13424 for details. Systems reporting an expired CRL can work around the error by making a new CRL with a lower lifetime or by applying a patch on that Redmine issue.

- Click **Save**

The browser will be return to the CRL list, and the new entry will be shown there.

13.5.2 Import an Existing Certificate Revocation List

To import a CRL from an external source:

- Navigate to **System > Certificates**, **Certificate Revocation** tab
- Select a CA from the drop-down menu under the **Create or Import a New Certificate Revocation List**
- Click  **Add** at the end of the row to create a new CRL
- Set the **Method** to *Import an Existing Certificate Revocation List*
- Enter a **Descriptive Name** for the CRL

  This is used to identify this CRL in lists around the GUI. It's usually best to include a reference to the name of the CA and/or the purpose of the CRL.
- Enter the **CRL data**

  This is typically in a file ending in `.crl`. It would be plain text data enclosed in a block such as:

```
-----BEGIN X509 CRL-----
[A bunch of random-looking base64-encoded data]
-----END X509 CRL-----
```

- Click **Save** to finish the import process.

If an error appears, follow the on-screen instructions to correct the problem and then try again. The most common error is not pasting in the right portion of the CRL data. Make sure to enter the entire block, including the beginning header and ending footer around the encoded data.

Warning: New entries cannot be added to imported CRLs. To update an imported CRL, see *Updating an Imported Certificate Revocation List*.

13.5.3 Export a Certificate Revocation List

- Navigate to **System > Certificates**, **Certificate Revocation** tab
- Locate the CRL to delete in the list
- Click the  icon

The file will download with the descriptive name of the CRL as the file name, and the extension `.crl`.

13.5.4 Delete a Certificate Revocation List

- Check areas that can use a CRL, such as IPsec and OpenVPN

Note: In most cases, the areas using a CRL are noted in the **In Use** column of the CRL list. This does not necessarily include all areas, especially if the CRL is used by a package.

- Remove entries using the CRL, or choose another CRL instead
- Navigate to **System > Certificates**, **Certificate Revocation** tab
- Locate the CRL to delete in the list
- Click the  icon at the end of the row for the CRL

Note: This icon will only be present if the CRL is not in use.

- Click OK on the confirmation dialog

If an error appears, follow the on-screen instructions to correct the problem and then try again.

13.5.5 Revoke a Certificate

A CRL isn't useful unless it contains revoked certificates. A certificate is revoked by adding the certificate to a CRL, or by entering its serial number.

- Navigate to **System > Certificates**, **Certificate Revocation** tab
- Locate the CRL to edit in the list
- Click the  icon at the end of the row for the CRL

  The GUI lists any revoked certificates on the CRL, and a control to add new ones.
- Select a **Reason** from the drop-down list to indicate why the certificate is being revoked

  This information doesn't affect the validity of the certificate it is merely informational in nature. This option may be left at the default value.
- To revoke by certificate, select the certificate(s) from the **Revoke Certificates** list

---

Note: Multiple certificates can be revoked at once by selecting all of them in the list.

- To revoke by serial number, enter one or more certificate serial numbers separated by spaces in the **Revoke by Serial** field
- Click  **Add** and the certificate(s) will be added to the CRL

Note: Certificates can be revoked by selection and by serial at the same time.

After adding a certificate, the CRL will be re-written if it is currently in use by any VPN instances so that the CRL changes will be immediately active.

13.5.6 Removing a Certificate from a CRL

Certificates can be removed from the CRL when editing a CRL:

- Navigate to **System > Certificates**, **Certificate Revocation** tab
- Locate the CRL to edit in the list
- Click the  icon at the end of the row for the CRL
- Find the certificate in the list and click the  icon to remove it from the CRL
- Click **OK** on the confirmation dialog

After removing a certificate, the CRL will be re-written if it is currently in use by any VPN instances so that the CRL changes will be immediately active.

13.5.7 Updating an Imported Certificate Revocation List

To update an imported CRL:

- Navigate to **System > Certificates**, **Certificate Revocation** tab.
- Locate the CRL to edit in the list
- Click the  icon at the end of the row for the CRL
- Enter a new copy of the **CRL Data**
- Click **Save**

After updating the imported CRL, it will be re-written if it is currently in use by any VPN instances so that the CRL changes will be immediately active.

13.6 DH Parameters

DH parameters are extra bits of randomness that help out during the key exchange process. They do not have to match on both sides of the tunnel, and new DH parameters can be made at any time. DH parameters are not specific to a given setup in the way that certificates or keys are. There is no need to import an existing set of DH parameters because generating new parameters is a better practice.

pfSense® software ships with a default set of DH parameter files so that new firewalls do not have to spend significant CPU resources to build them when they are needed. These pre-generated parameters are stored in files named `/etc/dh-parameters.<length>`. Selecting a specific length in the GUI will use the DH parameter set from the corresponding file. These DH parameters are not stored in config.xml.

To generate a new set of DH parameters, which can take quite a long time depending on the hardware in use, run the following commands:

```
/usr/bin/openssl dhparam -out /etc/dh-parameters.2048 2048
/usr/bin/openssl dhparam -out /etc/dh-parameters.4096 4096
/usr/bin/openssl dhparam -out /etc/dh-parameters.8192 8192
```

CPU time used to generate the parameters increases significantly with length.

The GUI allows longer DH parameters to be selected if they exist in `/etc/` in the format specified above.

Available lengths are: `2048`, `3072`, `4096`, `7680`, `8192`, `15360`, and `16384`.

For example, to generate a new set of DH parameters of length `16384`, run:

```
/usr/bin/openssl dhparam -out /etc/dh-parameters.16384 16384
```

The **Certificate Manager** under **System > Certificates**, creates and maintains certificate authority (CA), certificate, and certificate revocation list (CRL) entries for use by the firewall.

Entries in the Certificate Manager are used by the firewall for purposes such as TLS for the GUI, VPNs, LDAP, various packages, and more.

13.7 Basic Introduction to X.509 Public Key Infrastructure

One authentication option for VPNs is to use X.509. An in depth discussion of X.509 and Public Key Infrastructure (PKI) is outside the scope of this documentation, and is the topic of a number of entire books for those interested in details. This chapter provides a basic understanding necessary for creating and managing certificates.

With PKI, a CA is the source of trust and is the first entity of a PKI structure. This CA then signs all the individual certificates in a set. The certificate of the CA is used on VPN servers and clients to verify the authenticity of server and client certificates. The certificate for the CA can be used to verify signing on certificates, but not to sign certificates. Signing certificates requires the private key for the CA. The secrecy of the CA private key is what ensures the security of a PKI. Anyone with access to the CA private key can generate certificates to be used on a PKI, hence it must be kept secure. This key is never distributed to clients or servers.

Warning: Never copy more files to clients than are needed, as this may compromise the security of the PKI structure.

A certificate is considered valid if it has been trusted by a given CA. In the case of a VPN, this means that a certificate made from a specific CA would be considered valid for any VPN using that CA. For that reason the best practice is to create a unique CA for each VPN that has a different level of security. For instance, if there are two mobile access

VPNs with the same security access, using the same CA for those VPNs is OK. However, if one VPN is for users and another VPN is for remote management, each with different restrictions, then it is best for each VPN to have a unique CA.

Certificate revocation lists (CRLs) are lists of certificates that have been compromised or otherwise invalidated. Revoking a certificate will cause it to be considered untrusted so long as the application using the CA also uses a CRL. CRLs are generated and signed against a CA using its private key, so in order to create or add certificates to a CRL in the GUI the private key for a CA must be present.

CHAPTER

FOURTEEN

FIREWALL

One of the primary functions performed by pfSense® software is filtering traffic, deciding which traffic to pass or block between networks. This section covers fundamentals of firewall behavior, best practices, and required information necessary to configure firewall rules.

14.1 Managing Firewall Rules

Firewall rules control traffic passing through the firewall. These topics describe how to create and manage rules, plus settings related to rules.

14.1.1 Firewall Fundamentals

This section deals primarily with introductory firewall concepts and lays the groundwork for understanding how to configure firewall rules using pfSense® software.

Basic Terminology

Rule and *ruleset* are two terms used throughout this chapter:

Rule
: Refers to a single entry on the **Firewall > Rules** screen. A rule instructs the firewall how to match or handle network traffic.

Ruleset
: Refers to a group of rules collectively. Either *all* firewall rules as a whole, or a set of rules in a specific context such as the rules on an interface tab. The complete firewall ruleset is the sum of all user configured and automatically added rules, which are covered further throughout this section.

Rulesets on the **Interface** tabs are evaluated on a **first match** basis. This means that reading the ruleset for an interface from top to bottom, the first rule that matches will be the one used by the firewall. Evaluation stops after reaching this match and then the firewall takes the action specified by that rule. Always keep this in mind when creating new rules, especially when crafting rules to restrict traffic. The most permissive rules should be toward the bottom of the list, so that restrictions or exceptions can be made above them.

Note: The **Ethernet** and **Floating** tabs are exceptions to this rule processing logic. See *Ethernet (Layer 2) Rules* and *Floating Rules* for details.

Stateful Filtering

pfSense software is a stateful firewall, which means it remembers information about connections flowing through the firewall so that it can automatically allow reply traffic. This data is retained in the **State Table**. The connection information in the state table includes the source, destination, protocol, ports, and more: Enough to uniquely identify a specific connection.

Using this mechanism, traffic need only be permitted on the interface where it **enters** the firewall. When a connection matches a pass rule the firewall creates an entry in the state table. Reply traffic to connections is automatically allowed back through the firewall by matching it against the state table rather than having to check it against rules in both directions. This includes any related traffic using a different protocol, such as ICMP control messages that may be provided in response to a TCP, UDP, or other connection.

See also:

See *Firewall & NAT* and *State Type* for more information about state options and types.

Note: This does not apply to *Ethernet (Layer 2) Rules* which do not keep state, or to rules which have manually disabled the keep state option.

State Policy

The behavior of state matching can be fine-tuned by changing the State Policy to either strictly bind states to interfaces and only match on those interfaces ("Interface Bound States") or it can be more relaxed and ignore the interface when matching packets to states ("Floating States").

This behavior can be changed globally (*Firewall & NAT Advanced Options,*) and on a per-rule basis (*Firewall Rule Configuration, Advanced Options, State Policy*).

State table size

The firewall state table has a maximum size to prevent memory exhaustion. Each state takes approximately 1 KB of RAM. The default state table size in pfSense is calculated by taking about 10% of the RAM available in the firewall by default. On a firewall with 1GB of RAM, the default state table size can hold approximately 100,000 entries.

See also:

See *Large State Tables* for more information on state table sizing and RAM usage.

Each user connection typically consists of two states: One created as it enters the firewall, and one as it leaves the firewall. Therefore, with a state table size of 1,000,000, the firewall can handle approximately 500,000 user sessions actively traversing the firewall before any additional connections will be dropped. This limit can be increased as needed so long as it does not exceed the available amount of RAM in the firewall.

To increase the state table size:

- Navigate to **System > Advanced** on the **Firewall & NAT** tab
- Enter the desired number for **Firewall Maximum States**, or leave the box blank for the default calculated value. See Figure *Increased State Table Size to 2,000,000*
- Click **Save**

Historical state table usage is tracked by the firewall. To view the graph:

- Navigate to **Status > Monitoring**

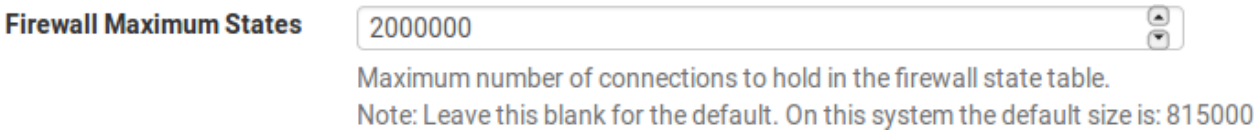


Fig. 1: Increased State Table Size to 2,000,000

- Click  to expand the graph options
- Set **Category** for the **Left Axis** to *System*
- Set the **Graph** for the **Left Axis** to *States*
- Click  **Update Graphs**

Block vs. Reject

There are two ways to disallow traffic using firewall rules on pfSense: **Block** and **reject**.

A rule set to **block** will silently drop traffic. A blocked client will not receive any response and thus will wait until its connection attempt times out. This is the behavior of the default deny rule in pfSense software.

A rule set to **reject** will respond back to the client for denied TCP and UDP traffic, letting the sender know that the connection was refused. Rejected TCP traffic receives a TCP RST (reset) in response, and rejected UDP traffic receives an ICMP unreachable message in response. Though reject is a valid choice for any firewall rule, IP protocols other than TCP and UDP are not capable of being rejected; These rules will silently drop other IP protocols because there is no standard for rejecting other protocols.

Deciding Between Block and Reject

There has been much debate amongst security professionals over the years as to the value of block vs. reject. Some argue that using block makes more sense, claiming it "slows down" attackers scanning the Internet. When a rule is set to reject, a response is sent back immediately that the port is closed, while block silently drops the traffic, causing the attacker's port scanner to wait for a response. That argument does not hold water because every good port scanner can scan hundreds or thousands of hosts simultaneously, and the scanner is not stalled waiting for a response from closed ports. There is a minimal difference in resource consumption and scanning speed, but so slight that it shouldn't be a consideration.

If the firewall blocks all traffic from the Internet, there is a notable difference between block and reject: Nobody knows the firewall is online. If even a single port is open, the value of that ability is minimal because the attacker can easily determine that the host is online and will also know what ports are open whether the blocked connections have been rejected by the firewall. While there isn't significant value in block over reject, the best practice is to use block on WAN rules. There is some value in not actively handing information to potential attackers, and it is also a bad practice to automatically respond to an external request unnecessarily.

For rules on internal interfaces the best practice is to use reject in most situations. When a host tries to access a resource that is not permitted by firewall rules, the application accessing it may hang until the connection times out or the client program stops trying to access the service. With reject the connection is immediately refused, and the client avoids these hangs. This is usually nothing more than an annoyance, but it is still a good idea to use reject to avoid potential application problems induced by silently dropping traffic inside a network.

14.1.2 Introduction to the Firewall Rules screen

This section provides an introduction and overview of the Firewall Rules screen located at **Firewall > Rules**. This page lists the WAN ruleset to start with, which by default has no entries other than those for **Block private networks** and **Block bogon networks** if those options are active on the WAN interface, as shown in Figure *Default WAN Rules*.

Tip: Click the to the right of the **Block private networks** or **Block bogon networks** rules to reach the WAN interface configuration page where these options can be enabled or disabled. See *Block Private Networks* and *Block Bogon Networks* for more details.

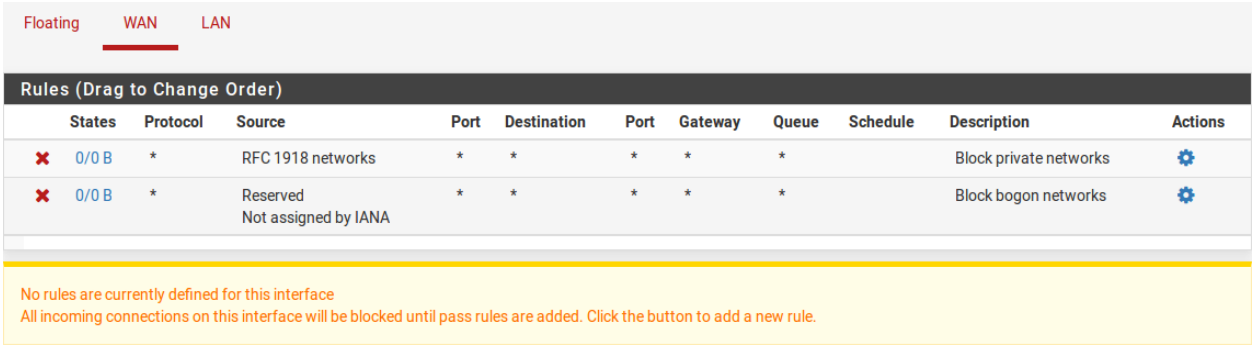


Fig. 2: Default WAN Rules

Click the **LAN** tab to view the LAN rules. By default, the only entries are the *Default allow LAN to any* rules for IPv4 and IPv6 as seen in Figure *Default LAN Rules*, and the **Anti-Lockout Rule** if it is active. The anti-lockout rule is designed to prevent administrators from accidentally locking themselves out of firewall management services. Click



next to the anti-lockout rule to reach the page where this rule can be disabled.

See also:

For more information on how the Anti-Lockout Rule works and how to disable the rule, see *Anti-lockout Rule* and *Anti-lockout*.

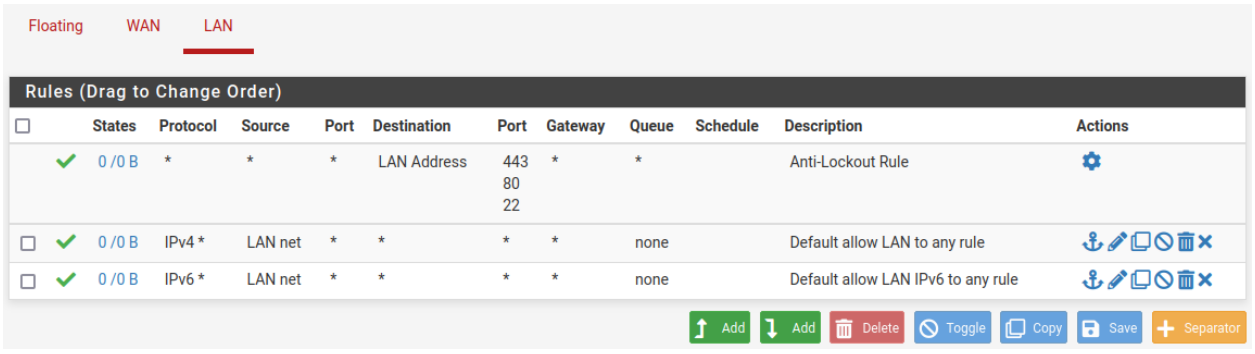


Fig. 3: Default LAN Rules

To display rules for other interfaces, click their respective tabs. OPT interfaces will appear with their descriptive names, so if the OPT1 interface was renamed `DMZ`, then the tab for its rules will also say **DMZ**.

To the left of each rule is a set of indicator icons, including:

- The action of the rule: pass (), block (), or reject ().
- Logging status: If logging is enabled for the rule,  is present.
- Advanced options: If the rule has any advanced options enabled, an  icon is present.

Hovering the mouse cursor over any of these icons will display text explaining their meaning. The same icons are shown for disabled rules, except the icon and the rule are a lighter shade of their original color.

Adding a firewall rule

To add a rule to the top of the list, click  **Add**.

To add a rule to the bottom of the list, click  **Add**.

Editing Firewall Rules

To edit a firewall rule, click  to the right of the rule, or double click anywhere on the line.

The edit page for that rule will load, and from there adjustments are possible. See *Configuring Firewall Rules* for more information on the options available when editing a rule.

Reordering Firewall Rules

The order of the rules on an interface can be changed in two different ways: Drag-and-drop or select-and-click.

To reorder rules using the drag-and-drop method:

- Move the mouse over the firewall rule to move, the cursor will change to indicate movement is possible.
- Click and hold the mouse button down
- Drag the mouse to the desired location for the rule
- Release the mouse button
- Click  **Save** to store the new rule order

Warning: Attempting to navigate away from the page after moving a rule, but before saving the order, will result in the browser presenting an error confirming whether to exit the page. If the browser navigates away from the page without saving, the rule will still be in its original location.

To move rules in the list in groups or by selecting them first, use the select-and-click method:

- Select the rules to move

Note: Select rules by single clicking anywhere on their line or by checking the box at the start of the row.

- Click  on the row **below** where the rule should be moved.

Tip: Hold **Shift** before clicking the mouse on  to move the rule below the selected rule instead of above.

When moving rules using the select-and-click method, the new order is stored automatically.

Copying Firewall Rules

To make a new rule that is similar to an existing rule, click  to the right of the existing rule. The edit screen will appear with the existing rule's settings pre-filled, ready to be adjusted. When duplicating an existing rule, the new rule will be added directly *below* the original rule. For more information about how to configure the new rule, see *Configuring Firewall Rules*.

To copy multiple rules:

- Select the rules to copy

Note: Select rules by single clicking anywhere on their line or by checking the box at the start of the row.

- Click the  **Copy** button below the rule list

  The firewall will open a new modal dialog with options to set before copying.
- Select the **Destination Interface**
- Select **Convert interface definitions** to automatically adjust the source of the rule to match the target interface, if necessary
- Click  **Paste** to complete the operation

Warning: When copying rules to different interfaces, they may fall at the start or the end of the target interface rule list depending on the order of the interface rules in the configuration. Be prepared to reorder the rules on the target interface before applying changes.

Deleting Firewall Rules

To delete a single rule, click  to the right of the rule. The firewall will present a confirmation prompt before deleting the rule.

To delete multiple rules:

- Select the rows to remove

> **Note:** Select rules by single clicking anywhere on their line or by checking the box at the start of the row.

- Click the  **Delete** button below the rule list
- Confirm the action

Checking Rule Usage

The **States** column contains usage counters for each rule. It shows the number of active states created by a rule and the amount of traffic consumed by those states.

Hovering the mouse over these counters shows additional detailed statistics:

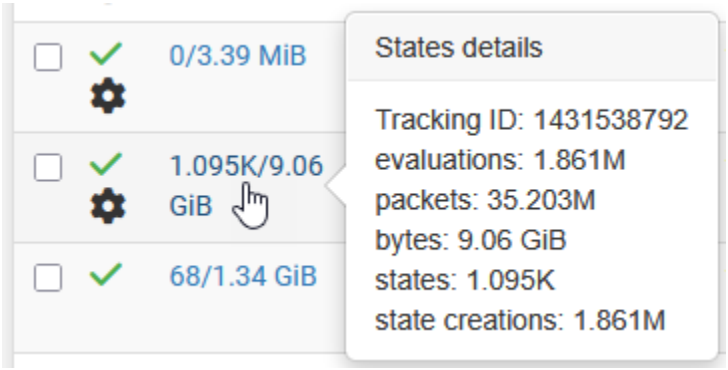


Fig. 4: Firewall Rule Usage Statistics

The statistics displayed by this view include:

Tracking ID
: The firewall *Rule Tracking ID* which uniquely identifies this rule.

Evaluations
: The number of times the firewall has evaluated this rule when processing packets.

Packets
: The number of packets passed by this rule.

Bytes
: The amount of traffic, in bytes, passed by this rule.

States
: The number of active state table entries created by this rule.

State Creations
The total number of state table entries created by this rule.

Note: Though the firewall makes an effort to maintain these statistics, the values can reset over time depending on firewall ruleset reloads and other similar actions.

Clicking the value in this column displays a list of states created by the rule.

See also:

Viewing Firewall States in the GUI

Clearing States Created by a Rule

Click the  icon to the right of a rule and then confirm the action to clear all active states created by that rule.

Note: This only affects states on this interface created by this rule directly. It does not clear states on other interfaces where traffic may have exited the firewall.

Disabling and Enabling Firewall Rules

To disable a rule, click  at the end of its row. The appearance of the rule will change to a lighter shade to indicate that it is disabled and the  icon changes to  .

To enable a rule which was previously disabled, click  at the end of its row. The appearance of the rule will return to normal and the enable/disable icon will return to the original  .

A rule may also be disabled or enabled by editing the rule and toggling the **Disabled** checkbox.

To disable or enable multiple rules at once:

- Select the rules to disable

  Note: Select rules by single clicking anywhere on their line or by checking the box at the start of the row.

- Click the  **Toggle** button below the rule list

Rule Separators

Firewall Rule Separators are colored bars in the ruleset that contain a small bit of text, but do not take any action on traffic. They are useful for visually separating or adding notes to special parts of the ruleset. Figure *Firewall Rule Separators Example* shows how they can group rules and document the ruleset.

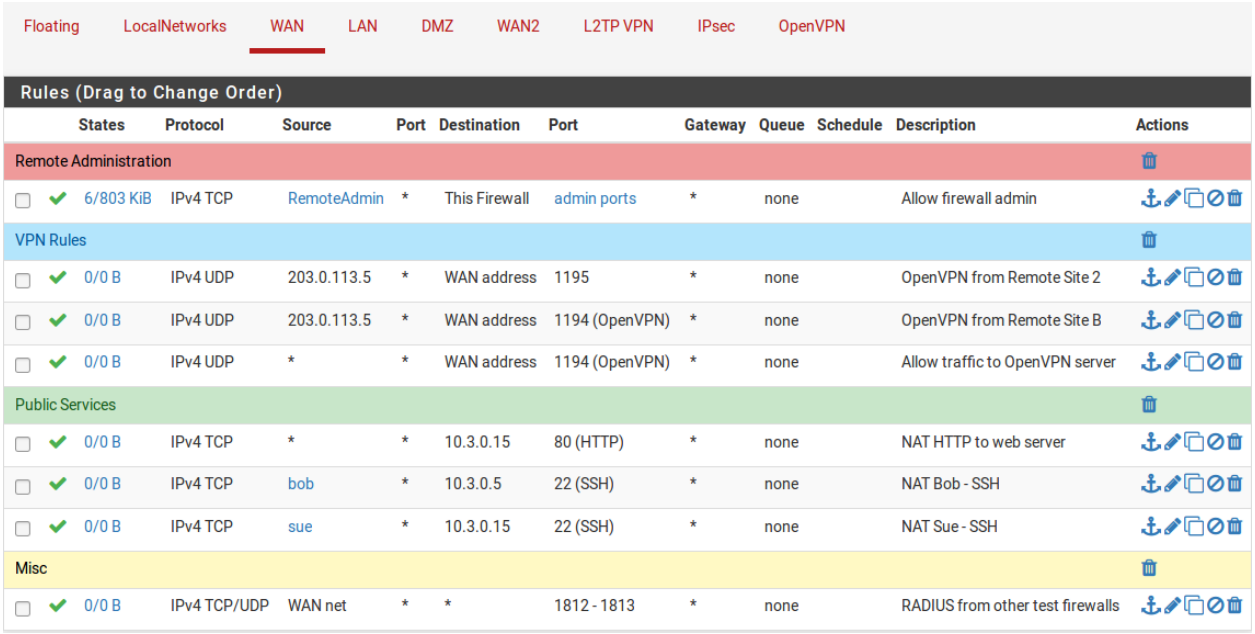


Fig. 5: Firewall Rule Separators Example

To create a new Rule Separator:

- Open the firewall rule tab where the Rule Separator will reside
- Click  **Separator**
- Enter description text for the Rule Separator
- Choose the color for the Rule Separator by clicking the  icon of the desired color
- Click and drag the Rule Separator to its new location
- Click  **Save** inside the Rule Separator to store its contents
- Click  **Save** at the bottom of the rule list

To move a Rule Separator:

- Open the firewall rule tab containing the Rule Separator
- Click and drag the Rule Separator to its new location
- Click  **Save** at the bottom of the rule list

To delete a Rule Separator:

- Open the firewall rule tab containing the Rule Separator
- Click  inside the Rule Separator on the right side
- Click  **Save** at the bottom of the rule list

Rule Separators cannot be edited. If a change in text or color is required, create a new Rule Separator and delete the existing entry.

Tracking Firewall Rule Changes

The firewall tracks rule creation and changes with data stored on each rule. These timestamps are visible when editing an existing rule.

See also:

Rule Information

14.1.3 Ingress Filtering

Ingress filtering refers to the concept of filtering traffic entering a network from an external source such as the Internet. In deployments with multi-WAN, the firewall has multiple ingress points. The default ingress policy on pfSense® software is to **block all traffic** as there are no allow rules on WAN in the default ruleset. Replies to traffic initiated from inside the local network are automatically allowed to return through the firewall by the state table.

14.1.4 Egress Filtering

Egress filtering refers to the concept of filtering traffic initiated inside the local network, destined for a remote network such as the Internet. pfSense, like nearly all similar commercial and open source solutions, comes with a LAN rule allowing everything from the LAN out to the Internet. This isn't the best way to operate, however. It has become the de facto default in most firewall solutions because it is what most people expect. The common misperception is "Anything on the internal network is 'trustworthy', so why bother filtering"?

Why employ egress filtering?

From our experience in working with countless firewalls from numerous vendors across many different organizations, most small companies and home networks do not employ egress filtering. It can increase the administrative burden as each new application or service may require opening additional ports or protocols in the firewall. In some environments it is difficult because the administrators do not completely know what is happening on the network, and they are hesitant to break things. In other environments it is impossible for reasons of workplace politics. The best practice is for administrators to configure the firewall to allow only the minimum required traffic to leave a network where possible. Tight egress filtering is important for several reasons:

Limit the Impact of a Compromised System

Egress filtering limits the impact of a compromised system. Malware commonly uses ports and protocols that are not required on most business networks. Some bots rely on IRC connections to phone home and receive instructions. Some will use more common ports such as TCP port 80 (normally HTTP) to evade egress filtering, but many do not. If access to TCP port 6667, the usual IRC port, is not permitted by the firewall, bots that rely on IRC to function may be crippled by the filtering.

Another example is a case where the inside interface of a pfSense software installation was seeing 50-60 Mbps of traffic while the WAN had less than 1 Mbps of throughput. There were no other interfaces on the firewall. Some investigation showed the cause as a compromised system on the LAN running a bot participating in a distributed denial of service (DDoS) attack against a Chinese gambling website. The attack used UDP port 80, and in this network UDP port 80 was not permitted by the egress ruleset so all the DDoS was accomplishing was stressing the inside interface of the firewall with traffic that was being dropped. In this situation, the firewall was happily chugging along with no performance degradation and the network's administrator did not know it was happening until it was discovered by accident.

The attack described in the above paragraph likely used UDP port 80 for two main reasons:

- UDP allows large packets to be sent by the client without completing a TCP handshake. With stateful firewalls being the norm, large TCP packets will not pass until the handshake is successfully completed, and this limits the effectiveness of the DDoS.
- Those who do employ egress filtering are commonly too permissive, allowing TCP and UDP where only TCP is required, as in the case of HTTP.

These types of attacks are commonly launched from compromised web servers. With a wide open egress ruleset, the traffic will go out to the Internet, and has the potential to overflow the state table on the firewall, cost money in bandwidth usage, and/or degrade performance for everything on the Internet connection.

Outbound SMTP is another example. Only allow SMTP (TCP port 25) to leave any network from a mail server. Or if a mail server is externally hosted, only allow internal systems to talk to that specific outside system on TCP port 25. This prevents every other system in the local network from being used as a spambot, since their SMTP traffic will be dropped. Many mail providers have moved to using only authentication submission from clients using TCP port 587, so clients should not need access to port 25. This has the obvious benefit of limiting spam, and also prevents the network from being added to numerous black lists across the Internet that will prevent that site from sending legitimate e-mail to many mail servers. This may also prevent the ISP for that site from shutting off its Internet connection due to abuse.

The ideal solution is to prevent these types of things from happening in the first place, but egress filtering provides another layer that can help limit the impact if other measures fail.

Prevent a Compromise

Egress filtering can prevent a compromise in some circumstances. Some exploits and worms require outbound access to succeed. An older but good example of this is the Code Red worm from 2001. The exploit caused affected systems to pull an executable file via TFTP (Trivial File Transfer Protocol) and then execute it. A web server almost certainly does not need to use the TFTP protocol, and blocking TFTP via egress filtering prevented infection with Code Red even on unpatched servers. This is largely only useful for stopping completely automated attacks and worms as a real human attacker will find any holes that exist in egress filtering and use them to their advantage. Again, the correct solution to prevent such a compromise is to fix the network vulnerabilities used as an attack vector, however egress filtering can help.

Limit Unauthorized Application Usage

Many applications such as VPN clients, peer-to-peer software, instant messengers, and more rely on atypical ports or protocols to function. While a growing number of peer-to-peer and instant messenger applications will port hop until finding a port which is allowed out of the local network, many will be prevented from functioning by a restrictive egress ruleset, and this is an effective means of limiting many types of VPN connectivity.

Prevent IP Address Spoofing

This is a commonly cited reason for employing egress filtering, but pfSense software automatically blocks spoofed traffic via the "anti-spoof" functionality of PF, so it isn't applicable here. Preventing IP address spoofing means that malicious clients cannot send traffic with obviously falsified source addresses.

Prevent Information Leaks

Certain protocols should never be allowed to leave a local network. Specific examples of such protocols vary from one environment to another, but a few common examples are:

- Microsoft RPC (Remote Procedure Call) on TCP port 135
- NetBIOS on TCP and UDP ports 137 through 139
- SMB/CIFS (Server Message Block/Common Internet File System) on TCP and UDP port 445.

Stopping these protocols can prevent information about the internal network from leaking onto the Internet, and will prevent local systems from initiating authentication attempts with Internet hosts. These protocols also fall under *Limit the Impact of a Compromised System* as discussed previously since many worms have relied upon these protocols to function. Other protocols that may be relevant are syslog, SNMP, and SNMP traps. Restricting this traffic will prevent misconfigured network devices from sending logging and other potentially sensitive information out to the Internet. Rather than worry about what protocols can leak information out of a local network and need to be blocked, the best practice is to only allow the traffic that is required.

Approaches for implementing egress filtering

On a network that has historically not employed egress filtering, it can be difficult to know what traffic is absolutely necessary. This section describes some approaches for identifying traffic and implementing egress filtering.

Allow what is known, block the rest, and work through the fallout

One approach is to add firewall rules for known required traffic to be permitted. Start with making a list of things known to be required such as in Table *Egress Traffic Required*.

Table 1: Egress Traffic Required

Description	Source	Destination	Destination port
HTTP and HTTPS from all hosts	LAN Network	Any	TCP 80 and 443
SMTP from mail server	Mail Server	Any	TCP 25
DNS queries from internal DNS servers	DNS Servers	Any	TCP and UDP 53

After making the list, configure firewall rules to pass only that traffic and let everything else hit the default deny rule.

Log Traffic and Analyze Logs

Another alternative is to enable logging on all pass rules and send the logs to a syslog server. The logs can be analyzed by the syslog server to see what traffic is leaving the network. pfSense software uses a custom log format, so the logs typically need be parsed by a custom script unless the server has some knowledge of the pfSense software filter log format. Analysis of the logs will help build the required ruleset with less fallout as it will yield a better idea of what traffic is necessary on the local network.

14.1.5 Firewall Rule Best Practices

This section covers general best practices for firewall rule configuration.

Default Deny

There are two basic philosophies in computer security related to access control: default allow and default deny. A default deny strategy for firewall rules is the best practice. Firewall administrators should configure rules to permit only the bare minimum required traffic for the needs of a network, and let the remaining traffic drop with the default deny rule built into pfSense® software. In following this methodology, the number of deny rules in a ruleset will be minimal. They still have a place for some uses, but will be minimized in most environments by following a default deny strategy.

In a default two-interface LAN and WAN configuration, pfSense software utilizes default deny on the WAN and default allow on the LAN. Everything inbound from the Internet is denied, and everything out to the Internet from the LAN is permitted. All home grade routers use this methodology, as do all similar open source projects and most similar commercial offerings. It's what most people expect out of the box, therefore it is the default configuration. That said, while it is a convenient way to start, it is not the recommended means of long-term operation.

pfSense software users often ask "What bad things should I block?" but that is the wrong question as it applies to a default allow methodology. Noted security professional Marcus Ranum includes default permit in his "Six Dumbest Ideas in Computer Security" paper, which is recommended reading for any security professional. Permit only what a network requires and avoid leaving the default allow all rule on the LAN and adding block rules for "bad things" above the permit rule.

Keep it short

The shorter a ruleset, the easier it is to manage. Long rulesets are difficult to work with, increase the chances of human error, tend to become overly permissive, and are significantly more difficult to audit. Utilize aliases to keep the ruleset as short as possible.

Review Firewall Rules

The best practice is a manual review of the firewall rules and NAT configuration on a periodic basis to ensure they still match the minimum requirements of the current network environment. The recommended frequency of such reviews varies from one environment to another. In networks that do not change frequently, with a few firewall administrators and good change control procedures, quarterly or semi-annually is usually adequate. For fast changing environments or those with poor change control and several people with firewall access, review the configuration at least on a monthly basis.

Quite often when reviewing rules with customers, Netgate TAC asks about specific rules, and they respond with "We removed that server six months ago." If something else had taken over the same internal IP address as the previous server, then traffic would have been allowed to the new server that may not have been intended.

Document The Configuration

In all but the smallest networks, it can be hard to recall what is configured where and why. The best practice is to use the **Description** field in firewall and NAT rules to document the purpose of the rules. In larger or more complex deployments, create and maintain a more detailed configuration document describing the entire pfSense software configuration. When reviewing the firewall configuration in the future, this will help determine which rules are necessary and why they are there. This also applies to any other area of the configuration.

It is also important to keep this document up to date. When performing periodic configuration reviews, also review this document to ensure it remains up-to-date with the current configuration. Ensure this document is updated whenever configuration changes are made.

Reducing Log Noise

By default, pfSense software logs packets blocked by the default deny rule. This means all the noise getting blocked from the Internet will be logged. Sometimes there will not be much noise in the logs, but in many environments there will inevitably be something incessantly spamming the logs.

On networks using large broadcast domains – a practice commonly employed by cable ISPs – this is most often NetBIOS broadcasts from clue-deficient individuals who connect Windows machines directly to their broadband connections. These machines will constantly pump out broadcast requests for network browsing, among other things. ISP routing protocol packets may also be visible, or router redundancy protocols such as VRRP or HSRP. In co-location environments such as data centers, a combination of all of those things may be present.

Because there is no value in knowing that the firewall blocked 14 million NetBIOS broadcasts in the past day, and that noise could be covering up logs that are important, it is a good idea to add a block rule on the WAN interface for repeated noise traffic. By adding a block rule *without logging enabled* on the WAN interface, this traffic will still be blocked, but no longer fill the logs.

The rule shown in Figure *Firewall Rule to Prevent Logging Broadcasts* is configured on a test system where the "WAN" is on an internal LAN behind an edge firewall. To get rid of the log noise to see the things of interest, we added this rule to block – but not log – anything with the destination of the broadcast address of that subnet.

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
☐	✖	0/0 B	IPv4 *	*	*	10.0.64.255	*	*	none		Do not log broadcasts

Fig. 6: Firewall Rule to Prevent Logging Broadcasts

The best practice is to add similar rules, matching the specifics of any log noise observed in an environment. Check the firewall logs under **Status > System Logs**, **Firewall** tab to see what kind of traffic the firewall is blocking, and review how often it appears in the log. If any particular traffic is consistently being logged more than 5 times a minute, and the traffic is not malicious or noteworthy, add a block rule for it to reduce log noise.

Logging Practices

Out of the box, pfSense software does not log any passed traffic and logs all dropped traffic. This is the typical default behavior of almost every open source and commercial firewall. It is the most practical, as logging all passed traffic is rarely desirable due to the load and log levels generated. This methodology is a bit backwards, however, from a security perspective. Blocked traffic cannot harm a network, so its log value is limited, while traffic that gets passed could be very important log information to have if a system is compromised. After eliminating any useless block noise as described in the previous section, the remainder is of some value for trend analysis purposes. If significantly more or less log volume than usual is observed, it is probably good to investigate the nature of the logged traffic. OSSEC, an open source host-based intrusion detection system (IDS), is one system that can gather logs from a firewall via syslog and alert based on log volume abnormalities.

14.1.6 Rule Methodology

In pfSense® software, rules on interface tabs are applied on a per-interface basis, always in the inbound direction on that interface. This means traffic initiated from hosts connected to the LAN is filtered using the LAN interface rules. Traffic initiated from hosts on the Internet is filtered with the WAN interface rules. Because all rules in pfSense software are stateful by default, a state table entry is created when traffic matches an allow rule. All reply traffic is automatically permitted by this state table entry.

The exception to this is Floating rules (*Floating Rules*), which can act on any interface using the inbound, outbound, or both directions. Outbound rules are never required, because filtering is applied on the inbound direction of every interface. In some limited circumstances, such as a firewall with numerous internal interfaces, having them available can significantly reduce the number of required firewall rules. In such a case, apply egress rules for Internet traffic as outbound floating rules on the WAN interface to avoid having to duplicate them for every internal interface. The use of inbound and outbound filtering makes a configuration more complex and more prone to user error, but it can be desirable in specific applications.

Interface Groups

Interface groups, discussed in *Interface Groups*, are a method to place rules on multiple interfaces at the same time. This can simplify some rule configurations if similar rules are required on many interfaces in the same way. Interface group rules, like interface rules, are processed in the inbound direction only. The VPN tabs for OpenVPN, L2TP, and the PPPoE server are all special Interface groups that are automatically created behind the scenes.

For example, a group may be used for a collection of interfaces including all LAN or DMZ type interfaces, or for a group of VLANs.

Note: Interface groups are not effective with Multi-WAN because group rules cannot properly handle `reply-to`. Due to that deficiency, traffic matching a group rule on a WAN that does not have the default gateway will go back out the WAN with the default gateway, and not through the interface which it entered.

Rule Processing Order

There are three main classes of Layer 3 rules: Regular interface rules, Floating rules, and Interface Group rules (including VPN tab rules). The order of processing of these types is significant, and it works like so:

1. Floating Rules
2. Interface Group Rules
3. Interface Rules

The rules are ordered in that way in the actual ruleset, keep that in mind when crafting rules. For example, if an interface group contains a rule to block traffic, that rule cannot be overridden with an interface tab rule because the traffic has already been acted upon by the group rule, which was matched first in the ruleset.

The rules are processed until a match is found, however, so if a packet is *not* matched in the group rules, it can still be matched by an interface rule.

Another significant place this comes into play is with assigned OpenVPN interfaces. If an "allow all" rule is in place on the OpenVPN tab, it is matched with the group rules. This means the rules on the interface tab will not apply. This can be a problem if OpenVPN rules need to have `reply-to` in order to ensure certain traffic exits back via the VPN.

See also:

See *Ordering of NAT and Firewall Processing* for a more detailed analysis of rule processing and flow through the firewall, including how NAT rules come into play.

Automatically Added Firewall Rules

pfSense software automatically adds internal firewall rules for a variety of reasons. This section describes automatically added rules and their purpose.

Anti-lockout Rule

To prevent locking an administrator out of the web interface, pfSense enables an anti-lockout rule by default. This is configurable on the **System > Advanced** page under **Anti-lockout**. This automatically added rule allows traffic from any source inside the network containing the rule, to any firewall administration protocol listening on the LAN IP address. For example, it grants access to TCP port 443 for the GUI, TCP port 80 for the GUI redirect, and TCP port 22 if SSH is enabled. If the GUI port has been changed, the configured port is the one allowed by the anti-lockout rule.

In security-conscious environments, the best practice is to disable this rule and configure the LAN rules, so only an alias of trusted hosts can access the administrative interfaces of the firewall. A better practice yet is to not allow access from the LAN but only from an isolated administrative management network.

Restricting access to the administrative interface from LAN

First, to configure the firewall rules as desired to restrict access to the required management interface(s). In this typical use case example, both SSH and HTTPS are used for management, so create a "ManagementPorts" alias containing these ports (Figure *Alias for Management Ports*).

Then create an alias for hosts and/or networks that will have access to the management interfaces (Figure *Alias For Management Hosts*).

The resulting aliases are shown in Figure *Alias List*.

Then the LAN firewall rules must be configured to allow access by the previously defined hosts, and deny access to all else. There are numerous ways to accomplish this, depending on specifics of the environment and how egress filtering

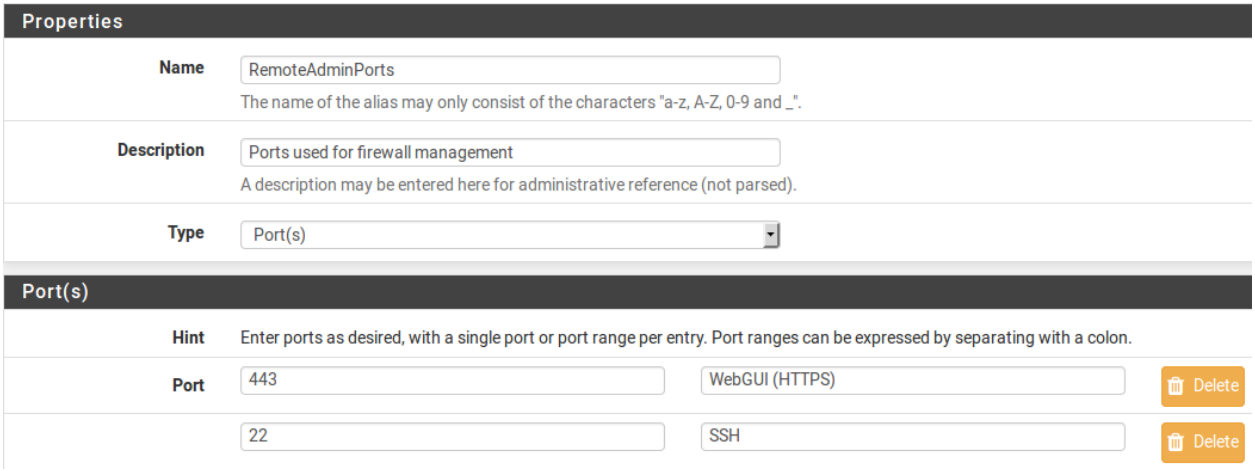


Fig. 7: Alias for Management Ports

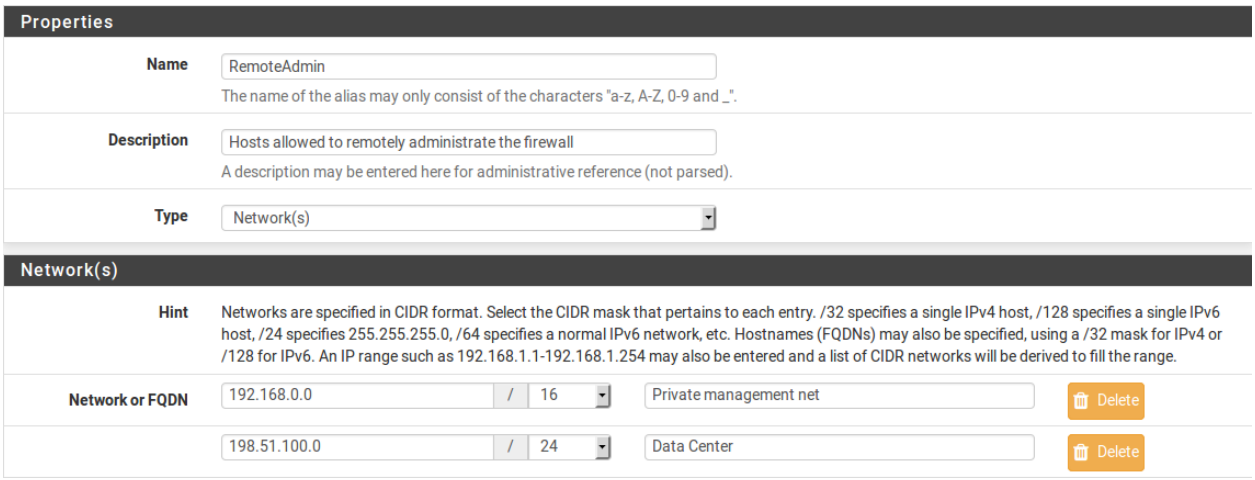


Fig. 8: Alias For Management Hosts

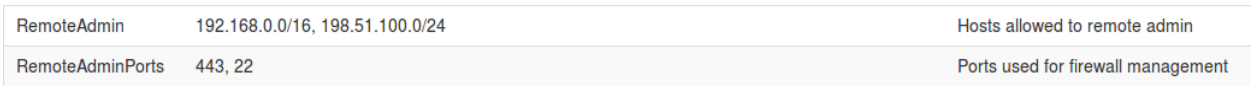


Fig. 9: Alias List

is handled. Figure *Example Restricted Management LAN Rules* show two examples. The first allows DNS queries to the LAN IP address, which is needed if the DNS Resolver or DNS Forwarder are enabled, and also allows LAN hosts to ping the LAN IP address. It then rejects all other traffic. The second example allows access from the management hosts to the management ports, then rejects all other traffic to the management ports. Choose the methodology that works best for the network environment in question. Remember that the source port is not the same as the destination port.

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
☐ ✔	0/0 B	IPv4 TCP/UDP	10.0.0.0/8	*	LAN address	53 (DNS)	*	none		Allow internal network to query the DNS Resolver
☐ ✔	0/0 B	IPv4 ICMP echoreq	10.0.0.0/8	*	LAN address	*	*	none		Allow internal network to ping the LAN IP Address
☐ ✔	0/0 B	IPv4 TCP	RemoteAdmin	*	LAN address	RemoteAdminPorts	*	none		Allow access to firewall management
☐ ✋	0/0 B	IPv4 *	*	*	LAN address	*	*	none		Reject everything else to the LAN IP address
☐ ✔	0/2.59 MiB	IPv4 *	10.0.0.0/8	*	*	*	*	none		LAN Traffic

Fig. 10: Example Restricted Management LAN Rules

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
☐ ✔	0/0 B	IPv4 TCP	RemoteAdmin	*	LAN address	RemoteAdminPorts	*	none		Allow access to firewall management
☐ ✋	0/0 B	IPv4 TCP	*	*	LAN address	RemoteAdminPorts	*	none		Reject access to firewall management from other host
☐ ✔	0/2.59 MiB	IPv4 *	10.0.0.0/8	*	*	*	*	none		LAN Traffic

Fig. 11: Restricted Management LAN Rules Alternate Example

Once the firewall rules are configured, disable the GUI anti-lockout rule on the **System > Advanced** page (Figure *Anti-Lockout Rule Disabled*). Check the box and click **Save**.

Note: If the management interface can no longer be accessed after disabling the anti-lockout rule, the firewall rules were not configured appropriately. Re-enable the anti-lockout rule by using the **Set Interface(s) IP address** option at the console menu, then choose to reset the LAN IP address. Set it to its current IP address, and the rule will automatically be re-enabled.

Anti-lockout

☒ Disable webConfigurator anti-lockout rule

When this is unchecked, access to the webConfigurator on the LAN interface is always permitted, regardless of the user-defined firewall rule set. Check this box to disable this automatically added rule, so access to the webConfigurator is controlled by the user-defined firewall rules (ensure a firewall rule is in place that allows access, to avoid being locked out!) *Hint: the "Set interface(s) IP address" option in the console menu resets this setting as well.*

Fig. 12: Anti-Lockout Rule Disabled

Anti-spoofing Rules

pfSense software uses the `antispoof` feature in PF to block spoofed traffic. This provides Unicast Reverse Path Forwarding (uRPF) functionality as defined in RFC 3704. The firewall checks each packet against its routing table, and if a connection attempt comes from a source IP address on an interface where the firewall knows that network does not reside, it is dropped. For example, a packet coming in WAN with a source IP address of an internal network is dropped. Anything initiated on the internal network with a source IP address that does not reside on the internal network is dropped.

Block Private Networks

The **Block private networks** option on the WAN interface automatically puts in a block rule for RFC 1918 subnets. Unless private IP space is in use on the WAN, enable this option. This only applies to traffic initiated on the WAN side. Local clients may still reach hosts on private networks from the inside of the firewall. This option is available for any interface, but is generally only used on WAN type interfaces. A similar rule can be created manually to block private networks on interfaces by creating an alias containing the RFC 1918 subnets and adding a firewall rule to the top of the interface rules to block traffic with a source matching that alias. (See *Private IP Addresses* for more information about private IP addresses.)

Block Bogon Networks

Bogon networks are those which should never be seen on the Internet, including reserved and unassigned IP address space. The presence of traffic from these networks can indicate either spoofed traffic or an unused subnet that has been hijacked for malicious use. Bogon lists are intended to filter invalid traffic from the Internet (e.g. on WANs) coming to the firewall for cases where the source cannot be otherwise filtered or validated, such as for public services. If rules on an interface only allow from specific remote sources, bogon blocking does not offer any benefit. pfSense software provides two bogons lists that are updated as needed, one for IPv4 bogon networks and one for IPv6 bogon networks.

Warning: Blocking bogon networks is not suited for use on local/private interfaces such as LAN. Blocking bogon networks on local interfaces can be harmful as they will block traffic which is necessary for proper local network operations, especially for IPv6. If local interfaces have proper rules which only allow from specific local sources, bogon blocking is unnecessary.

The firewall fetches an updated bogons list on the first day of each month from Netgate servers. The script runs at 3:00 a.m. local time, and sleeps a random amount of time up to 12 hours before performing the update. This list does not change frequently, and new IP address assignments are removed from the bogons list months before they are allocated, so a monthly update is adequate. To automatically update the list more frequently, change the **Update Frequency** for bogons under **System > Advanced** on the **Firewall & NAT** tab.

Note: The bogons list for IPv6 is quite large, and may not load if there is not enough memory in the system, or if the maximum number of table entries is not large enough to contain it. See *Firewall Maximum Table Entries* for information on changing that value.

See also:

For information on troubleshooting bogon updates and forcing manual updates, see *Troubleshooting Bogon Network List Updates*.

IPsec

When a site to site IPsec connection is enabled the firewall automatically adds rules which pass traffic necessary for the tunnel to connect and pass traffic. See *IPsec and firewall rules* for details.

Due to the nature of policy routing any traffic that matches a rule specifying a gateway will be forced out to the Internet and will bypass IPsec processing. Rules are added automatically to negate policy routing for traffic destined to remote VPN subnets, but they do not always have the intended effect. To disable the automatic negation rules, see *Disable Negate rules* and add a firewall rule at the *top* of the rules on the internal interface to pass traffic to the VPN *without a gateway set*.

See also:

Automatically added IPsec rules are discussed in further depth in *IPsec and firewall rules*.

Default Deny Rule

Rules that do not match any user-defined rules nor any of the other automatically added rules are silently blocked by the default deny rule (as discussed in *Default Deny*).

14.1.7 Configuring Firewall Rules

When configuring firewall rules in the pfSense® software GUI under **Firewall > Rules**, many options are available to control how the firewall matches and controls packets. This section covers each of these options.

See also:

- *Ordering of NAT and Firewall Processing*
- *Floating Rules*

Action

This option specifies whether the rule will *pass*, *block*, or *reject* packets.

Pass
: The firewall will allow packets matching this rule to pass.

  If the rule has state tracking enabled, the firewall creates a state table entry for the first packet of a connection. This state table entry allows related return packets to pass back through.

  See also:

  Stateful Filtering

Block
: The firewall will discard packets matching this rule.

Reject
: The firewall will discard packets matching this rule and, for compatible protocols, the firewall will send a message back to the source indicating that it refused the connection.

See also:

See *Block vs. Reject* for a deeper description of these options and for assistance choosing between Block and Reject.

Disabled

Disables a rule without removing it from the rule list. This entry will appear faded in the rule list to indicate its inactive state.

Interface

The interface receiving packets to be matched by this rule. The GUI pre-sets this value to match the interface tab from which the user added or edited the rule. Changing this value will move the rule to the interface tab matching the new value.

Note: Rules on interface and group tabs only filter packets on the interface where packets *enter* the firewall (ingress). In other words, the interface where hosts *initiated* those packets.

For example, rules on the **LAN** tab match packets initiated from hosts on the LAN which pass through the firewall, such as connections to hosts on the Internet or other remote networks.

Address Family

The address family this rule will match: *IPv4*, *IPv6*, or both *IPv4+IPv6*. The rule will only match and act upon packets with the selected address family.

Tip: Firewall rules can utilize aliases which contain both IPv4 and IPv6 addresses. Rules will use whichever alias entries match its address family.

NAT64

When **Address Family** is set to *IPv6* the GUI displays a checkbox to **Enable NAT64** which changes this rule into a *NAT64* rule. NAT64 rules perform translation to allow IPv6-only hosts to reach IPv4 resources. Despite acting in a role similar to Outbound NAT, these rules are placed on **internal** interface *firewall* rules where packets *enter* the firewall, e.g. on the LAN tab.

When **Enable NAT64** is checked, the GUI displays an **Address Family Translation** section with the following option:

Source
: Configures the external **source** address to which this rule will translate IPv4-mapped connections. For example, the external WAN IPv4 address. This is typically left on *Automatic* which is the default, but can be set to a specific address or VIP as needed.

> **Warning:** While NAT64 is technically compatible with Outbound NAT, the NAT64 translation happens before Outbound NAT so it is not practical. See *NAT64 and other NAT* for details.

Tip: NAT64 is compatible with policy routing. See *NAT64 and Policy Routing* for details.

See also:

- *NAT64*

- *NAT64 Configuration Recipe*

Protocol

The IP protocol this rule will match, such as *TCP* or *UDP*. The drop-down list contains several common protocols.

Most options only match packets based on the IP protocol encoded in the packet, however some protocols have additional filtering behaviors:

TCP, UDP, SCTP
: These protocols have source and destination port numbers. When a rule is configured for one of these protocols, the GUI displays controls for ports.

TCP/UDP
: Matches both *TCP* and *UDP* packets with a single rule.

ICMP
: When a rule is configured for ICMP, the GUI displays an *ICMP Subtypes* option to filter specific types of ICMP messages.

Note: This field defaults to *TCP* for a new rule because it is a common choice and the GUI will display fields that users expect to see, such as source and destination ports.

To make the rule apply to any protocol, change this option to *Any*. A common mistake in creating a new ruleset is accidentally creating only TCP rules and then not being able to pass non-TCP packets such as ICMP, DNS, etc.

ICMP Subtypes

This multi-select list contains all possible ICMP subtypes which this rule can match. The GUI displays this option when a rule has the **Protocol** option set to *ICMP*.

When passing ICMP packets, the most secure practice is for rules to only pass required subtypes when feasible. The most common use case is for rules to pass only the *Echo Request* subtype which will pass ICMP "ping" packets.

Tip: Allowing an ICMP subtype of *any* is acceptable in modern networks. ICMP has a bad reputation in the past, but it is generally beneficial and has overcome that reputation on modern networks. Some older equipment may still have issues, however.

Source

The source IP address, subnet, or alias that this rule will match.

The **Source** option contains several different pre-defined types of sources:

Any
: Matches any address.

Address or Alias
: Matches a single IP address or *alias name*.

  When a rule is configured for this value, the **Source Address** field allows entering an alias name.

  See also:

  Aliases

Network
Matches a range of addresses defined by an IP address and CIDR mask/prefix length.

PPPoE Clients
Macro which will match packets from the client address range for the PPPoE server if the PPPoE server is enabled.

L2TP Clients
Macro which will match packets from the client address range for the L2TP server if the L2TP server is enabled.

Interface Address
Macros which match all IP addresses configured on a firewall interface. This includes IP addresses of any address family on the selected interface, such as static IP addresses, *Virtual IP Addresses*, and dynamic IP addresses obtained from DHCP or PPPoE.

This list contains entries for each interface on the firewall.

Interface Subnets
Macros which match networks directly attached to a firewall interface. This includes networks for IP addresses of any address family including static IP addresses, *Virtual IP Addresses*, and dynamic IP addresses.

This list contains entries for each interface on the firewall.

> **Warning:** The *Interface Subnets* macro **does not** match networks reachable *through* an interface, only networks for IP addresses configured on an interface.

> **Warning:** The *WAN subnets* value **does not** mean "The Internet" or "any remote host", it only matches the network(s) of the WAN interface IP address(es).

Invert Match

Invert Match will negate the matching behavior, so the rule will match all packets *except* those with the specified **Source** value.

Note: **Invert Match** can lead to undesired rule behavior when configured in combination with *<interface> subnets* macros, such as *LAN subnets*, when the interface also uses Virtual IP addresses.

When configured in this manner the firewall will match packets against the interface network **OR** the VIP networks. For example, given a Subnet of `192.168.1.0/24`, a VIP of `10.0.0.1/32`, and a rule with a negated interface macro such as `pass on $LAN from any to ! $LAN_net`, packets destined to `192.168.1.100` will pass because the destination IP address does not match the VIP.

Source Port Range

Rules configured to match the TCP, UDP, or SCTP protocols can also match based on the source port of a packet.

The GUI hides this option behind a  **Display Advanced** button as it is rarely correct to match based on source port. In nearly all cases the source port must remain set to *any* as nearly all clients source TCP, UDP, and SCTP connections from a random port in the ephemeral port range (typically between 1024 through 65535).

The source port is almost never the same as the destination port, and it should never be configured as such unless the application in use is known to employ this atypical behavior.

It is typically safe to define a source port as a range from `1024` to `65535`.

Destination

The destination IP address, subnet, or alias that this rule will match. This operates the same as the **Source** option, but it checks the packet destination.

See also:

Source

The **Destination** field contains one additional macro:

This firewall (self)
: Matches all IP addresses on all firewall interfaces.

Destination Port Range

Rules configured to match TCP, UDP, or SCTP protocols can also match based on the destination port, port range, or an alias containing ports.

Configuring a destination port is necessary in many cases as it is more secure than allowing *any* port and the destination port is usually known in advance based on specific network services.

The drop-down lists contain common port values. Select *(other)* to enter a numerical port manually or use a port alias.

Tip: A rule will match a continuous range of ports starting at the lower port in the **From** field and ending at the higher port value in the **To** field, inclusive.

Log

When checked, if this rule matches a packet it will create a log entry in the firewall log.

See also:

Logging Practices

Description

Text describing the rule, such as its intended behavior or name of a service. The best practice is to clearly describe the purpose of the rule in this field.

The description is optional and does not affect functionality of the rule. The maximum length is 52 characters.

Advanced Options

Options in this section are less common or have functionality confusing to new users. As such, the GUI hides them by default. Click  **Display Advanced** to show all the advanced options. If an option in this section of the page has been set, the GUI will automatically display this section when the rule is loaded in the future.

Source OS

Attempts to match a packet by guessing which operating system (OS) is running on the host initiating a connection. This is only possible on rules which match TCP packets.

This task is handled by passive OS fingerprinting ("p0f"). The p0f feature of PF determines the OS in use by comparing characteristics of the TCP SYN packet which initiates TCP connections with a fingerprints file.

Note: It is possible for users to change the fingerprint of one OS to mimic a different OS, especially with an open source OS such as a BSD or Linux variant. This isn't easy, but if a network contains technically proficient users with administrator or root level access to devices, it is possible.

Diffserv Code Point

Differentiated Services Code Point is a packet header a device or application can use to indicate how it prefers routers treat the packet. The most common use of this value is for quality of service or traffic shaping purposes.

Differentiated Services Code Point is often shortened to *Diffserv Code Point* or abbreviated as *DSCP* and sometimes referred to as the *TOS field*.

The **Diffserv Code Point** drop-down field sets the DSCP value that this rule will match. There are numerous options, each with special meaning. Consult the documentation for the device or application originating the packets for more detail on which values rules must match.

Note: This option only reads and matches the DSCP value. It **does not** set a value in packets.

The device or application generating the packets will set the DSCP field value in the packets it creates. For example, Asterisk sets DSCP values via its `tos_sip` and `tos_audio` configuration parameters. After that point it is up to the firewall and other interim routers to match and queue or act on the packets.

The downside of DSCP is that it assumes routers are compatible with, or act on, the DSCP field, which may or may not be the case. Different routers may treat the same DSCP value in unintended or mismatched ways. Worse yet, some routers will clear the DSCP field in packets entirely before passing them to the next hop. Also, the way PF matches packets, the DSCP value must be set on the first packet of a connection creating a state, as PF does not inspect each packet individually once it has created a state.

IP Options

When set, packets matching this rule can have values set for IP options. By default, PF does not match packets which contain IP options in order to deter OS fingerprinting, among other reasons.

Tip: Check this box to pass IGMP or other multicast packets containing IP options.

Disable Reply-To

When set, prevents the firewall from adding the `reply-to` keyword on this rule.

By default, the firewall adds the `reply-to` keyword to rules on WAN type interfaces to ensure that packets which enter a WAN will also leave via that same WAN. In certain cases this behavior is undesirable, such as when the firewall must route packets via a separate firewall/router within the network on the WAN interface. In these cases, check this option rather than disabling `reply-to` globally.

Tag and Tagged

The **Tag** and **Tagged** fields are useful in concert with floating rules. Using tags, the firewall can mark a packet with a specific string as it enters an interface and then act differently on a matched packet on the way out with a floating rule.

See also:

Marking and Matching

Maximum state entries this rule can create

Limits the maximum number of connection states, total, that this rule will allow. If this rule matches a packet for a new connection while it is at its connection limit, the firewall will skip this rule during evaluation. If a later rule matches, the firewall applies the action of that rule to the packet, otherwise it hits the default deny rule. Once the number of active connections permitted by this rule drops below this connection limit, the rule can match packets for new connections again.

Maximum number of unique source nodes

Limits the total number of unique source IP addresses which may simultaneously have connection states created by this rule. Each source IP address is allowed an unlimited number of connections, but the total number of distinct source IP addresses is restricted to this value.

Maximum number of established connections per host

Limits access based on established connection states per host. This value can limit a rule to a specific number of connections per source host (e.g. `10`), instead of a global connection total. This option controls how many fully established (completed handshake) TCP connections the rule will allow per host. This option is only available for use with TCP connections.

Maximum state entries per host

Limits connection states in total per host without considering if the connections are established. This setting works similar to the established count above, but it can work with any protocol.

Maximum new connections / per second

Limits the amount of new connection states this rule can create in a given time period. This option is only available for use with TCP connections. The **Max. src. conn. Rate** field sets the number of states and the **Max. src. conn. Rates** field sets the time period.

Rules can use this feature to attempt preventing a high TCP connection rate from overloading a server or the state table on the firewall. For example, a rule can limit incoming connections to a mail server, reducing the burden of being overloaded by spambots. It can also be used on outbound rules to set limits that can prevent any single host from filling up the state table on the firewall or making too many rapid connections, behaviors which are common with viruses.

The firewall will block any IP address exceeding the specified number of connections within the given time frame for one hour. Behind the scenes, this is handled by the `virusprot` table, named for its typical purpose of virus protection.

State timeout in seconds

Overrides the default timeout for states created by this rule. Any inactive connections will be closed when the connection has been idle for this amount of time.

The default state timeout depends on the firewall optimization algorithm in use. The optimization choices are covered in *Firewall Optimization Options*

Note: This option only controls states created in the inbound direction on the given interface, so it has limited usefulness on its own. Outbound packets for a connection will still have the default state timeout as they egress. To fully apply this new timeout, create a matching floating rule in the outbound direction with a similar state timeout value.

TCP Flags

By default, pass rules for TCP packets only check for the TCP SYN flag to be set, out of a possible list of SYN and ACK flags. This means that for a packet to match the rule, the SYN flag **must** be set and the ACK flag **must not** be set. Other flags are ignored.

This set of controls enables rules to match different flag combinations to account for more complex scenarios, such as working around asymmetric routing or other non-traditional combinations of packet flow.

The **Set** row controls which flags must be **set** to match the rule. The **Out of** row defines the list of flags that will be consulted on the packet to look for a match.

Tip: Flags which are **not checked** in the **Set** row but are **checked** in the **Out of** row **must not** be set in packets to match this rule.

The meanings of the most commonly used flags are:

SYN
: Synchronize sequence numbers. Indicates a new connection attempt.

ACK
Indicates ACKnowledgment of data. These are replies to let the sender know data was received OK.

FIN
Indicates there is no more data from the sender, closing a connection.

RST
Connection reset. This flag is set when replying to a request to open a connection on a port which has no listening daemon. Can also be set by firewall software to turn away undesirable connections.

PSH
Indicates that data should be pushed or flushed, including data in this packet, by passing the data up to the application.

URG
Indicates that the urgent field is significant, and this packet should be sent before data that is not urgent.

To allow TCP packets with any combination of flags set, check **Any Flags**.

No pfsync

Prevents states created by this rule from synchronizing data to high availability (HA) peers via *pfsync*. This isolates states for this rule to only this HA node.

State Policy

Optionally overrides the default *Firewall State Policy* for connection states created by this rule. This is only effective when the *State Type* for this rule is a type that creates states.

The available options are:

Use Global Default
Uses the global default state policy configured at **System > Advanced**, **Firewall & NAT** tab, in the **Advanced Options** section (*Firewall State Policy*).

Interface Bound States
Binds states to interfaces so that when a packet is inspected to determine if it matches an existing state, it must be on the interface where the state was created.

Floating States
Does not bind states to interfaces, allowing packets matching the source and destination of a state to pass no matter which interface they are traversing.

See also:

For more information on how these policies work, see *Firewall State Policy*.

State Type

Controls how this rule will perform state tracking for connections created by packets matching this rule.

Keep
The firewall will create and maintain a state table entry for packets passed by this rule.

This is the default, and the best choice in most situations.

Sloppy State
Similar to *Keep*, but the firewall will perform less strict state comparison checks. This is intended for scenarios with asymmetric routing.

When the firewall can only see half the packet flow of a connection, the default validity checks for state tracking will fail, and the firewall will block packets which may be part of active connections. Mechanisms in PF that prevent certain kinds of attacks will not trigger during a sloppy state check.

Synproxy
Instructs PF to proxy incoming TCP connections.

TCP connections start with a three-way handshake. The first packet of a TCP connection is a SYN from the source, which elicits a SYN ACK response from the destination, then an ACK in return from the source to complete the handshake. Normally the host behind the firewall will handle this on its own, but synproxy state has the firewall complete this handshake instead. This helps protect against one type of Denial of Service attack, SYN floods. This is typically only used with rules on WAN interfaces.

This type of attack is best handled at the target OS level today, as every modern operating system includes capabilities of handling this on its own. Because the firewall can't know what TCP extensions are compatible with the back-end host, when using synproxy state, it announces no TCP extensions. This means connections created using synproxy state will not use window scaling, SACK, nor timestamps which will lead to significantly reduced performance in most all cases.

This option can be useful when opening TCP ports to hosts that do not handle network abuse well, where top performance isn't a concern.

None
This rule will not create states for packets matching the rule. This is only necessary in highly specialized advanced scenarios, which are not covered in this documentation as they are exceedingly rare.

Note: Setting *None* on interface tab or group rules only affects packets in the inbound direction. As such, it is not very useful on its own since the outbound rules will still create a state as the packets egress. Rules on interface or group tabs must be paired with a floating rule in the outbound direction which also has the same option set.

Packet Flow Data

Optionally overrides the default *Firewall Packet Flow Data* configuration for tracking data for states created by this rule. This is only effective when the *State Type* for this rule is a type that creates states.

Note: pfSense® Plus software version 24.03 or later is required to use the Packet Flow Data feature. This feature is not available on pfSense CE Software.

Global configuration for Packet Flow Data is at **Firewall > Packet Flow Data**.

See also:

Read the option descriptions in *Global Packet Flow Options* for information on how this behaves by default.

There are three options available for **Packet Flow Data** here in the rule configuration:

Use global default
Honor the global default value for tracking packet flow data.

See also:

Global Packet Flow Options

Always track Packet Flow Data
Always tracks packet flow data for states created by this rule, no matter what default behavior is configured. This is useful for tracking specific data when the default is *not* to track.

Never track Packet Flow Data
Never tracks packet flow data for states created by this rule, no matter what default behavior is configured. This is useful for excluding data from monitoring when the default is to track everything.

No XMLRPC Sync

Checking this box prevents this rule from synchronizing to other High Availability cluster members via XMLRPC.

See also:

High Availability

Warning: This does not prevent a rule on a secondary node from being overwritten by the primary.

VLAN Priority (Match and Set)

As described in *VLAN Priority*, these values are 802.1p Priority Code Point (PCP) values which can be used to match and tag packets with a specific quality of service priority. Unlike DSCP, 802.1p operates at layer 2 with VLANs. However, like DSCP, the upstream router must also be compatible with 802.1p for it to be useful.

There are two options in this section. The first matches an 802.1p field, so the firewall can act on it. The second injects an 802.1p tag into a packet as it passes through this firewall. Some ISPs may require an 802.1p tag to be set in certain areas, such as France, in order to properly handle voice/video/data on segregated VLANs at the correct priority to ensure quality.

See also:

Refer to *VLAN Priority* for an explanation of each code and associated priority value.

Schedule

Configures a schedule specifying the days and times this rule will be active. Outside the scheduled time, the rule is effectively disabled and the firewall skips the rule.

The default value is *none* which means the rule will always be enabled.

See also:

Time Based Rules

Gateway

A gateway or gateway group through which the firewall will deliver packets matching this rule.

When set to *Default*, the firewall will consult the routing table to determine the next hop for a packet.

See also:

- *Policy routing*
- *Routing*

In/Out Pipe (Limiters)

Traffic shaper Limiters this rule will use to apply bandwidth limits to packets entering this interface (**In**) and exiting this interface (**Out**).

See also:

Limiters.

Ackqueue/Queue

ALTQ traffic shaper queues into which this rule will place matching packets entering and exiting this interface.

See also:

Traffic Shaper

Rule Information

When a firewall rule, port forward rule, or outbound NAT rule is created or updated the firewall records the login name of the user who modified the rule, the IP address from which they logged in, and timestamps on the rule to track when that user created and/or last changed the rule in question. If the firewall automatically created the rule, the note includes which action created the rule.

An example of a rule update tracking block is shown in Figure *Firewall Rule Time Stamps*, which is visible when editing a firewall rule at the bottom of the rule editing screen.



Fig. 13: Firewall Rule Time Stamps

Rule Tracking ID

Figure *Firewall Rule Time Stamps* also displays the rule **Tracking ID** which is a unique identifier the firewall assigns to a rule. This identifier allows the firewall to locate the rule and correlate actions between the ruleset, firewall logs, and other data tracking purposes.

Note: This value is created and maintained by the firewall itself, the tracking ID is not intended to be edited by administrators.

The firewall generates some rules automatically and/or dynamically, and tracking IDs for such rules may change when the ruleset reloads. Tracking IDs for manually created rules are static.

14.1.8 Ethernet (Layer 2) Rules

pfSense® Plus software versions 23.05 and later include rule-based pass/block filtering of packets based on Ethernet (Layer 2) header attributes. These are known as **Ethernet Rules**.

Processing of these rules is **not** enabled by default and can be togged under **System > Advanced**, **Firewall & NAT** tab.

When enabled, Ethernet rules are managed at **Firewall > Rules**, **Ethernet** tab.

See also:

- *WAN Connectivity with 802.1X Authentication Bridging and VLAN 0 PCP Tagging*
- *Ethernet Rules on Bridge Interfaces*

Ethernet Rules Overview

Layer 2 Interfaces

Ethernet rules are capable of operating on Layer 2 (L2) header information which is not visible to traditional firewall rules. To accomplish this, Ethernet rules operate at Layer 2 (L2) and are **only** processed on interfaces which carry L2 data.

Traditional Ethernet and VLAN interfaces will work, but certain types of VPN and tunneling interfaces will not. For example, OpenVPN in TAP mode can carry L2 information, while IPsec, WireGuard, and OpenVPN in TUN mode cannot.

As a general rule, if an interface has a MAC address, then it is capable of carrying L2 data.

Warning: The firewall will not generate an error if the user attempts to apply an Ethernet rule on an interface which is not capable of L2. These rules can never be matched as a non-L2 interface is not capable of triggering Ethernet rules.

Captive Portal

Captive Portal uses Ethernet rules behind the scenes to pass users through the portal. This works no matter what the current state of the Ethernet rules option is as these rules are managed automatically and not via the **Ethernet** tab.

Warning: Be careful adding manual Ethernet rules to interfaces involved in Captive Portal. Manual Ethernet rules are processed before Captive Portal rules, so there is a potential for interference.

Stateless

Ethernet rules do not keep state. As such, while block rules can work on their own, when making exceptions to blocks it is best to add rules in pairs to cover both the inbound and outbound direction, with the source and destination values on the rule reversed for the opposing direction.

Default Behavior

When Ethernet rules are inactive, all L2 traffic (other than Captive Portal) is passed by default so it can then be processed at L3 by regular rules. This **does not** affect the behavior of L3 rules. Traffic is still blocked there by default as usual.

When Ethernet rules are active the firewall does not have an automatic hidden rule to block Ethernet traffic by default. This preserves the existing behavior and makes it less prone to be easily broken.

If an administrator wants to block all L2 traffic by default, they can first craft a ruleset with appropriate pass rules and then follow that with a rule to block any other Ethernet traffic. As the potential for mistakes and disruption is extremely high, this is not considered a best practice at this time.

Aliases

Ethernet rules can use Aliases for L3 source/destination matching, but there are no MAC Address type aliases at this time.

Enabling Ethernet Rules

To enable Ethernet rules:

- Navigate to **System > Advanced**, **Firewall & NAT** tab
- Locate the **Advanced Options** section
- Check **Enable Ethernet Filtering**
- Click **Save**

Managing Ethernet Rules

To manage Ethernet rules, navigate to **Firewall > Rules**, **Ethernet** tab. From there, rules are managed using the list view similar to other rules.

Note: The **Ethernet** tab only appears while the Ethernet Rules function is enabled. If the tab is not visible, *enable Ethernet Rules*.

Configuring Ethernet Rules

When editing an Ethernet rule the available options are similar to those found on *firewall rules* and *floating rules* with the following differences:

Protocol
A protocol specific to layer 2 for which this rule will apply.

ARP
Address Resolution Protocol

IPv4
IPv4 traffic

IPv6
IPv6 traffic

IEEE 802.1X
Network authentication traffic

VLAN (C-Tag)
Customer VLAN tag (e.g. first level)

VLAN (S-Tag)
Service VLAN tag (e.g. second level, double tagged)

Other
A protocol not listed in the drop-down, set manually in **Protocol Value**.

Protocol Value
To specify a protocol not in the list, enter its 16-bit hexadecimal **EtherType** (e.g. `0xffff`).

Source/Destination
Though Ethernet rules operate at L2, they can still act on the contents of the L3 source/destination (e.g. IPv4 or IPv6 addresses) in a packet using these fields.

Note: When setting a **Protocol** in addition to a **Source** or **Destination**, ensure the protocol matches the address family of the source/destination. For example, when using an IPv4 address in **Source**, either set the **Protocol** to *IPv4* or *Any*.

MAC Filtering
Match a packet based on the L2 **Source MAC Address** or **Destination MAC Address**.

This option is in the **Advanced Options** section of the page.

Bridge To
When set, packets matching this rule will be sent out of the chosen interface without further processing. This can be used to send certain L2 packets out another interface, bypassing L3 rules (e.g. 802.1X authentication from an ISP).

This option is in the **Advanced Options** section of the page.

Package Rules

There is a plugin hook available for packages to add their own Ethernet rules.

In the package metadata, define a filter rule callback:

```
<filter_rules_needed>package_rules</filter_rules_needed>
```

And then in the package PHP include file, add a function which returns the rules:

```
function package_rules($ruletype) {
    if ($ruletype === 'ether') {
            return '# add an ether rule' . PHP_EOL;
    }
}
```

14.1.9 Floating Rules

Floating Rules are a special type of advanced rule that can perform complicated actions not possible with rules on interface or group tabs. Floating rules can act on multiple interfaces in the inbound, outbound, or both directions. The use of inbound and outbound filtering makes designing rules highly complex and prone to user error, but they can be desirable in certain challenging scenarios.

Tip: Most firewall configurations will never have floating rules, or only have floating rules added by the traffic shaper.

See also:

- *Configuring Firewall Rules*
- *Ordering of NAT and Firewall Processing*
- *Traffic Shaper*

Precautions/Caveats

Floating rules can be a lot more powerful than other rules, but also more confusing. With floating rules it is easier for administrators to make an error with unintended consequences when passing or blocking traffic, which can be dangerous.

The firewall does not automatically add `reply-to` on floating rules in the inbound direction as it does for individual interface rules. Thus, floating rules have the same problem as interface groups: Return traffic passed by states created from floating rules will always exit the WAN with the default gateway, a reply packet cannot automatically return out a non-default WAN through which it entered the firewall.

Given the relative unfamiliarity of most administrators with floating rules, they may not think to look on the **Floating** tab for rules when maintaining the firewall, which increases the difficulty of firewall administration.

Take care when considering the source and destination of packets depending on the inbound and outbound direction. For example, rules in the outbound direction on a WAN typically have a local source of the firewall (after NAT) and remote destination.

Potential Uses

The most common use of Floating rules is for ALTQ traffic shaping. Floating rules are the only type of rules which can match and queue traffic without explicitly passing the traffic.

Another use of floating rules is to control traffic egressing from the firewall itself. Floating rules can prevent the firewall from reaching specific IP addresses, ports, and so on.

Other common uses are to ensure that no traffic can exit from other paths into a secure network, no matter what rules exist on other interfaces. Blocking outbound toward a secure network from all but approved sources reduces the likelihood of later accidentally allowing traffic in through another unintended path. Similarly, floating rules can be used to prevent traffic destined for private networks from exiting a WAN interface, to prevent VPN traffic from leaking.

Floating rules are useful for completely enacting state timeouts, tag/match operations, "no state" rules, and "sloppy state" rules for asymmetric routing.

Processing Order

In the inbound direction, floating rules work essentially the same as interface or group rules except that they are processed first. Processing in the outbound direction is more complicated.

The firewall processes floating rules after NAT rules, so rules in the outbound direction on a WAN can never match a private IP address source if the firewall also applies outbound NAT to connections on that interface. By the time a packet hits the floating rule, the source address of the packet is the post-NAT WAN IP address. In most cases this limitation can be overcome by applying a tag to a packet inbound on the LAN and then matching that tag in an outbound floating rule (*Marking and Matching*).

The firewall processes floating rules before interface group rules and interface rules, so that must also be taken into consideration.

See also:

See *Ordering of NAT and Firewall Processing* for a more detailed analysis of rule processing and flow through the firewall, including how NAT rules come into play.

Floating Rule Configuration

Most options available for floating rules are identical to those found on interface and group tab rules. However, floating rules have a few differences in available options and available choices.

See also:

- *Configuring Firewall Rules*

Match Action

The *match* action is unique to floating rules. A rule with the *match* action will not pass or block a packet, but only match it for purposes of assigning traffic to queues or limiters for traffic shaping.

Quick

Quick controls whether the firewall stops processing rules when a packet matches this rule. Interface and group tab rules always behave in this manner, but on floating rules this behavior is optional. Without *Quick* checked, the rule will only take effect if no other rules match the packet. In other words, this option reverses the behavior of "first match wins" to be "last match wins".

Using this mechanism, administrators can craft a default action of sorts which will take effect only when no other rules match a packet, similar to the implicit default block rules on interfaces.

In most situations, the best practice is to check **Quick**. There are certain specific scenarios where leaving **Quick** unchecked is necessary, but they are rare.

Interface

The **Interface** selection for floating rules a multi-select control. With this control a rule can apply to one, multiple, or all possible interfaces. Ctrl-click on interfaces to select them one by one, or use other combinations of click/drag or shift-click to select multiple interfaces.

Direction

Floating rules are not limited to the inbound direction like interface rules, they have the following direction choices:

any
: The firewall will process this rule for both inbound and outbound packets.

in
: The firewall will process this rule for inbound packets.

out
: The firewall will process this rule for outbound packets.

  The *out* direction is useful for filtering traffic from the firewall itself, for matching other undesirable traffic trying to exit an interface, or for fully configuring "sloppy state" rules, "no state" rules, or alternate state timeouts.

Marking and Matching

Using the **Tag** and **Tagged** fields, an administrator can mark a connection with an interface tab rule and then match that connection in the outbound direction with a floating rule. This is a useful way to act on outbound WAN connections from a specific internal host which the firewall could not otherwise match due to NAT masking the source. It can also be used similarly to apply traffic shaping outbound on WAN for connections specifically tagged on the way into the firewall.

For example, on a LAN rule, use a short string in the **Tag** field to mark a packet from a source of `10.3.0.56`. Then on a floating rule, quick, outbound on WAN, use **Tagged** with the same string to act on the traffic matched by the LAN rule.

14.1.10 Time Based Rules

Time based rules allow firewall rules to activate during specified days and/or time ranges. Time based rules function the same as any other rule, except they are effectively not present in the ruleset outside their scheduled times.

Time Based Rules Logic

When dealing with time-based rules, the schedule determines when to apply the action specified in the firewall rule. When the current time or date is not covered by the schedule, the firewall acts as if the rule is not there. For example, a rule that passes traffic on Saturdays will only block it on other days if a separate block rule exists underneath it. The rules are processed from the top-down, the same as other firewall rules. The first match is used, and once a match is found, that action is taken if the rule is in schedule, and no other rules are evaluated.

Tip: Remember when using schedules that the rule will have **no effect** outside their scheduled times. The rule **will not** have its action reversed because the current time is not within the scheduled time. Failing to account for this behavior could result in giving clients unintended access outside the defined time ranges in a schedule.

Configuring Schedules for Time Based Rules

Schedules must be defined before they can be used on firewall rules. Schedules are defined under **Firewall > Schedules**, and each schedule can contain multiple time ranges. In the following example, a company wants to deny access to HTTP during business hours, and allow it all other times of the day.

Defining Times for a Schedule

To add a schedule:

- Navigate to **Firewall > Schedules**
- Click  **Add** to bring up the schedule editing screen, as seen in Figure *Adding a Time Range*.
- Enter a **Schedule Name**. This is the name that will appear in the selection list for use in firewall rules. Much like alias names, this name must only contain letters and digits, no spaces. For example: `BusinessHours`
- Enter a **Description** of this schedule, such as `Normal Business Hours`.
- Define one or more time ranges:
  - Set the **Month** by selecting a specific month and days, or by clicking the day of the week header for weekly recurring schedules.
  - Choose a **Start Time** and **Stop Time** which control when the rule is active on the selected days. The time cannot cross midnight on any day. A full day is *0:00* to *23:59*.
  - Enter an optional **Time Range Description** for this specific range, e.g. `Work Week`
  - Click  **Add Time** to add the choice as a range
  - Repeat **Month**, **Time**, and  steps for additional ranges
- Click **Save**

A schedule can apply to specific days, such as September 2, 2016, or to days of the week, such as Monday-Wednesday. To select any given day within the next year, choose the Month from the drop-down list, then click on the specific day or day numbers on the calendar. To select a day of the week, click its name in the column headers.

For this example, click on **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. This will make the schedule active for any Monday-Friday, regardless of the month. Now select the time for this schedule to be active, in 24-hour format. The hours for this example business are *9:00* to *17:00* (5pm). All times are given in the local time zone.

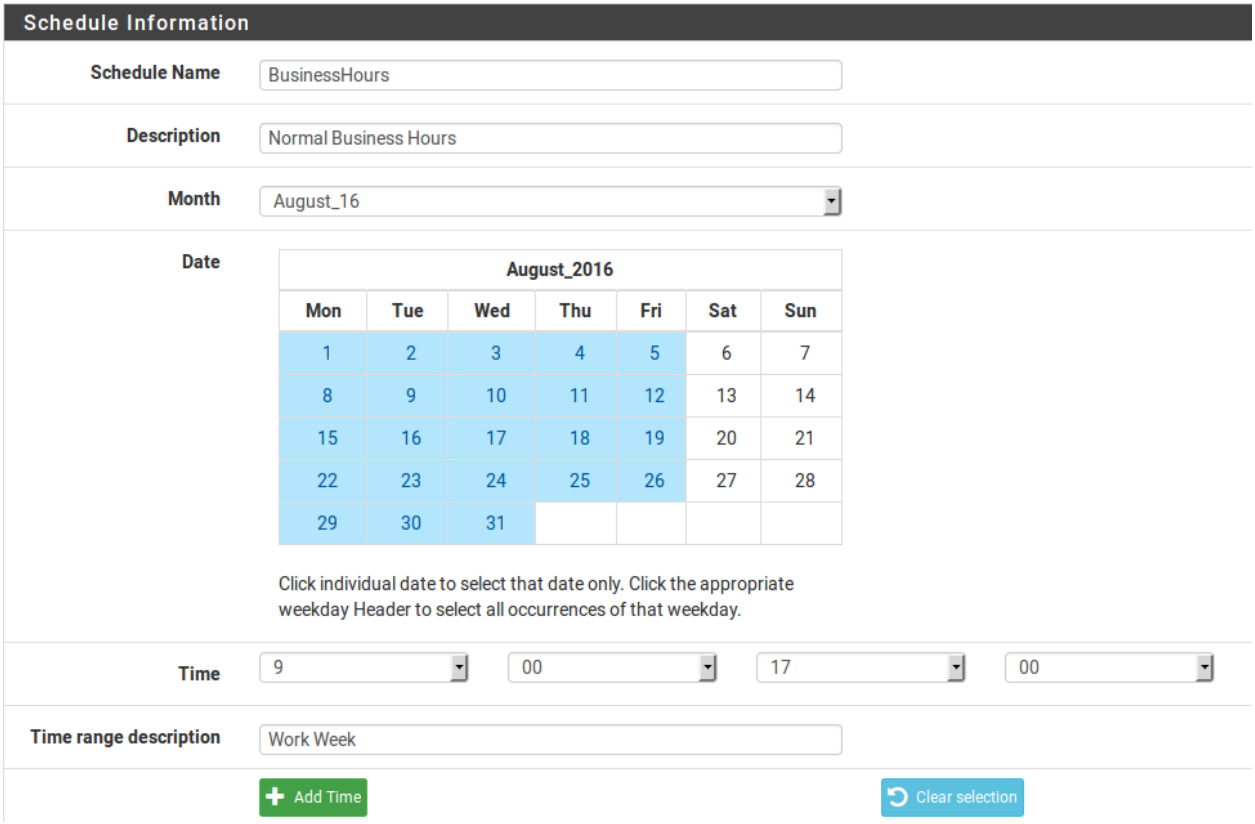


Fig. 14: Adding a Time Range

Once the time range has been defined, it will appear in the list at the bottom of the schedule editing screen, as in Figure *Added Time Range*.

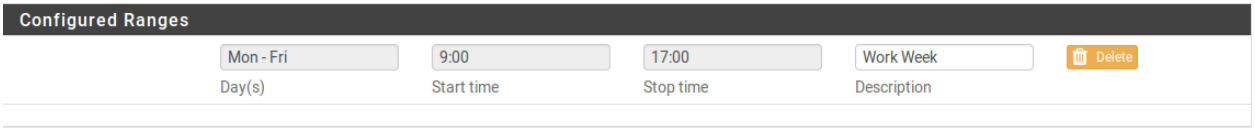


Fig. 15: Added Time Range

To expand on this setup, there may be a half day on Saturday to define, or maybe the shop opens late on Mondays. In that case, define a time range for the identical days, and then another range for each day with different time ranges. This collection of time ranges will be the full schedule.

Once the schedule entry has been saved, the browser will return to the schedule list, as in Figure *Schedule List After Adding*. This schedule will now be available for use in firewall rules.

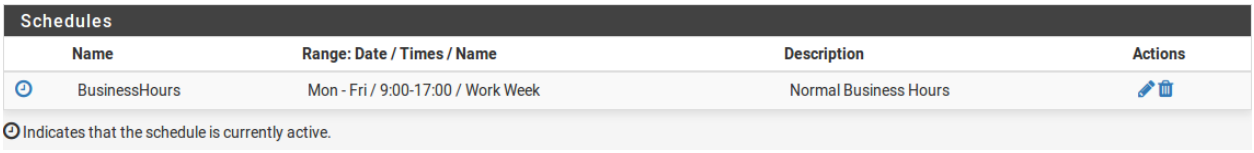


Fig. 16: Schedule List After Adding

Using the Schedule in a Firewall Rule

To create a firewall rule employing this schedule, create a new rule on the desired interface. See *Adding a firewall rule* and *Configuring Firewall Rules* for more information about adding and editing rules. For this example, add a rule to reject TCP traffic on the LAN interface from the LAN subnet to any destination on the HTTP port. In the advanced options for the rule, locate the **Schedule** setting and choose the *BusinessHours* schedule, as in Figure *Choosing a Schedule for a Firewall Rule*.

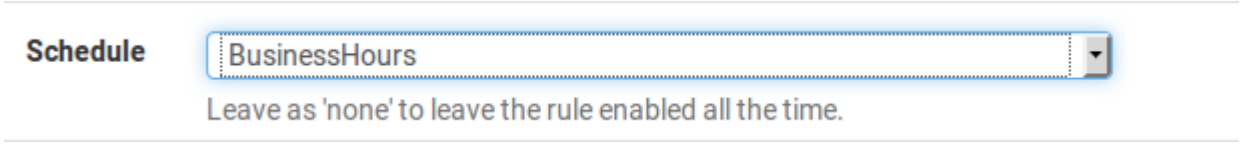


Fig. 17: Choosing a Schedule for a Firewall Rule

After saving the rule, the schedule will appear in the firewall rule list along with an indication of the schedule's active state. As shown in Figure *Firewall Rule List with Schedule*, this is a reject rule, and the schedule column indicates that the rule is currently in its active blocking state because it is being viewed at a time within the scheduled range. If the mouse cursor hovers over the schedule state indicator, a tooltip is displayed by the firewall showing how the rule will behave at the current time. Since this is being viewed inside the times defined in the "BusinessHours" schedule, this will say "Traffic matching this rule is currently being denied". If there is a pass rule that would match the traffic out on port 80 from the LAN net after this rule, then it would be allowed outside the scheduled hours.



Fig. 18: Firewall Rule List with Schedule

Now that the rule is defined, test it both inside and outside the scheduled times to ensure that the desired behavior is enacted.

Tip: By default, states are cleared for active connections permitted by a scheduled rule when the schedule expires. This shuts down access for anyone allowed by the rule while it was active. To allow these connections to remain open, check **Do not kill connections when schedule expires** under **System > Advanced** on the **Miscellaneous** tab.

14.1.11 Viewing the PF ruleset

pfSense® software handles translating the firewall rules in the GUI into a set of rules which can be interpreted by the packet filter (PF).

Generated Rules

The PF rules generated by the firewall are in `/tmp/rules.debug`. However, that file **cannot** be edited to make persistent changes - the firewall will overwrite it during the next filter reload event.

Note: There is rarely a need to manually edit firewall rules generated by the GUI. In most cases if it appears to be necessary, something is incorrect with the configuration.

If the generated rules truly must be edited, then the edits must be made to the source code which generates the ruleset in `/etc/inc/filter.inc`. Such changes will be lost when updating to a new version.

Interpreted Rules

PF can interpret the rules slightly differently than the way they were generated by the filter code. To view the rule set as has been interpreted by PF, use one of the following methods.

Using the *SSH console* or *Command Prompt* field in the GUI, run the following:

Show Firewall Rules:

```
# pfctl -sr
```

Show NAT rules:

```
# pfctl -sn
```

Show all:

```
# pfctl -sa
```

For more verbose output including rule counters, ID numbers, and so on, use:

```
# pfctl -vvsr
```

There may be additional rules in anchors from packages or features such as UPnP. To view these rules, use:

```
# pfSsh.php playback pfanchordrill
```

14.1.12 Methods of Using Additional Public IP Addresses

Methods of deploying additional public IP addresses vary depending on how the addresses are delegated, the size of the allocation, and the goals for the specific network environment. To use additional public IP addresses with NAT, for example, the firewall will need *Virtual IP Addresses*.

There are two options for directly assigning public IP addresses to hosts: Routed public IP subnets and bridging.

Choosing between routing, bridging, and NAT

Additional public IP addresses can be put to use by directly assigning them on the systems that will use them, or by using NAT. The available options depend on how the addresses are allocated by the ISP.

Additional static IP addresses

Methods of using additional static public IP addresses vary depending on the type of assignment. Each of the common scenarios is described here.

Single IP Subnet on WAN

With a single public IP subnet on WAN, one of the public IP addresses will be on the upstream router, commonly belonging to the ISP, and another one of the IP addresses will be assigned as the WAN IP address on pfSense® software. The remaining IP addresses can be used with either NAT, bridging or a combination of the two.

To use the addresses with NAT, add Proxy ARP, IP alias or CARP type Virtual IP addresses.

To assign public IP addresses directly to hosts behind the firewall, a dedicated interface for those hosts must be bridged to WAN. When used with bridging, the hosts with the public IP addresses directly assigned must use the same default gateway as the WAN of the firewall: the upstream ISP router. This will create difficulties if the hosts with public IP addresses need to initiate connections to hosts behind other interfaces of the firewall, since the ISP gateway will not route traffic for internal subnets back to the firewall.

Figure *Multiple Public IP addresses In Use Single IP Subnet* shows an example of using multiple public IP addresses in a single block with a combination of NAT and bridging.

See also:

For information on configuration, NAT is discussed further in *Network Address Translation*, and bridging in *Bridging*.

Small WAN IP Subnet with Larger LAN IP Subnet

Some ISPs will allocate a small IP subnet as the "WAN side" assignment, sometimes called a transport or interconnect network, and route a larger "inside" subnet to the firewall. Commonly this is a /30 on the WAN side and a /29 or larger for use inside the firewall. The service provider router is assigned one end of the /30, typically the lowest IP address, and the firewall is assigned the higher IP address. The provider then routes the second subnet to the WAN IP address of the firewall. The additional IP subnet may be used by the firewall on a routed LAN or OPT interface with public IP addresses directly assigned to hosts, with NAT using Other type VIPs, or a combination of the two. Since the IP addresses are routed to the firewall, ARP is not needed so VIP entries are not necessary for use with NAT.

Because pfSense software is the gateway on the local segment, routing from the public local subnet hosts to LAN is much easier than in the bridged scenario required when using a single public IP subnet. Figure *Multiple Public IP Addresses Using Two IP Subnets* shows an example that combines a routed IP subnet and NAT. Routing public IP addresses is covered in *Routing Public IP Addresses*, and NAT in *Network Address Translation*.

WAN IP Address Assignments	
WAN Subnet	192.0.2.128/29
WAN Gateway	192.0.2.129
Usable IP Addresses	192.0.2.129-192.0.2.134
WAN IP Address	192.0.2.130
IP Alias VIP	192.0.2.131
Bridge Host	192.0.2.134

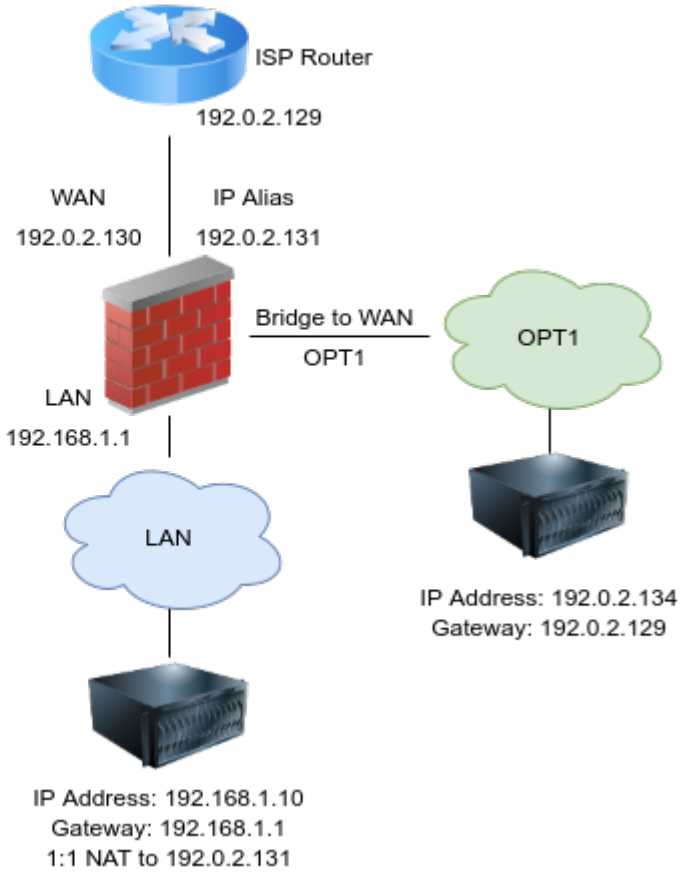


Fig. 19: Multiple Public IP addresses In Use Single IP Subnet

WAN IP Address Assignments	
WAN Subnet	192.0.2.128/30
WAN Gateway	192.0.2.129
WAN IP Address	192.0.2.130
Inside Subnet	192.0.3.0/26
OPT1 Address	192.0.3.1/26
OPT1 Usable Addresses	192.0.3.2-192.0.3.62

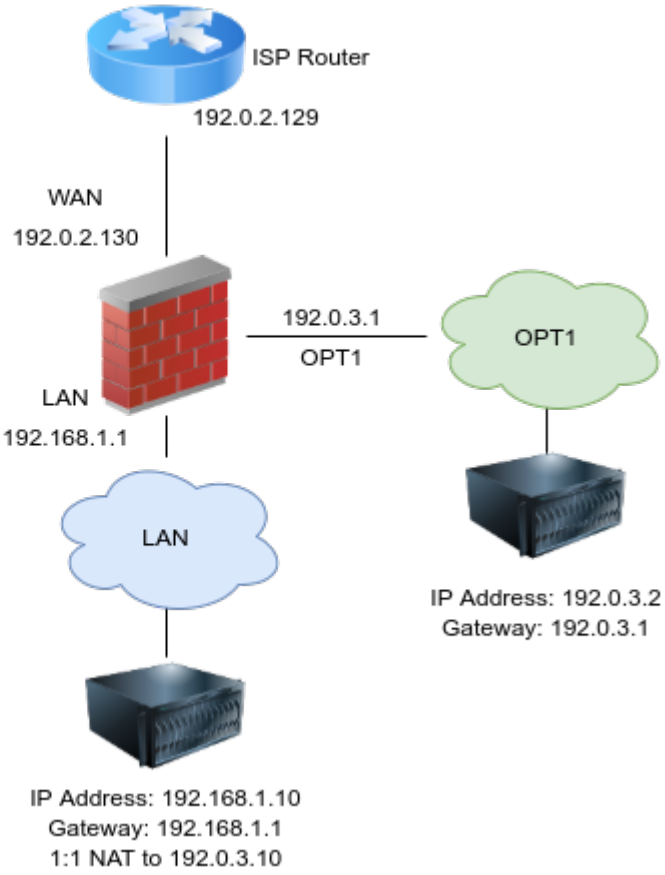


Fig. 20: Multiple Public IP Addresses Using Two IP Subnets

If the firewall is part of a High Availability cluster using CARP, the WAN side subnet will need to be at least a /29 size network, so each firewall has its own WAN IP address plus a CARP VIP. The provider will route the larger inside subnet to the WAN CARP VIP in this type of configuration. The inside IP subnet must be routed to an IP address that is always available regardless of which firewall is up, and the smallest subnet usable with CARP is a /29. Such a setup with CARP is the same as illustrated above, with the OPT1 gateway being a CARP VIP, and the provider routing to a CARP VIP rather than the WAN IP address. CARP is covered in *High Availability*.

Multiple IP subnets

In other cases, a site may be allocated multiple IP subnets from the ISP. Usually when this happens, the site started with one of the two previously described arrangements, and later when requesting additional IP addresses the site was provided with an additional IP subnet. Ideally, this additional subnet will be routed to the firewall by the ISP, either to its WAN IP address in the case of a single firewall, or to a CARP VIP when using HA. If the provider refuses to route the IP subnet to the firewall, but rather routes it to their router and uses one of the IP addresses from the subnet as a gateway IP address, the firewall will need to use Proxy ARP VIPs, IP Alias VIPs, or a combination of IP Alias and CARP VIPs for the additional subnet. If at all possible, the provider should route the IP subnet to the firewall as it makes it easier to work with regardless of the firewall being used. It also eliminates the need to burn 3 IP addresses in the additional subnet, one for the network and broadcast addresses and one for the gateway IP address. With a routed subnet, the entire subnet is usable in combination with NAT.

Where the IP subnet is routed to the firewall, the scenario described in *Small WAN IP Subnet with Larger LAN IP Subnet* applies for an additional internal subnet. The subnet can be assigned to a new OPT interface, used it with NAT, or a combination of the two.

Additional IP Addresses via DHCP

Some ISPs require additional IP addresses to be obtained via DHCP. This is not a good means of obtaining multiple public IP addresses, and must be avoided in any serious network. A business-class connection should not require this. pfSense software is one of the few firewalls which can be used in any capacity with additional IP addresses from DHCP. This offers limited flexibility in what the firewall can do with these addresses, leaving only two feasible options.

Bridging

If the additional IP addresses from DHCP must be directly assigned to the systems that will use them, bridging is the only option. Use an OPT interface bridged with WAN for these systems, and the systems must be configured to obtain their addresses using DHCP.

Pseudo multi-WAN

The only option for having the firewall pull these DHCP addresses as leases is a pseudo multi-WAN deployment. Install one network interface per public IP address, and configure each for DHCP. Plug all the interfaces into a switch between the firewall and the modem or router. Since the firewall will have multiple interfaces sharing a single broadcast domain, enable **Suppress ARP messages** on **System > Advanced**, **Networking** tab to eliminate ARP warnings in the system log, which are normal in this type of deployment.

The only use of multiple public IP addresses assigned in this fashion is for port forwarding. Port forwards can be used on each WAN interface that uses an IP address assigned to that interface by the ISP DHCP server. Outbound NAT to the OPT WANs will not work because of the limitation that each WAN must have a unique gateway IP address to properly direct traffic out of that WAN. This is discussed further in *Multiple WAN Connections*.

14.1.13 Virtual IP Addresses

Some types of interfaces on pfSense® software can utilize more than one IP address at a time. The primary IP address for an interface comes from the interface settings, while Virtual IP (VIP) addresses facilitate the use of additional IP addresses in conjunction with NAT or local services.

VIP Types

There are four types of Virtual IP addresses available in pfSense: *IP Alias*, *CARP*, *Proxy ARP*, and *Other*. Each is useful in different situations. In most circumstances, pfSense software will need to answer ARP request for a VIP which means that IP Alias, Proxy ARP or CARP must be used. In situations where ARP is not required, such as when additional public IP addresses are routed by a service provider to the WAN IP address on the firewall, Other type VIPs can make it easier to use those addresses in NAT rules.

pfSense software will not respond to ICMP echo requests (pings) destined to Proxy ARP and Other type VIPs regardless of firewall rule configuration. With Proxy ARP and Other VIPs, NAT must be present on the firewall, forwarding traffic to an internal host for ping to function. See *Network Address Translation* for more information.

IP Alias

IP Aliases work like any other IP address on an interface, such as the actual interface IP address. They will respond to layer 2 (ARP) and can be used as binding addresses by services on the firewall. They can also be used to handle multiple subnets on the same interface. pfSense software will respond to ping on an IP Alias, and services on the firewall that bind to all interfaces will also respond on IP Alias VIPs unless the VIP is used to forward those ports in to another device (e.g. 1:1 NAT).

IP Alias VIPs can use *Localhost* as their interface to bind services using IP addresses from a block of routed addresses without specifically assigning the IP addresses to an interface. This is primarily useful in HA with CARP scenarios so that IP addresses do not need to be consumed by a CARP setup (one IP each per node, then the rest as CARP VIPs) when the subnet exists only inside the firewall (e.g. NAT or firewall services such as VPNs).

IP Aliases on their own do not synchronize to XMLRPC Configuration Synchronization peers because that would result in an IP address conflict. One exception to this is IP Alias VIPs using a CARP VIP "interface" for their interface. Those do not result in a conflict so they will synchronize. Another exception is IP Alias VIPs bound to Localhost as their interface. Because these are not active outside the firewall itself, there is no chance of a conflict so they will also synchronize.

CARP

CARP VIPs are primarily used with High Availability redundant deployments utilizing CARP (*CARP Overview*). CARP VIPs each have their own unique MAC address derived from their VHID, which can be useful even outside a High Availability deployment.

See also:

For information on using CARP VIPs, see *High Availability*.

CARP VIPs may also be used with a single firewall. This is typically done in cases where the pfSense deployment will eventually be converted into an HA cluster node, or when having a unique MAC address is a requirement. In rare cases a provider requires each unique IP address on a WAN segment to have a distinct MAC address, which CARP VIPs provide.

CARP VIPs and IP Alias VIPs can be combined in two ways:

- To reduce the amount of CARP heartbeats by stacking IP Alias VIPs on CARP VIPs. See *Using IP Aliases to Reduce Heartbeat Traffic*.
- To use CARP VIPs in multiple subnets on a single interface. See *High Availability*.

Proxy ARP

Proxy ARP VIPs function strictly at layer 2, providing ARP replies for the specified IP address or CIDR range of IP addresses. This allows pfSense software to accept traffic targeted at those addresses inside a shared subnet. For example, pfSense software can forward traffic sent to an additional address inside its WAN subnet according to its NAT configuration. The address or range of addresses are not assigned to any interface on pfSense, because they don't need to be. This means no services on pfSense software itself can respond on these IP addresses.

Proxy ARP VIPs do not sync to XML-RPC Configuration Sync peers because doing so would cause an IP address conflict.

Other

Other type VIPs define additional IP addresses for use when ARP replies for the IP address are not required. The only function of adding an *Other* type VIP is making that address available in the NAT configuration drop-down selectors. This is convenient when the firewall has a public IP block routed to its WAN IP address, IP Alias, or a CARP VIP.

VIP Configuration Options

The options available for each VIP vary based on the selected type. This list contains all available options for all types, with the restrictions noted as needed.

Type
: Sets the type of VIP this will be, and changes the available options on the page. See *VIP Types* for a description of each type.

Interface
: The parent interface upon which this VIP will reside.

  For IP Alias type VIPs, this can also be a CARP VIP (*Using IP Aliases to Reduce Heartbeat Traffic*).

Address Type
: For Proxy ARP and Other type VIPs this option declares whether the VIP defines a single IP address or an entire subnet.

Addresses
: The IP address and subnet mask for this VIP.

Expansion
: For IPv4 Proxy ARP and Other type VIPs, this specifies whether a subnet address type VIP will be expanded into its individual IP address components in drop-down menus where VIPs can be selected, such as on NAT rules.

> **Warning:** For large subnets this can cause rendering issues in browsers as they may not properly handle large drop-down lists.

CARP Options
: These options are only available for CARP type VIPs:

Virtual IP Password
The secret key with which to protect the CARP heartbeat traffic. If peers do not agree on the key, they cannot coordinate their actions and will all appear as MASTER status.

VHID Group
The group identifier used to indicate between multiple hosts that they are coordinating control of the same address. VHIDs must be unique per layer 2.

For IPv4, this is commonly set to match the last octet of the IPv4 VIP address.

The VHID also influences the resulting MAC address of the VIP, where the last two characters of the MAC are the VHID in hexadecimal.

Advertising Frequency
Controls how often the master node sends out CARP heartbeat advertisements. Hosts transmitting faster than their peers assume control of the VIP.

Base
The base value of whole seconds between advertisements. The default of `1` is typically ideal, but can be raised in special cases.

Skew
1/256th fractions of a second added to the base. Typically, the preferred primary node will have a skew of `1` or `0`, and secondary nodes will be around `100` or higher.

When XMLRPC is configured to synchronize VIPs, this value is automatically adjusted for the secondary node by adding `100` to this value.

CARP Mode (Plus Only)
pfSense Plus software has an option to choose how CARP operates, either in multicast or unicast mode.

Multicast
CARP heartbeats are advertised to the entire Layer 2 network using multicast, and multiple peers can monitor for heartbeats and participate. This requires that all peers be connected to multicast-capable networks (e.g. Ethernet switches or equivalent). It also requires that the underlying network does not manipulate or limit multicast.

Unicast
CARP heartbeats are sent directly to the specified **Peer Address** and do not use multicast. This allows CARP to operate on networks which cannot carry multicast traffic, such as cloud provider networks and VPNs. However, this limits a cluster to two nodes since both nodes must send heartbeats to each other directly.

> **Warning:** The best practice is to avoid using unicast mode on traditional switched L2 infrastructure when possible. In unicast mode, CARP heartbeats are sourced from the interface MAC address and not the CARP MAC. Since switches do not see packets sourced from the CARP MAC, they may flood packets for unicast CARP VIPs to all ports. This behavior has significant security and performance concerns which are eliminated by using multicast mode instead.

Peer Address
The IP address of the peer node to which the firewall will send its unicast CARP heartbeats for this VIP.

> **Warning:** At this time the **Peer Address** cannot be an IPv6 link local address. This will be corrected in a future release. See Redmine #14385 for details.

Note: When VIPs are synchronized via XMLRPC, the primary node will adjust the VIP such that its own address on the parent interface is configured on the peer VIP. In scenarios where this is not correct, VIP synchronization may need to be managed manually.

See also:

Switch/Layer 2 Concerns

Description
Optional text to describe the VIP and/or its purpose, e.g. "VIP for PBX" or "NAT Subnet VIP".

Feature Comparison

Virtual IP Address Feature Comparison

This document summarizes and compares capabilities of the different Virtual IP Address types.

See *Virtual IP Addresses* for detailed information about each type of VIP.

VIP Features Table

Table 2: Virtual IP Address Feature Comparison

VIP Type	NAT	Binding	ARP/L2	Clustering	Subnet Mask	ICMP	Single/Range
IP Alias	Yes	Yes	Yes	See Notes	See Notes	Yes	Single
CARP	Yes	Yes	Yes	Yes	Yes	Yes	Single
Proxy ARP	Yes	No	Yes	No	n/a	No (1)	Either
Other	Yes	No	No	Yes (2)	n/a	No (1)	Either

Notes:

1. The ICMP column represents responses from the firewall itself without NAT. With 1:1 NAT or port forwards, any VIP will pass ICMP through to the target device.
2. "Other" type VIPs are for routed subnets, and CARP is irrelevant, so they are compatible with HA (See below)

Virtual IP Feature Summary

It is difficult to express all details of VIP capabilities in a table format, so this section contains a more thorough overview of the various types and what they can/cannot do a bullet point format.

IP Alias

- **Can** be used for NAT.
- **Can** be used by the firewall itself to bind/run services.
- Adds extra IP addresses to an interface.
- Generates ARP (Layer 2) responses for the VIP address.
- Can be in a **different subnet** than the **real** interface IP address when used directly on an interface.
- **Will** respond to ICMP ping if allowed by firewall rules.
- Must be added individually
- Subnet mask should match the interface IP, or `/32`. Matching the interface subnet is best. For IP addresses in different subnets at least one IP alias VIP **must** have the correct mask for the new subnet.
- Can be stacked on top of a CARP VIP to bypass VHID limits and lower the amount of CARP heartbeat traffic.
  - Stacked IP Alias VIPs will synchronize via XMLRPC.
  - Stacked IP Alias VIPs must be inside the same subnet as the CARP VIP upon which they are placed.
- Can be added to localhost for binding services in routed subnets. IP Alias VIPs bound to localhost will synchronize via XMLRPC

CARP

- **Can** be used for NAT.
- **Can** be used by the firewall itself to bind/run services.
- Generates ARP (Layer 2) traffic for the VIP.
- Can be used for clustering (master firewall and standby failover firewall.)
- CARP VIPs may be in other subnets.
- **Will** respond to ICMP ping if allowed by firewall rules.
- Must be added individually.
- Subnet mask **must** match the interface IP address.
- Generates its own MAC address for the VIP. This MAC is different from its physical parent interface.

Proxy ARP

- **Can** be used for NAT.
- **Cannot** be used by the firewall itself to bind/run services.
- Generates ARP (Layer 2) traffic for the VIP.
- Can be in a **different subnet** than the real interface IP.
- **Will not** respond to ICMP ping.
- Can be added individually or as a subnet to make a group of VIPs.

Other

- **Can** be used for NAT.
- **Cannot** be used by the firewall itself to bind/run services.
- Can be used if the address is routed to the firewall without needing ARP/Layer 2 messages. (e.g. Upstream provider routes a subnet to the WAN IP address)
- Can be in a **different subnet** than the real interface IP address.
- **Will not** respond to ICMP echo requests.
- Can be added individually or as a subnet to make a group of VIPs.
- Can be used with CARP, e.g. subnet routed to external CARP VIP.

14.1.14 Using EasyRule to Manage Firewall Rules

The EasyRule function found in the GUI and on the command line can add firewall rules quickly.

EasyRule in the GUI

In the pfSense® software GUI, this function is available in the Firewall Log view (**Status > System Logs**, **Firewall** tab).

The  icon next to the source IP address adds a **block** rule for that IP address on the interface. To be more precise, it creates or adds to an alias containing IP addresses added from Easy Rule and blocks them on the selected interface.

The  icon next to the destination IP address works similar to the block action, but it adds a more precise **pass** rule. This **pass** rule allows traffic on the interface, but it must match the same protocol, source IP address, destination IP address, and destination port.

EasyRule in the Shell

The shell version of Easy Rule, `easyrule`, can manage EasyRule firewall rules and entries from a shell prompt. When the `easyrule` command is run without parameters, it prints a usage message to explain its syntax.

The way `easyrule` adds a block rule using an alias, or a precise pass rule specifying the protocol, source, and destination, work the same as the GUI version.

The general form of the command is:

```
# easyrule <action> <interface> <parameters>
```

action
: The action can be one of `pass`, `block`, `showblock`, or `unblock`. Each one takes different parameters and is explained later in this section.

interface
: The descriptive name of the interface, as seen in the GUI on the interface configuration page. For example: `WAN`, `LAN`, `DMZ`, `OFFICEVPN`. When using the descriptive names, it is not case-sensitive.

  The interface value can also be the internal designation for the interface, such as `wan` or `opt2`.

Special names for certain groups are also available here: `openvpn` for OpenVPN tab rules, `ipsec` for IPsec tab rules, `pppoe` for PPPoE server tab rules, and `l2tp` for L2TP server tab rules.

Pass

Passing requires several details so it does not create an overly permissive rule. The destination port is optional if the protocol does not require a port (e.g. ICMP, OSPF, etc.).

```
# easyrule pass <interface> <protocol> <source address> <destination address>␣
↪[destination port]
```

protocol
The name of the protocol to pass, or `any` to pass any protocol.

source address
The source of traffic for the pass rule.

Can be an IPv4/IPv6 address, subnet, alias name, or special network name such as `any`, `pppoe` or `l2tp`.

destination address
The destination of traffic for the pass rule.

Can be an IPv4/IPv6 address, subnet, alias name, or special network name such as `any`, `pppoe` or `l2tp`.

destination port
The destination port number if the protocol requires ports (`TCP`, `UDP`).

To pass traffic to any port, use `any`.

Note: The address family of the source and destination must match.

Example pass rule for a protocol that uses ports:

```
# easyrule pass wan tcp 1.2.3.4 192.168.0.4 80
```

Example pass rule for a protocol without ports:

```
# easyrule pass wan icmp 1.2.3.4 192.168.0.4
```

Block

Blocking only requires a source IP address to block:

```
# easyrule block <interface> <source address>
```

source address
The source of traffic to block.

Can be an IPv4/IPv6 address, subnet, alias name, or special network name such as `any`, `pppoe` or `l2tp`.

Block example:

```
# easyrule block wan 1.2.3.4
```

Show a Block

This program can also display the contents of addresses currently blocked by easyrule on an interface.

```
# easyrule showblock <interface>
```

```
# easyrule showblock wan
1.2.3.4/32
5.6.7.8/32
9.10.11.0/24
```

Remove a Block

```
# easyrule unblock <interface> <source address>
```

source address
The source of traffic to unblock. The address must already be blocked *by EasyRule*.

Note: This action **will not** remove block rules or entries that were not created by EasyRule.

Can be an IPv4/IPv6 address, subnet, alias name, or special network name such as `any`, `pppoe` or `l2tp`.

```
# easyrule showblock wan
1.2.3.4/32
5.6.7.8/32
9.10.11.0/24
# easyrule unblock wan 5.6.7.8
Host unblocked successfully
# easyrule showblock wan
1.2.3.4/32
9.10.11.0/24
```

14.1.15 Firewall Packet Flow Data

Starting with pfSense Plus software version 24.03 the firewall can directly export NetFlow v5 and IPFIX traffic flow data to one or more collectors using the `pflow(4)` feature in PF. The data is collected directly from firewall states and does not require a separate daemon, service, or add-on package.

Note: pfSense® Plus software version 24.03 or later is required to use the Packet Flow Data feature. This feature is not available on pfSense CE Software.

As this is a function of the firewall, this feature is located at **Firewall > Packet Flow Data**.

Flows can be tracked by default or only for specific rules, so the user has the flexibility to control the scope of exported data.

The feature is capable of up to 16 different export configurations which send data to multiple collectors and/or using different data formats.

Note: As this feature relies upon data from firewall states to function, it requires the firewall rules to keep state to generate that data. This is the default behavior, so it is unlikely to be a blocking issue for most users.

Warning: Packet Flow Data changes only happen for **new** states created after the feature has been activated or deactivated. When activating the feature, states that already exist will not start sending data, only new connections will. Similarly, when deactivating the feature, old states may continue to trigger export data.

Rebooting the device or resetting the states after changing the configuration will ensure that all states are handled appropriately.

Global Packet Flow Options

There are two options available on the main Packet Flow Data configuration page at **Firewall > Packet Flow Data**:

Enables Packet Flow Data Tracking and Exporting
: This option loads the required kernel module for `pflow(4)` and configures the packet flow data exporters in the OS.

  This option must be checked for the firewall to export flow data.

Enable Packet Flow tracking on all rules by default
: When checked, all firewall rules will track packet flow data for new states by default. This can be overridden on a per-rule basis if necessary.

  When unchecked, nothing is tracked by default and flow tracking must be enabled on each firewall rule manually.

  See also:

  For information on managing Packet Flow Data in firewall rules, see *Packet Flow Data*.

Packet Flow Exporters

Under the global options section the page contains a list of current exporter configurations.

The list displays the configuration properties of each current entry. Also on each line are several indicators and action icons:

- At the start of a line is a checkmark. If it's dark, the entry is enabled. If It's light, the entry is disabled. Clicking the checkmark icon will toggle the entry between enabled and disabled states.

  The text color for the entire line changes to a lighter shade for disabled entries as well.

-  - Edits the current entry
-  - Duplicates the current entry
-  - Deletes the current entry after a confirmation prompt.

Under the list is an  **Add** button to create new entries. This button is hidden when the list is full.

Exporter Options

Exporter entries define remote hosts capable of receiving and processing packet flow data (e.g. NetFlow or IPFIX). Up to 16 distinct exporters can be defined.

When creating or editing an exporter entry, the following options are available:

Description
Some text describing this entry (e.g. Collector host name, purpose, etc.).

Enable
A checkbox which controls whether the exporter is active.

When checked, the firewall will export packet flow data to this host. When unchecked, this host is ignored.

Source IP Address
A drop-down menu containing local interfaces, VIPs, and other valid traffic sources. If the collector expects flow data packets to come from a specific source address, select it in this list.

If the destination is over a VPN, this likely should either be set to the corresponding VPN interface or another local interface.

Source Port
Source port for packet flow export. Leave blank to use a random port (default).

> **Warning:** This port must not already be in use anywhere on the firewall, including other exporters.
>
> If this option is set, the **Source IP Address** must also be set to a specific value.

Destination IP Address
IP address (IPv4 or IPv6) of the remote flow collector host, e.g.: `192.168.100.100` or `fd00:abcd::1`.

This outer transport can be either IPv4 or IPv6 no matter which **Flow Protocol** is active, but the address family must match the **Source IP Address**. For example, if the **Destination IP address** is set to an IPv6 address, the selected **Source IP Address** entry must also have an IPv6 address.

Destination Port
Destination port on the collector where it is listening for packet flow data.

Leave blank to use the default port, `2055`.

Flow Protocol
Format for packet flow data.

NetFlow v5
NetFlow v5 is more widely compatible with collectors but is more limited in the type of traffic it can include in flow data.

Note: The NetFlow v5 specification does not include IPv6. To track IPv6 flows, use IPFIX instead.

IPFIX
IPFIX is capable of carrying IPv6 flow data, but is not compatible with all collectors.

This protocol also includes RFC 8158 NAT44 flow information which can be utilized for centralized logging of NAT translation data.

Consult the documentation for the collector to determine the best available format.

Observation Domain
Observation Domain for flows to this exporter. This is an unsigned non-zero 32-bit integer (`1-4294967295`). Not all collectors use this value, but a collector can use it to differentiate flows from devices in similar roles or locations.

Leave blank to use the default of `1`.

Packet Flow Data Example Configuration

This is a brief example of configuring the Packet Flow Data feature to export all flow data to a collector at on LAN at an address of `10.1.2.3` with otherwise default settings.

- Navigate to **Firewall > Packet Flow Data**
- Check **Enable Packet Flow Data Tracking and Exporting**
- Check **Enable Packet Flow tracking on all rules by default**
- Click **Save**
- Click  **Add** to create a new exporter
- Configure the settings as described in *Exporter Options*:

  Description
  Local NetFlow Collector

  Enable
  Checked

  Source IP Address
  LAN

  Source Port
  Empty

  Destination IP Address
  `10.1.2.3`

  Destination Port
  Empty

  Flow Protocol
  IPFIX

  Observation Domain
  Empty

- Click **Save**
- Click **Apply Changes**

At this point, new states will begin exporting flow data.

See also:

- *Ordering of NAT and Firewall Processing*
- *Viewing the Firewall Log*
- *Filter Reload Status*

14.2 Aliases

Aliases are collections of addresses that allow many hosts to be acted upon by a few firewall rules. They can greatly simplify a ruleset and make it easier to understand and manage.

14.2.1 Aliases

Aliases define groups of ports, hosts, or networks. They can be referenced by firewall rules, port forwards, outbound NAT rules, and several other areas. Using aliases results in configurations and rulesets which are significantly shorter, self-documenting, and easier to manage.

Note: These aliases are collections of items for use by the firewall. Despite similar names, this is a completely different concept than "IP alias" type *Virtual IP Addresses*, which are a means of adding additional IP addresses to a network interface.

Aliases are located at **Firewall > Aliases** in the GUI. The page is divided into separate tabs for each general *type of alias*:

IP
: Lists *Host Aliases* and *Network Aliases*.

Ports
: Lists *Port Aliases*.

URLs
: Lists *URL Aliases* and *URL Table Aliases*.

All
: Shows all types of aliases, including *Built-In System Aliases*, in one large list.

When creating an alias, the GUI will sort the entry to the correct location based on the type of alias content, no matter which tab was used to create the alias.

Alias Features and Limitations

Administrators have a lot of flexibility when defining aliases. These abilities are a large part of how aliases can make firewall rulesets easier to manage.

Alias Sizing Concerns

When the firewall loads alias data it copies the contents into internal tables which it uses to quickly perform address matches.

The total size of all aliases/tables must fit in roughly **half** the amount of **Firewall Maximum Table Entries**, which defaults to `400000`.

See *Firewall Maximum Table Entries* for more information on this behavior.

Nested Aliases

Most aliases can be nested inside other aliases of similar types to collect entries into larger groups. For example, one `Servers` alias can nest an alias containing web servers, an alias containing mail servers, and an alias containing database servers all together.

To nest, aliases must be either the same or compatible types. For example, *Network Aliases* cannot nest *Port Aliases* since they are not the same type of alias. However, *Host Aliases* and *Network Aliases* can nest each other since they are compatible. *URL Table Aliases* can nest other URL table aliases, and *URL Aliases* can nest other URL aliases.

Hostnames in Aliases

Host and network type aliases can contain entries consisting of fully qualified domain name (FQDN) style hostnames (e.g. `host.example.com`) in regular or IDN format.

Tip: This feature is also useful for tracking dynamic DNS entries to allow users to access services from dynamic IP addresses.

For these entries to function, the firewall **must** be able to resolve hostnames using A or AAAA type DNS queries. This means that the firewall must have working DNS, the FQDN must exist, and the firewall must be able to resolve these hostnames using the DNS servers present in its configuration.

Warning: This feature only performs **forward** name resolution of FQDNs using A and AAAA records such as `host.example.com`.

Aliases **do not** perform pattern matches, wildcard matches (e.g. `*.example.com`), SRV record lookups, or any other style of record comparison.

If a DNS query for a hostname returns multiple IP addresses, the firewall adds **all** IP addresses in the result **at the time of the query** to the alias.

Warning: This feature is *not* useful for controlling access to hostnames for large public websites such as those served by content delivery network (CDN) providers. Such sites tend to have round-robin, localized, or randomized responses to DNS queries so the contents of the alias on the firewall do not necessarily match the response a client receives when it resolves the same hostname. This feature can work for smaller sites which have single addresses or sites which always return complete sets of addresses in their DNS responses.

The firewall periodically resolves and updates hostname entries in host or network type aliases. The default interval is `300` seconds (5 minutes). This behavior can be changed by adjusting the *Aliases Hostnames Resolve Interval*.

Mixing IPv4 and IPv6 Addresses in Aliases

IPv4 and IPv6 addresses can coexist inside the same alias. The firewall uses the appropriate type of addresses from the alias content as needed when a rule references the alias. This allows a single alias to function in IPv4 rules, IPv6 rules, and even IPv4+IPv6 rules.

Uses Beyond Firewall and NAT Rules

pfSense® software allows the use of aliases in several places outside of firewall and NAT rules. For example, they can be used with certain fields in OpenVPN and in static routes. GUI pages *will indicate* if and when a field can utilize aliases.

Alias Types

The next sections describe the behavior of each alias type in detail. Any type-specific *configuration settings* are also covered in these sections.

Built-In System Aliases

pfSense® software includes several built-in **System Aliases** which are accessible to users. These include stock collections of networks and IP addresses as well as automatically generated and/or maintained collections.

These are visible on the **All** tab under **Firewall > Aliases** in the **System Aliases** section.

Note: Some of these System Aliases have multiple variations for IPv4, IPv6, and combined IPv4+IPv6. These are denoted by the number `4`, `6` or `46` near the end of the name, respectively. In most cases it is safe to use the IPv4+IPv6 version as the *firewall will use whichever address family is appropriate for a rule*, but some users may prefer to configure different rules for IPv4 and IPv6.

The current list includes the following entries:

bogons
: *Bogon networks*.

sshguard
: Hosts blocked by anti-brute-force login protection for SSH and the GUI.

snort2c
: Hosts blocked by IDS/IPS software.

virusprot
: Hosts which tripped firewall rate limit protections in advanced options (e.g. maximum source hosts for a rule).

_nexus_vpn_port_
: Netgate Nexus VPN port (*General Options*).

vpn_networks
: Networks for IPsec, OpenVPN, and PPPoE servers.

  Note: This may not be complete as it cannot detect networks made available via dynamic routing, pushed routes from OpenVPN, advanced/custom options, etc.

  ---

negate_networks
Networks to exclude from policy routing rules with any destination.

tonatsubnets
Networks for which the firewall will apply outbound NAT in Automatic and Hybrid modes.

loopback(4|6|46)
Local loopback addresses.

linklocal(4|6|46)
Link-local addresses which must not leave their own network.

private(4|6|46)
"Private" network allocations common for local private networks.

multicast(4|6|46)
Multicast networks.

reserved(4|6|46)
All current known reserved networks.

Host Aliases

Host type aliases contain entries consisting of individual IP addresses or fully qualified domain names (FQDNs).

When creating an alias, users may enter an IP address range such as `192.168.1.1-192.168.1.10` or a small subnet such as `192.168.1.16/28`. When saving the alias, the firewall translates that specification into a list of individual IP addresses.

Figure *Example Hosts Alias* shows an example of a host type alias which contains a list of public web servers.

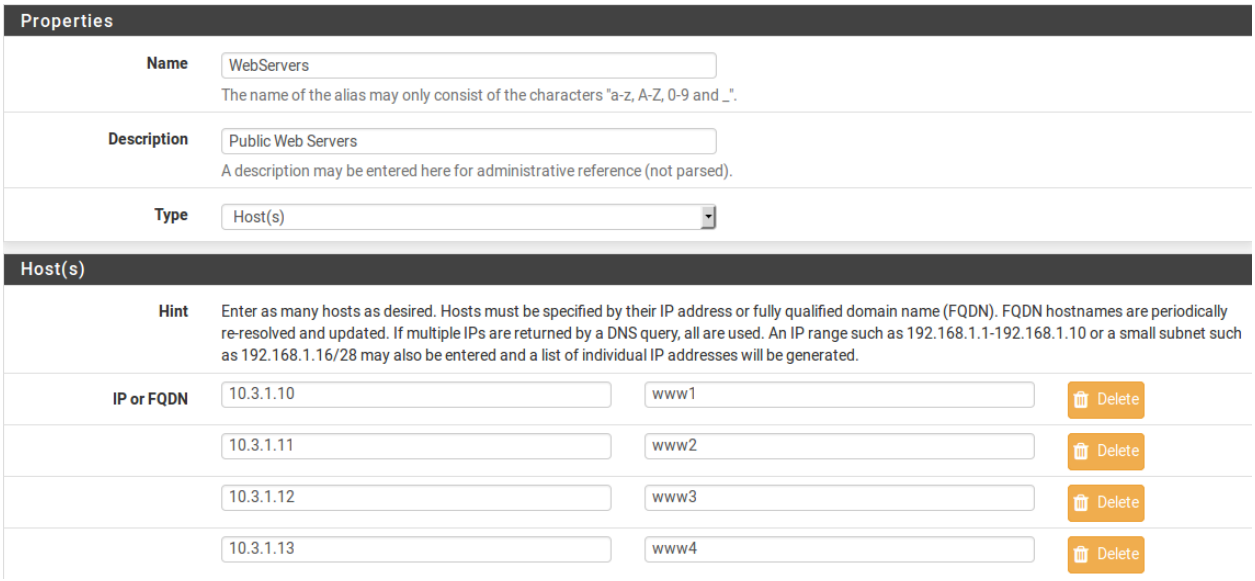


Fig. 21: Example Hosts Alias

Network Aliases

Network type aliases contain entries consisting of CIDR format networks/prefixes or fully qualified domain names (FQDN) for single addresses.

For subnets, select the CIDR mask that pertains to each entry. `/32` specifies a single IPv4 host, `/128` specifies a single IPv6 host, `/24` specifies `255.255.255.0`, `/64` specifies a normal IPv6 network, etc.

Hostnames (FQDNs) may also be specified, using a `/32` mask for IPv4 or `/128` for IPv6. This mask **is not** applied to addresses returned by DNS.

Figure *Example Network Alias* shows an example of a network alias.

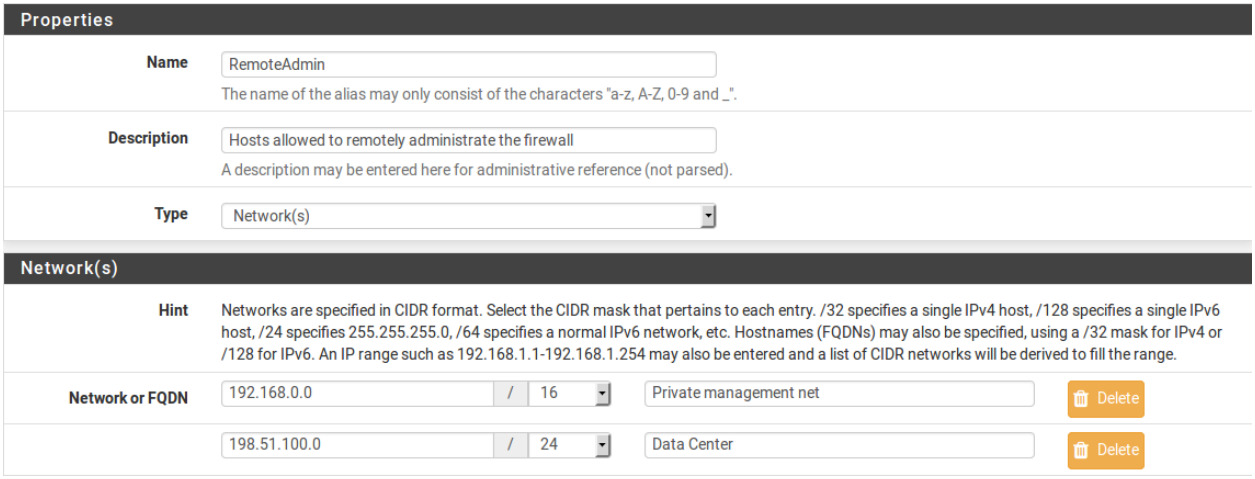


Fig. 22: Example Network Alias

When an alias entry contains an IPv4 range the firewall automatically translates the range to an equivalent set of IPv4 CIDR networks which exactly contain the range. This is shown in Figure *Example IP Range After*.



Fig. 23: Example IP Range Before

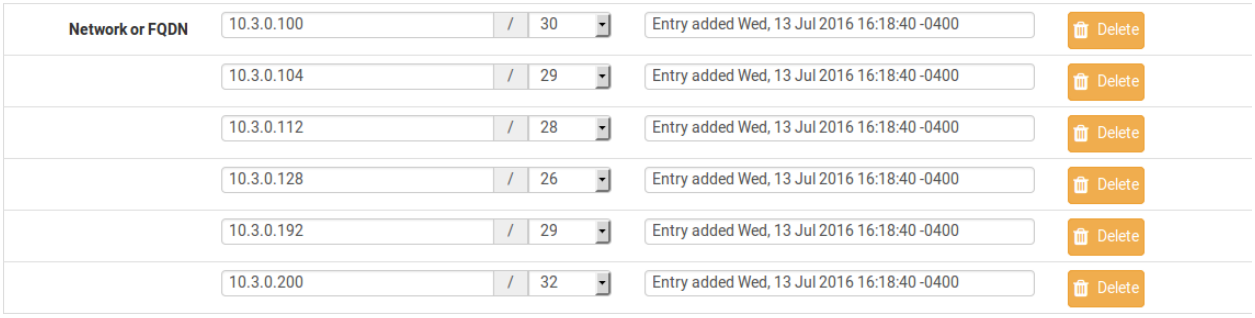


Fig. 24: Example IP Range After

Port Aliases

Port type aliases contain entries consisting of ports numbers and port ranges. A single port is an integer from `1-65535`. A port range is two ports separated by a colon (`:`), for example, `1194:1199`. Port ranges match the specified start and ending port numbers and all ports in between.

Port aliases **do not** have a direct relationship with any protocol. Firewall rules using a protocol of TCP, UDP, or SCTP may use port aliases to match port numbers for the protocol on the rule.

Figure *Example Ports Alias* shows an example of a port type alias.

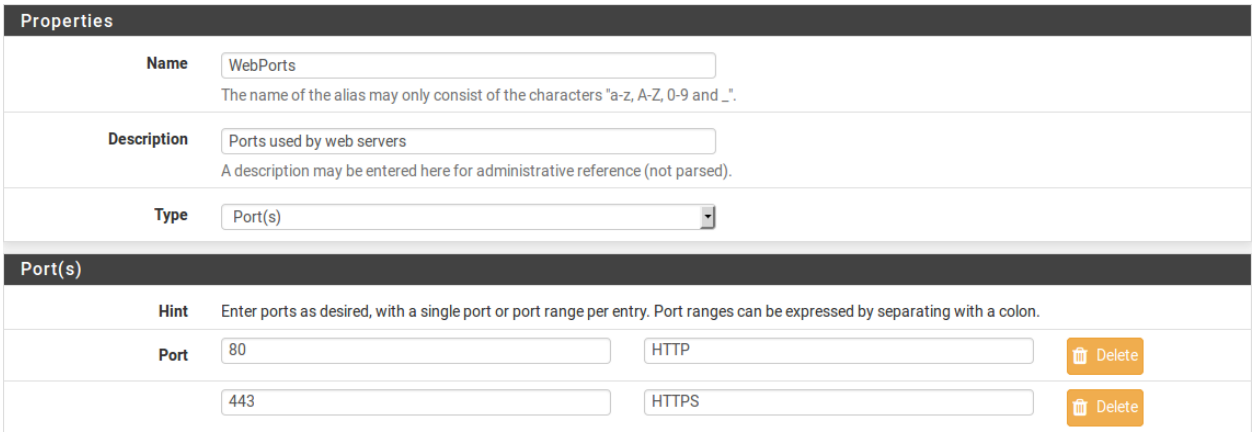


Fig. 25: Example Ports Alias

URL Aliases

URL type aliases contain one or more URLs which return a plain text a list of entries.

When creating the alias, the firewall reads up to **3,000** entries from each URL and imports them into an alias.

The exact behavior depends on the type of URL alias:

URL (IPs)
: The URLs must contain IP addresses, CIDR masked network entries, or FQDNs. The firewall creates a Network type alias from the content.

URL (Ports)
: The URL must contain only port numbers or port ranges. The firewall creates a Port type alias from the content.

URL Table Aliases

URL Table type aliases contain one or more URLs which return a plain text a list of entries, plus an update interval which indicates how frequently the firewall should re-fetch the URL contents.

URL Table type aliases download the content of the URLs into a special location on the firewall, then use that content for a `persist` table, also known as a file-based alias. The full contents of these alias are not directly editable in the GUI, but the GUI can display the *Firewall Table Contents*.

The drop-down list after the `/` in an entry of a URL Table alias controls the number of days after which the firewall re-fetches content from the stored URLs. The firewall checks once per day to determine if URL table aliases need updates.

URL Table aliases can be quite large, containing many thousands of entries. Some customers use them to hold lists of all IP address blocks in a given country or region, which can easily surpass 40,000 entries. The pfBlockerNG package uses this type of alias when handling country lists and other similar actions.

The exact behavior depends on the type of URL alias:

URL Table (IPs)
The URLs must contain IP addresses, CIDR masked network entries, or FQDNs. The firewall creates a Network type alias from the content.

URL Table (Ports)
The URL must contain only port numbers or port ranges. The firewall creates a Port type alias from the content.

Alias Configuration

Alias Settings

When editing an Alias entry, the following settings are available:

Name
A **Name** for the alias. The name may only consist of the characters `a-z`, `A-Z`, `0-9` and `_`.

Note: The name of an alias cannot conflict with reserved names for items such as interface names, gateway names, or internal PF keywords. Input validation will reject conflicting names.

Description
A **Description** for the alias.

Type
The **Type** for the alias, which alters the behavior of the alias and tells the firewall which types of entries can be added to the alias.

The following types are available:

Host
Host Aliases contain single IP addresses or FQDN hostnames.

Network
Network Aliases contain CIDR-masked lists of networks, FQDN hostnames, IP address ranges, or single IP addresses.

Port
Port Aliases contain lists of port numbers or ranges of ports for TCP or UDP.

URL (IP or Port)
URL Aliases contain items the firewall fetches from the specified URL(s) at the time the alias is created. Once created, the alias becomes a typical network or port type alias.

URL Table (IP or Port)
URL Table Aliases contain items the firewall fetches from the specified URL(s), but it periodically updates the content.

Entries
The lower section of the alias page contains the entries for the alias. The behavior of this section varies based on the selected *alias type*.

Creating an Alias

To create a new alias:

- Navigate to **Firewall > Aliases**
- Click  **Add**
- Enter settings as described in *Alias Settings*
- Enter the type-specific information for each member entry

  All *alias types* have a data field and a description field for each entry.

To add new entry to an alias, click  **Add** at the bottom of the list of entries.

To remove entries from an alias, click  **Delete** at the end of the row to remove.

When the alias is complete, click **Save** to store the alias contents.

Each manually-created alias is limited to **5,000 members**. Aliases which require larger amounts of entries should use *URL Table Aliases* instead.

> **Warning:** Some browsers have trouble displaying or using the alias editing view with more than around 3,000 entries.

Bulk Import Network Aliases

The GUI can add multiple entries for Host, Network, and Port type aliases in bulk using the import feature. This can be useful when importing long existing lists of addresses, such as for block lists or corporate networks.

To use the import feature:

- Navigate to **Firewall > Aliases**
- Navigate to the appropriate tab to match the desired alias type

  Use the **IP** tab for Host or Network aliases, use the **Port** tab for Port aliases.

- Click  **Import**
- Fill in the **Alias Name** and **Description**
- Enter the alias contents into the **Aliases to import** text area

  Each entry must be on a separate line.

  Importing from the **IP** tab allows entries containing IP addresses, CIDR masked networks, IP address ranges, or FQDNs.

  Importing from the **Port** tab allows entries containing port numbers and port ranges.

  The import process takes any text following a valid entry as the description for that entry.

- Click **Save**

The firewall imports the content into a normal alias which can be edited later.

Alias Usage

When a user types a letter into a form input which is capable of using aliases, the GUI displays a list of matching aliases. The user can then select the desired alias from the list or type its name out completely.

Note: Alias autocompletion is not case-sensitive, but it is restricted by type. For example, the GUI will autocomplete a Network or Host type alias for a network field, but not a Port alias. Likewise, the GUI will autocomplete a Port alias in a port field, but not a Network alias.

Figure *Autocompletion of Hosts Alias* shows how the `WebServers` alias, configured as shown in Figure *Example Hosts Alias*, can be used in the **Destination** field when adding or editing a firewall rule.

- Edit the firewall rule
- Set the **Destination** drop-down to *Address or Alias*
- Then type the first letter of the desired alias: Enter `W` and the alias appears as shown.

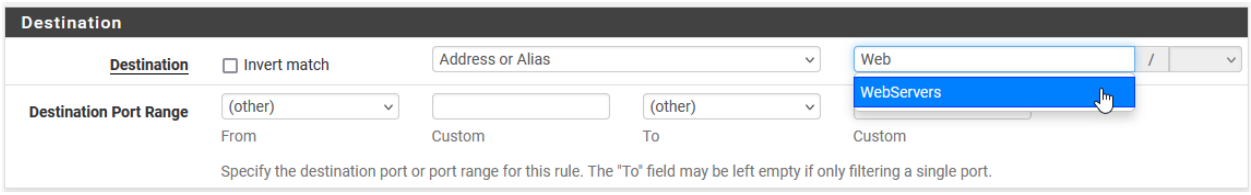


Fig. 26: Autocompletion of Hosts Alias

Figure *Autocompletion of Ports Alias* shows the autocompletion of the ports alias configured as shown in Figure *Example Ports Alias*. If multiple aliases match the letter entered, all matching aliases of the appropriate type are listed. Click on the desired alias to select it.

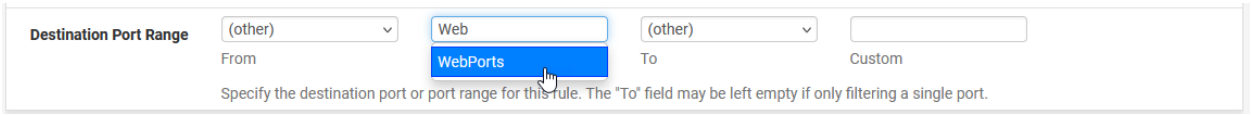


Fig. 27: Autocompletion of Ports Alias

Figure *Example Rule Using Aliases* shows the rule created using the `WebServers` and `WebPorts` aliases. This rule is on WAN, and allows any source to the IP addresses defined in the `WebServers` alias when using the ports defined in the `WebPorts` alias.



Fig. 28: Example Rule Using Aliases

Hovering the mouse cursor over an alias on the **Firewall > Rules** page shows a tooltip displaying the contents of the alias with the descriptions included in the alias. Figure *Hovering Shows Hosts Contents* shows this for the `WebServers` alias and Figure *Hovering Shows Ports Contents* for the `WebPorts` alias.



Fig. 29: Hovering Shows Hosts Contents



Fig. 30: Hovering Shows Ports Contents

Firewall Table Contents

The GUI page at **Diagnostics > Tables** displays the contents of tables defined by the firewall and by users.

The firewall stores aliases and other similar lists of addresses in a PF structure called a **table**. These tables can be relatively static, such as the bogons list or aliases, or dynamic for things like login protection lockout, IDS/IPS, or IP addresses exceeding connection limits.

An alias becomes a table once the firewall loads it into the ruleset.

If an alias contains a hostname the firewall populates the contents of the alias from the results of resolving hostnames using DNS. It periodically resolves the hostnames again and updates the table accordingly if the addresses change. Given this dynamic nature, viewing the table contents may be necessary to confirm which IP addresses are in a table at a given time.

Viewing Tables

To view the contents of a table:

- Navigate to **Diagnostics > Tables**
- Select the desired table from the **Table** drop-down

After making a selection the page will update to display the contents of the selected table.

The  icon at the end of each row in the table content removes individual entries from a table. Removing entries is best used for dynamic tables to remove an entry before it automatically expires. Tables defined manually or by a file will be refreshed when the system performs a filter reload, so it is best to edit an alias and remove an entry rather than removing it from this page.

Default Tables

The firewall includes several tables by default, depending on which features are enabled:

bogons/bogonsv6
: If any interface is configured with *Block Bogon Networks*, these tables will be present.

  The page offers an  **Update** button for the bogon tables which will immediately re-fetch the bogons data rather than waiting for the usual monthly update.

cpzoneid*
: Tables starting with this string are used internally by Captive Portal and are not meant to be managed manually.

negate_networks
: Networks for which the firewall has made policy route negation rules.

snort2c
: A dynamic table containing blocked offenders from IDS/IPS packages, Snort and Suricata.

sshguard
: A dynamic table containing clients that repeatedly failed login attempts for the GUI and SSH.

tonatsubnets
: When using automatic outbound NAT, this table contains the list of networks for which the firewall performs automatic outbound NAT.

  Inspecting the table can aid in diagnosing tricky NAT issues to confirm if a subnet will have automatic outbound NAT applied to its traffic.

virusprot
: A dynamic table containing addresses that have exceeded defined limits on firewall rules.

vpn_networks
: A list of remote networks reachable across VPNs.

14.3 Firewall Guides

How to perform various tasks with firewall rules.

See also:

- *Blocking Websites*
- *Allowing Remote Access to the GUI*
- *Preventing RFC 1918 Traffic from Exiting a WAN Interface*
- *Configuring pfSense Software for Online Gaming*

Troubleshooting problems with firewall behavior.

See also:

- *Troubleshooting Thread Errors with Hostnames in Aliases*
- *Troubleshooting Firewall Rules*
- *Troubleshooting Asymmetric Routing*
- *Troubleshooting Blocked Log Entries for Legitimate Connection Packets*

CHAPTER

FIFTEEN

NETWORK ADDRESS TRANSLATION

15.1 Port Forwarding

Port forwarding is a type of inbound NAT which enables access to a specific port, port range, or protocol on an internal network device. Port forward rules rewrite the destination of a packet and then forward those packets to the new destination. This functionality is known by different names in various products, including "Destination NAT" or "Inbound NAT".

To *configure port forwarding rules*, navigate to **Firewall > NAT**, on the **Port Forward** tab.

"Port Forward" is a term familiar to users, but it is an oversimplification. Port forward rules can redirect entire protocols such as GRE or ESP in addition to redirecting specific TCP, UDP, and SCTP ports.

The most common uses of port forward rules are for hosting servers or using applications which require inbound connections from the Internet.

See also:

Hangouts Archive to view the May 2016 hangout for NAT on pfSense software version 2.3, The June 2016 hangout on Connectivity Troubleshooting, and the December 2013 Hangout on Port Forward Troubleshooting, among others.

15.1.1 Port Forwarding Risks

In a default configuration, pfSense® software does not allow any connections initiated from hosts on the Internet. This provides protection from attackers scanning the Internet searching for targets. Using port forward rules, and corresponding firewall rules allows connections matching the rules through to local targets, potentially exposing them to attackers. The firewall does not know the difference between a packet with a malicious payload and one that is benign. However, the exposure is limited to the parameters set inside the port forward and corresponding firewall rule. Even if the firewall rules were to allow any inbound packets, the internal host would still only be exposed on ports/protocols set in the port forward rules.

Target systems must use local host-based controls to secure services forwarded through the firewall.

15.1.2 Port Forwarding and Local Services

Port forwards take precedence over services running locally on the firewall, such as the web interface, VPNs, and SSH. For example, say an administrator intended to allow remote web interface access on the WAN from a remote office using HTTPS on TCP port 443. A port forwarding rule on WAN for TCP 443 will take precedence and the web interface will no longer be accessible from the remote office through the WAN. This does not affect access on other interfaces, only the interface on the overlapping port forward rule.

15.1.3 Port Forward Rule Precedence

For inbound packets, port forward rules take precedence over 1:1 NAT rules. This allows port forwards to override 1:1 NAT rule behavior and forward specific ports to different internal targets or even the firewall itself.

See also:

- *Ordering of NAT and Firewall Processing*
- *1:1 NAT Rule Precedence*

15.1.4 Port Forward Rule Options

When creating or editing a port forward rule, the following settings are available:

Disable
Toggles whether this rule is active.

No RDR (NOT)
Prevents the firewall from redirecting packets matching the rule as they arrive. This is necessary if the packets would otherwise match a port forward rule, but those packets should not be redirected by the firewall.

This can override a forwarding action, which may be useful to allow access to a service on the firewall itself on an IP address used in 1:1 NAT rules, or another similar advanced scenario.

Most configurations will not use this field.

Interface
The ingress interface where this rule will apply.

Typically, this is WAN or a WAN-type interface, but in some special cases it could be LAN or another internal interface.

Address Family
The address family for which this port forward rule will apply, either *IPv4* or *IPv6*.

This must match the address family of the destination IP address and redirect target IP address.

When an interface contains addresses of both families, the firewall will use the appropriate matching address. Additionally, when selecting an interface it must have an IP address which matches this family.

Protocol
The IP protocol this port forward rule will match.

This must be set to match the type of service the rule is forwarding, whether it is *TCP*, *UDP*, or another available choice.

Most common services are *TCP* or *UDP*, but consult the documentation for the service to confirm the correct value. The *TCP/UDP* option forwards both TCP and UDP together in a single rule.

> **Warning:** Though it can be tempting to use the *TCP/UDP* choice when the exact value is unclear, the best practice is to only allow in the specific protocol the service requires. Otherwise, a different unapproved daemon could be run on the target host to receive connections for the other protocol.

The page will display controls for ports when the **Protocol** is set to *TCP*, *UDP*, *TCP/UDP*, or *SCTP*.

Source

The source network and port which this rule will match to limit redirection.

These options are hidden behind an **Advanced** button by default, and set to *any* source.

The **Source** options restrict which source IP addresses and ports can access this port forward rule. These restrictions are not typically necessary, but can offer increased security if remote sources are limited and static.

If the port forward must be reachable from any location on the Internet, the source must be *any*. For restricted access services, use an alias here so only a limited set of IP addresses may access the port forward.

Unless the service absolutely requires a specific source port, the **Source Port Range** must be left as *any* since nearly all clients will use randomized source ports.

Tip: If a role requires source restrictions, a better practice is to use a VPN to allow access to these services instead of port forwarding or firewall rules. However, that may not always be viable.

Destination

The destination address this rule will match.

This is the external IP address which is the original destination of an incoming packet which the rule will redirect.

For port forward rules on WAN, in most cases this is *WAN Address*. If multiple public IP addresses are available, the destination may be a *Virtual IP* on WAN.

Invert Match changes the behavior, so the port forward rule will match any packet that does not match the specified destination.

Destination port range

The destination port or range of ports this rule will match.

This is the external destination port which is the original destination port on an incoming packet which the rule will redirect.

Note: For a port forward rule to match a single port, enter the port in the **From port** box and leave the **To port** box blank.

The drop-down controls contain a list of common services. When set to *Other*, the fields can accept a port number or port alias name. When this field contains a port alias name, the same port alias name must be set as the **Redirect target port**.

Redirect target IP

The IP address where the rule will forward (redirect) incoming packets. The firewall will replace the original destination address on the packet with this address and then forward it to the new destination.

When using an IPv6 redirect target, it must be of the same scope as the destination. For example, it is not possible to forward between link-local scope addresses (`fe80::/10`) and local (`::1`).

Note: This field is only compatible with aliases that contain a **single IP address**.

An alias containing multiple addresses will cause the rule to redirect connections to target hosts in a round-robin fashion, but it is not ideally suited to that task. If one of the target hosts is down, the rule will still forward connections to the unreachable target.

For situations requiring forwarding to multiple hosts, such as load balancing or failover scenarios, use the *HAProxy package*.

Redirect target port
The port number where this rule will forward (redirect) incoming packets. The firewall will replace the original destination port on the packet with this port and then forward it to the new destination.

If the destination port is a range, this field defines where the forwarded port range begins. When a rule forwards a range of ports, e.g. `19000-19100`, a rule can only define the local starting point since the number of ports in the range must be identical.

This field allows a rule to redirect connections to a different port on the outside than the on the inside. For example, a port forward rule can redirect connections on external port `8888` to local port `80` for HTTP on an internal server. A list of common services is available to pick from in the drop-down box.

The drop-down controls contain a list of common services. When set to *Other*, the fields can accept a port number or port alias name. When this field contains a port alias name, the same port alias name must be set as the **Destination port range**.

Description
Text describing the rule, such as its intended behavior or name of a service. The best practice is to clearly describe the purpose of the rule in this field.

The description is optional and does not affect functionality of the rule.

No XMLRPC Sync
Prevents this rule from synchronizing to other High Availability cluster members via XMLRPC.

See also:

High Availability

Warning: This does not prevent a rule on a secondary node from being overwritten by the primary.

NAT Reflection
Overrides global NAT reflection behavior for this rule.

Use system default
Respects the global NAT reflection settings.

Enable (NAT+Proxy)
Always performs NAT reflection for this rule using the NAT+Proxy method.

Enable (Pure NAT)
Always performs NAT reflection for this rule using the pure NAT method.

Disable
Never performs NAT reflection for this rule.

See also:

- *NAT Reflection*
- *Accessing Port Forwards from Local Networks*

Filter Rule Association

Optionally create a firewall rule which passes connections which matches the same packet parameters as this port forward rule.

> **Warning:** This is *very* important. A port forwarding rule only defines which packets the firewall will redirect, a firewall rule allows those packets to *pass* through to the redirect target. Without a firewall rule to pass the packets, the firewall will drop them and there is nothing for it to forward.

The available choices are:

None
Does not create a firewall rule.

Add associated filter rule
Creates a firewall rule linked with this port forward rule.

If an administrator changes the port forward rule, the firewall automatically updates the firewall rule to match.

The firewall replaces this choice with an entry pointing to the new associated firewall rule on save.

This is the default behavior and the best choice for most use cases.

Add unassociated filter rule
Creates a firewall rule separate from this port forward rule.

If an administrator changes the port forward rule, they must also manually update the firewall rule to match.

This can be useful if the firewall requires options or restrictions that cannot be set on a port forward rule.

Pass
Uses a special PF keyword on the port forward rule that passes matching packets without the need of a firewall rule.

Because there is no firewall rule, the firewall forwards **all** packets matching this rule to the target host.

Note: Rules using *Pass* can only work on the interface containing the default gateway for the firewall, they do not work with Multi-WAN.

15.1.5 Configuring Port Forward Rules

Port forward rules are located at **Firewall > NAT**, on the **Port Forward** tab.

See also:

The list of port forward rules works the same as the list of firewall rules for management. See *Introduction to the Firewall Rules screen* for details.

To create a new port forward rule:

- Navigate to **Firewall > NAT**, **Port Forward** tab
- Click  **Add** to create a new port forward rule at the top of the list
- Configure the rule as described in *Port Forward Rule Options*
- Click **Save**
- Click **Apply Changes**

Tracking Changes to Port Forward Rules

The firewall tracks changes to port forward rules in the same manner as firewall rules.

See also:

Rule Information

Firewall rules automatically created by associated NAT rules are also marked as such in the tracking information on the associated firewall.

15.1.6 Port Forwarding Configuration Examples

This section contains example scenarios and configurations for port forwarding.

See also:

- *Accessing Port Forwards from Local Networks*
- *Configuring NAT for a VoIP PBX*
- *Redirecting Client DNS Requests*

Example Port Forward Rule for HTTP (TCP/80)

This example demonstrates how to configure a port forward rule to redirect HTTP (TCP port `80`) inbound on WAN destined to the WAN IP address to the internal host at `10.3.0.15`.

To configure this rule:

- Navigate to **Firewall > NAT**, **Port Forward** tab
- Click  **Add** to create a new port forward rule at the top of the list
- Configure the rule as described in *Port Forward Rule Options* with the following options:

  Interface
  WAN

Address Family
IPv4

Protocol
TCP

Destination
WAN Address

Destination Port Range
HTTP

Redirect Target IP
Address or Alias, `10.3.0.15`

Redirect Target Port
HTTP

Description
`HTTP to web server`

- Click **Save**
- Click **Apply Changes**

Figure *Port Forward Example* shows the port forward configuration page with these values:

After clicking **Save**, the GUI displays the port forward rule list with the new rule, as shown in Figure *Port Forward List*.

Double-check the firewall rule at **Firewall > Rules** on the tab matching the port forward rule interface. The rule should show that it will pass connections in to the target IP address on the expected target port, as shown in Figure *Port Forward Firewall Rule*.

If everything is configured correctly, the port forward rule will work when tested from outside the network.

See also:

Troubleshooting NAT Port Forwards

Accessing an Internal Target Using the External Destination Address

It takes additional work for clients behind the firewall to access services redirected by port forward rules using the external address. For example, when internal users must access public redirected services using hostnames from DNS.

This can be accomplished either by using NAT reflection or Split DNS, depending on the needs and capabilities of the clients.

See also:

- *NAT Reflection*
- *Accessing Port Forwards from Local Networks*

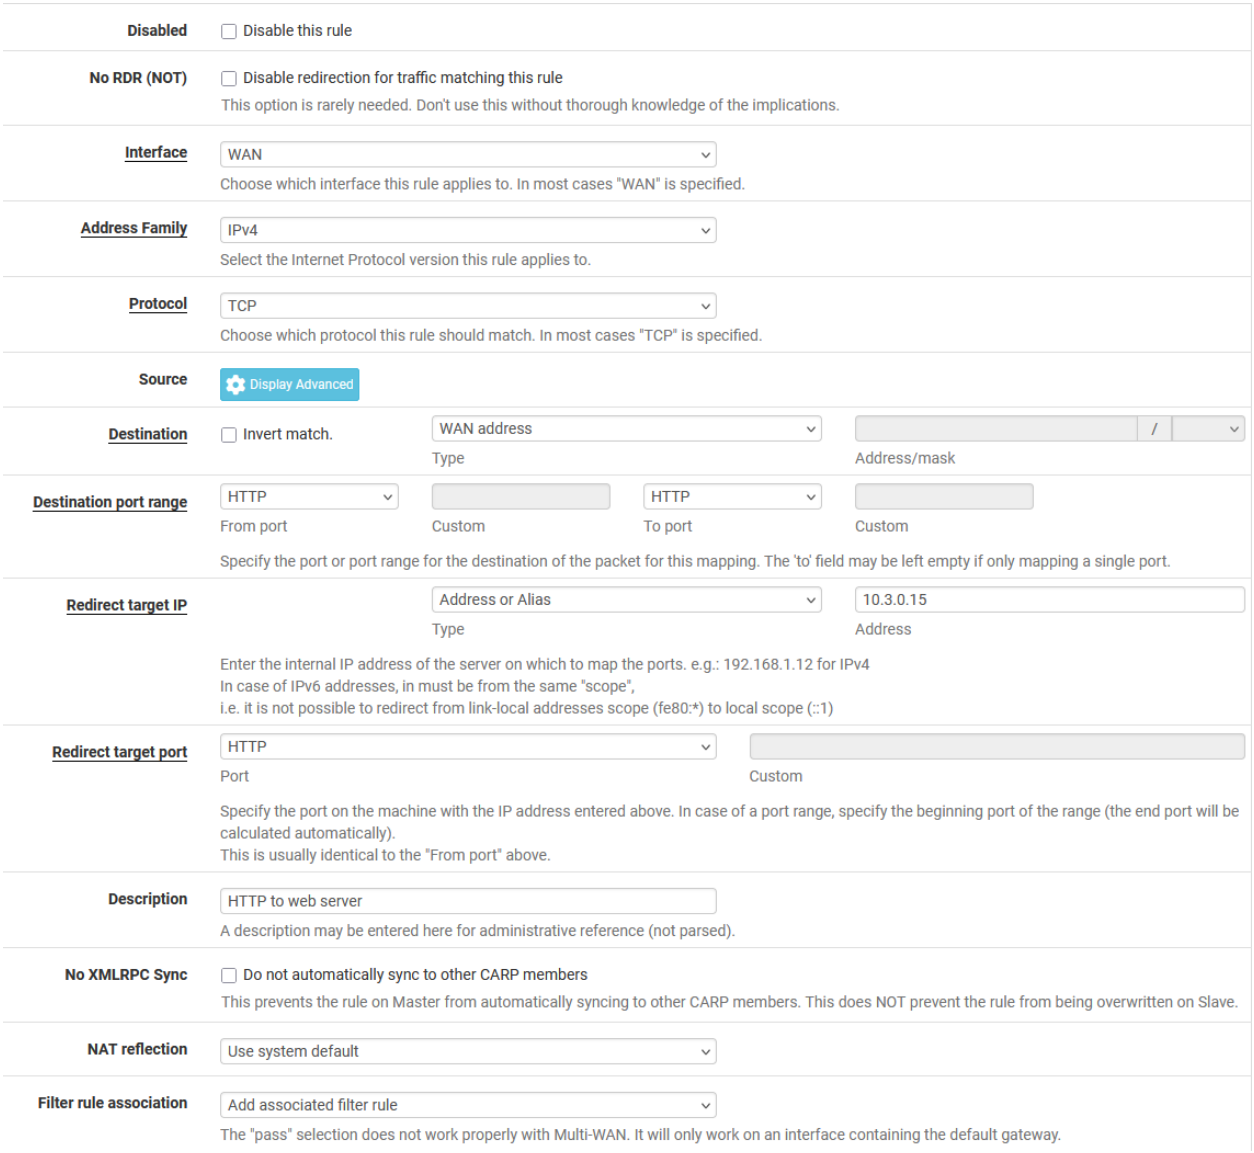


Fig. 1: Port Forward Example

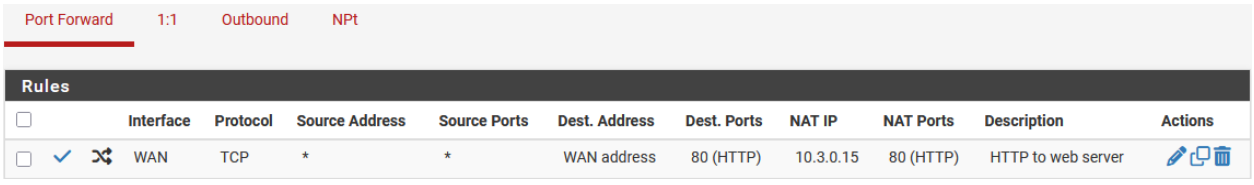


Fig. 2: Port Forward List

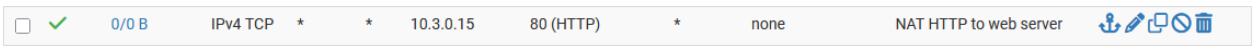


Fig. 3: Port Forward Firewall Rule

Outbound Redirection with Port Forwards

Port forward rules can transparently redirect packets attempting to egress from an internal network. Port forward rules which specify the *LAN* interface or another internal interface will redirect packets matching the rule to the specified target. This is most commonly used for redirecting all outbound DNS to one server.

See also:

Redirecting Client DNS Requests

Forwarding a Single Port to Multiple Targets

The firewall can only forward a single port to one internal host for each available public IP address when it forwards packets from any source.

For example, if a firewall has only one public IP address available, it can only forward TCP port `80` to one internal web server. Any additional servers must use alternate ports such as `8080`.

If a firewall has five public IP addresses available as Virtual IP addresses, then it can forward TCP port `80` to five internal web servers.

See also:

Virtual IP Addresses

Tip: For services such as HTTP and HTTPS, port sharing may be possible by using the *HAProxy package*. If the requests differ in some way, such as by different request hostnames, a proxy can make more advanced decisions about how to forward requests to multiple internal hosts.

The exception to this behavior is when port forward rules can match based on the source IP address of incoming connections.

If all sources are known in advance, rules can forward connections from those specific remote networks to different internal servers by matching those sources.

Figure *Port Forward Example with Different Sources*.

☐	✓	⇄	WAN	TCP	Bob	*	WAN address	22 (SSH)	10.3.0.5	80 (HTTP)	Redirect SSH from Bob to Bob's Server
☐	✓	⇄	WAN	TCP	Sue	*	WAN address	22 (SSH)	10.3.0.15	80 (HTTP)	Redirect SSH from Sue to Sue's Server

Fig. 4: Port Forward Example with Different Sources

If only some sources are known in advance, rules can forward connections from specific remote networks to one server while a later rule can forward connections from any source on the same port to a different target.

15.1.7 Service Self-Configuration With UPnP IGD & PCP

Some programs support Universal Plug-and-Play (UPnP IGD) or Port Control Protocol (PCP) to automatically configure port forward rules and firewall rules. That approach has even more security concerns, but in home use the benefits often outweigh the potential drawbacks.

See also:

UPnP IGD & PCP

15.2 1:1 NAT

1:1 NAT (pronounced "one-to-one NAT") maps one external IP address to one internal IP address. This usually maps a public IP address on a WAN type interface to a private IP address on a LAN type interface.

For outbound packets, 1:1 NAT rules translate the source address of all packets originating from an internal address to an external address as those packets exit the interface defined in the rule.

For inbound packets, 1:1 NAT rules translate the destination address of all packets initiated from external hosts (e.g. from the Internet) destined for the external IP address on the rule interface to the private IP address on the rule. The firewall then evaluates those packets against the firewall ruleset on the external interface (e.g. WAN). If firewall rules on the external interface permit packets matching a target of the internal IP address, those packets will pass to the internal IP address.

1:1 NAT rules can also translate entire subnets provided the subnets are the same size and align on proper subnet boundaries.

1:1 NAT rules do not change ports on packets, they remain static. For example, on outbound connections, 1:1 NAT rules preserve the source ports used by the local host, similar to using **Static Port** on outbound NAT rules.

15.2.1 1:1 NAT Risks

The risks of 1:1 NAT are essentially the same as the risks for port forwarding if firewall rules on the external interface pass packets, since they both may potentially admit harmful content into the local network.

There is some added risk with 1:1 NAT rules in that firewall rule mistakes can have more severe consequences. With port forward rules, packets are limited by constraints within the port forward rule and the firewall rule. If a port forward rule opens TCP port 80, then an allow all rule on WAN still only permits packets with a destination of TCP port 80 on that internal host. With 1:1 NAT rules in place and an allow all firewall rule on WAN, that rule exposes everything on that internal host to the Internet.

In either case, misconfigurations are always a potential hazard, and this alone is not a reason to avoid 1:1 NAT rules. Keep these behaviors in mind when configuring firewall rules and avoid permitting any connections not required by internal hosts.

15.2.2 1:1 NAT Rule Precedence

For outbound packets, 1:1 NAT rules take precedence over outbound NAT rules. This allows 1:1 NAT rules to override default behaviors defined in outbound NAT rules, including automatic outbound NAT.

For inbound packets, port forward rules take precedence over 1:1 NAT rules. This allows port forwards to override 1:1 NAT rule behavior and forward specific ports to different internal targets or even the firewall itself.

See also:

- *Ordering of NAT and Firewall Processing*
- *Outbound NAT Rule Precedence*
- *Port Forward Rule Precedence*

15.2.3 1:1 NAT Rule Options

When adding or editing a 1:1 NAT rule under **Firewall > NAT** on the **1:1** tab, each rule has the following options:

Disabled
Controls whether this 1:1 NAT rule is active.

No BINAT (NOT)
Excludes packets matching this 1:1 NAT rule from 1:1 NAT when packets would otherwise match another rule after this rule in the list.

Interface
The interface where the 1:1 NAT translation will take place, typically a WAN type interface.

The 1:1 NAT rule will only affect packets entering and exiting this specific interface. If there are multiple WAN type interfaces, nudging packets to use this interface may require static routing, policy routing, or equivalent configurations.

Address Family
The address family this rule will match: *IPv4* or *IPv6*. The rule will only match and act upon packets with the selected address family.

Tip: Though 1:1 NAT rules can be used with IPv6, in most cases *IPv6 Network Prefix Translation (NPt)* is a better fit for translating the prefix of IPv6 packets.

External subnet IP
The outside IP address to which this rule will translate the **Internal IP** address as packets enter or exit the **Interface**.

This is typically a Virtual IP address on **Interface**, or an IP address routed to the firewall on **Interface**.

Internal IP
The local IP address this rule will translate to the **External subnet IP** address as packets enter or exit the **Interface**.

This is typically an IP address behind this firewall. The device with this address must use this firewall as its gateway directly (attached) or indirectly (via static route).

Specifying a subnet mask here will translate the entire network matching the subnet mask. For example using `x.x.x.0/24` will translate anything in that subnet to its equivalent host address in the external subnet.

Destination
An optional restriction which limits packets the 1:1 NAT rule will match. The default value is *Any*. This field can utilize aliases when set to *Address or Alias*.

When a **Destination** value is present, the 1:1 NAT rule will only take effect for packets from or to addresses matching the value, depending on the direction:

- For outbound packets, the source must match the **Internal IP** address and the destination IP address must match this value.
- For inbound packets, the source IP address must match this value and the destination IP address must match the **External subnet IP** address.

Description
Text describing the rule, such as its intended behavior or name of a service. The best practice is to clearly describe the purpose of the rule in this field.

The description is optional and does not affect functionality of the rule.

NAT reflection
Overrides global NAT reflection behavior for this rule.

Use system default
Respects the global NAT reflection settings.

Enable
Always performs NAT reflection for this rule.

Disable
Never performs NAT reflection for this rule.

See also:

NAT Reflection

15.2.4 Configuring 1:1 NAT Rules

1:1 NAT rules are located at **Firewall > NAT**, on the **1:1** tab.

See also:

The list of 1:1 NAT rules works the same as the list of firewall rules for management. See *Introduction to the Firewall Rules screen* for details.

To create a new 1:1 NAT rule:

- Add a Virtual IP for the public IP address to be used for the 1:1 NAT rule as described in *Virtual IP Addresses*
- Navigate to **Firewall > NAT**, **1:1** tab
- Click  **Add** to create a new 1:1 NAT rule at the top of the list
- Configure the 1:1 NAT rule as described in *1:1 NAT Rule Options*
- Click **Save**
- Click **Apply Changes**

15.2.5 1:1 NAT Configuration Examples

This section contains example scenarios and configurations for 1:1 NAT.

See also:

- *Accessing Port Forwards from Local Networks*
- *Configuring NAT for a VoIP PBX*

Example Single IP Address 1:1 Configuration

This section demonstrates how to configure a 1:1 NAT rule with a single internal and external IP address.

In this example, a 1:1 NAT rule will map a local mail server to a public IP address. The external address is `198.51.100.210` which is a Virtual IP address on the WAN interface. The mail server resides on a DMZ segment using internal IP address `10.3.1.15`.

The 1:1 NAT rule to map `198.51.100.210` to `10.3.1.15` is shown in Figure *1:1 NAT Rule*.

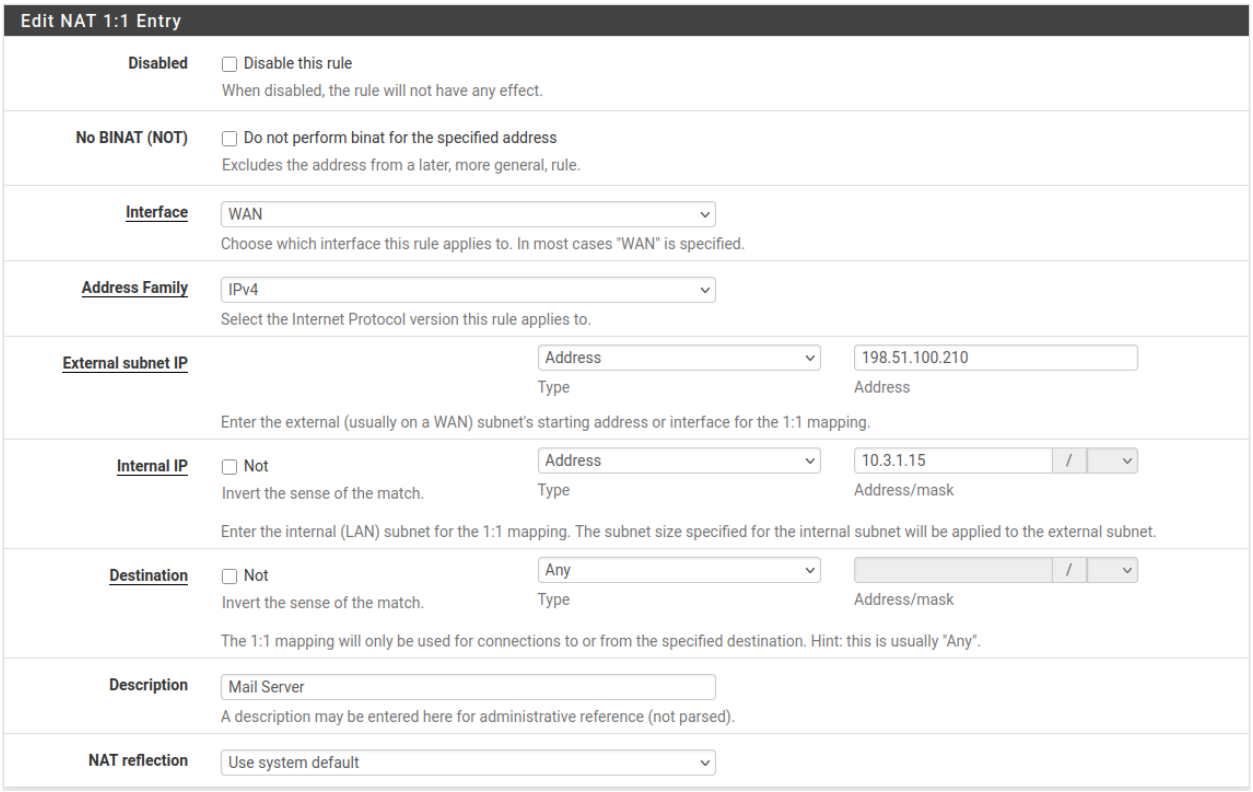


Fig. 5: 1:1 NAT Rule

Example IP Address Range 1:1 Configuration

A 1:1 NAT rule can map multiple external IP addresses using CIDR ranges. In this example, a 1:1 NAT rule maps a /30 CIDR range of IP addresses.

See also:

See *CIDR Summarization* for more information on summarizing networks or groups of IP addresses inside a larger subnet using CIDR notation.

Table 1: /30 CIDR Mapping Matching Final Octet

External IP	Internal IP
198.51.100.64/30	10.3.1.64/30
198.51.100.64	10.3.1.64
198.51.100.65	10.3.1.65
198.51.100.66	10.3.1.66
198.51.100.67	10.3.1.67

Figure *1:1 NAT rule for /30 CIDR range* shows how to configure 1:1 NAT to achieve the mapping listed in Table */30 CIDR Mapping Matching Final Octet*.

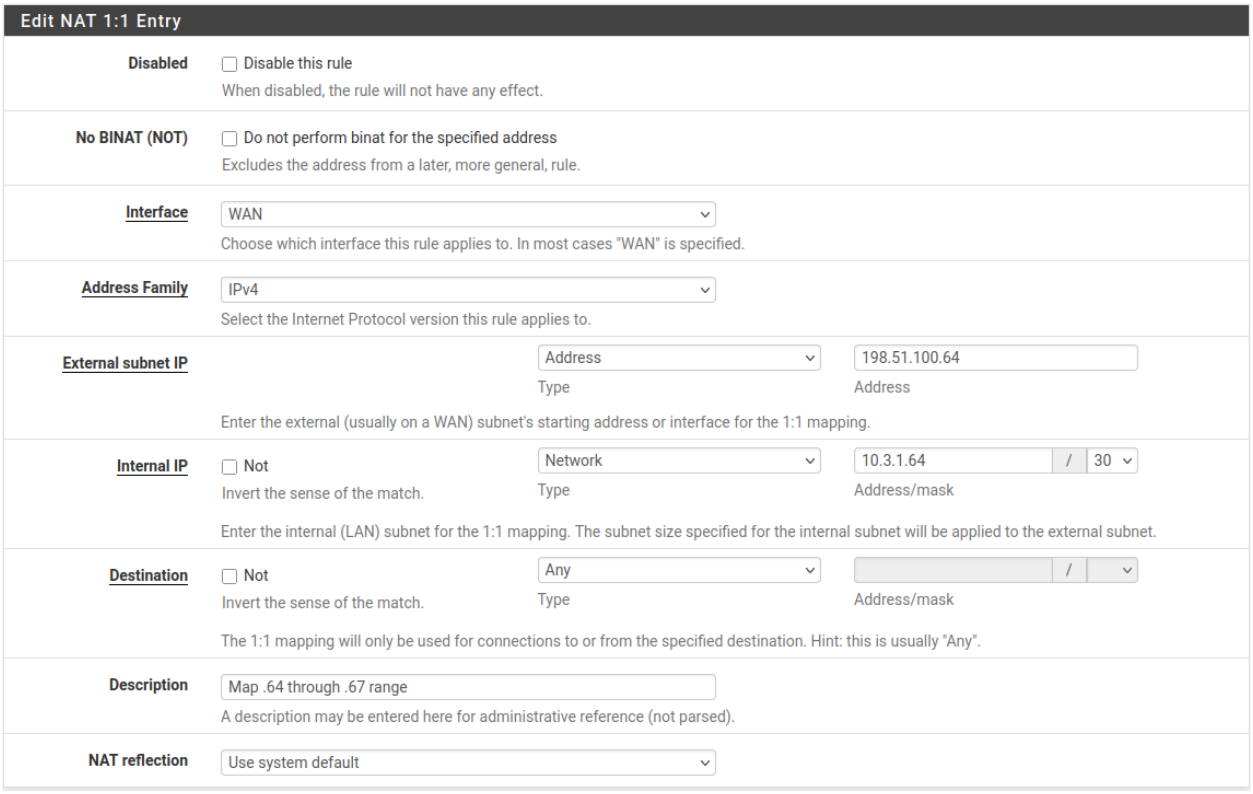


Fig. 6: 1:1 NAT rule for /30 CIDR range

The last octet of the IP addresses need not be the same on the inside and outside, but doing so makes it simpler to follow. For example, Table */30 CIDR Mapping Non-Matching Final Octet* is also valid.

Table 2: /30 CIDR Mapping Non-Matching Final Octet

External IP	Internal IP
198.51.100.64/30	10.3.1.200/30
198.51.100.64	10.3.1.200
198.51.100.65	10.3.1.201
198.51.100.66	10.3.1.202
198.51.100.67	10.3.1.203

Choosing an addressing scheme where the last octet matches makes the layout easier to understand and hence maintain.

1:1 NAT on the WAN IP, aka "DMZ" on Consumer Gateways

Some consumer gateways have what they call a "DMZ" feature which forwards all ports and protocols destined to the WAN IP address to a device on the LAN. This behavior is similar to a 1:1 NAT mapping between the WAN IP address and the IP address of the internal device.

However, the manner in which these devices use the term "DMZ" is unrelated to a DMZ in formal networking terminology. In fact, it is almost the opposite. A host in a traditional network DMZ is in an isolated network away from the other LAN hosts, secured away from the Internet and LAN hosts alike. In contrast, a "DMZ" host in the consumer gateway meaning is not only on the same network as the LAN hosts, but completely exposed to incoming Internet packets with little or no protection.

pfSense® software can perform 1:1 NAT using the WAN IP address to achieve a similar effect, with the caveat that doing so leaves all services running on the firewall itself inaccessible externally. As such, 1:1 NAT rules cannot use the WAN IP address in cases where pfSense software is hosting VPNs of any type, or cases where remote users must access other local services on the firewall. In some cases, this limitation can be mitigated by using a port forward for locally hosted services.

15.3 Ordering of NAT and Firewall Processing

Understanding the order in which PF performs firewall and NAT actions is important when configuring NAT and firewall rules. Figure *Ordering of NAT and Firewall Processing* illustrates the basic logical order. The figure also depicts where tcpdump ties in, since its use as a troubleshooting tool is described later in this documentation in *Packet Capturing*.

Each layer is not always hit in typical configurations, but the use of floating rules or manual outbound NAT or other more complicated configurations can hit each layer in both directions. The diagram only covers basic scenarios for inbound and outbound traffic.

In terms of how the ruleset is processed, the order is:

- Ethernet rules
- Outbound NAT rules
- Inbound NAT rules such as Port Forwards (including `rdr pass` and UPnP IGD & PCP)
- Rules dynamically received from RADIUS for IPsec and OpenVPN clients
- Internal automatic rules (pass and block for various items like lockout, snort, DHCP, etc.)
- User-defined rules:
  - Rules defined on the *floating tab*
  - Rules defined on interface group tabs (Including IPsec and OpenVPN)
  - Rules defined on interface tabs (WAN, LAN, OPTx, etc.)
- Automatic VPN rules

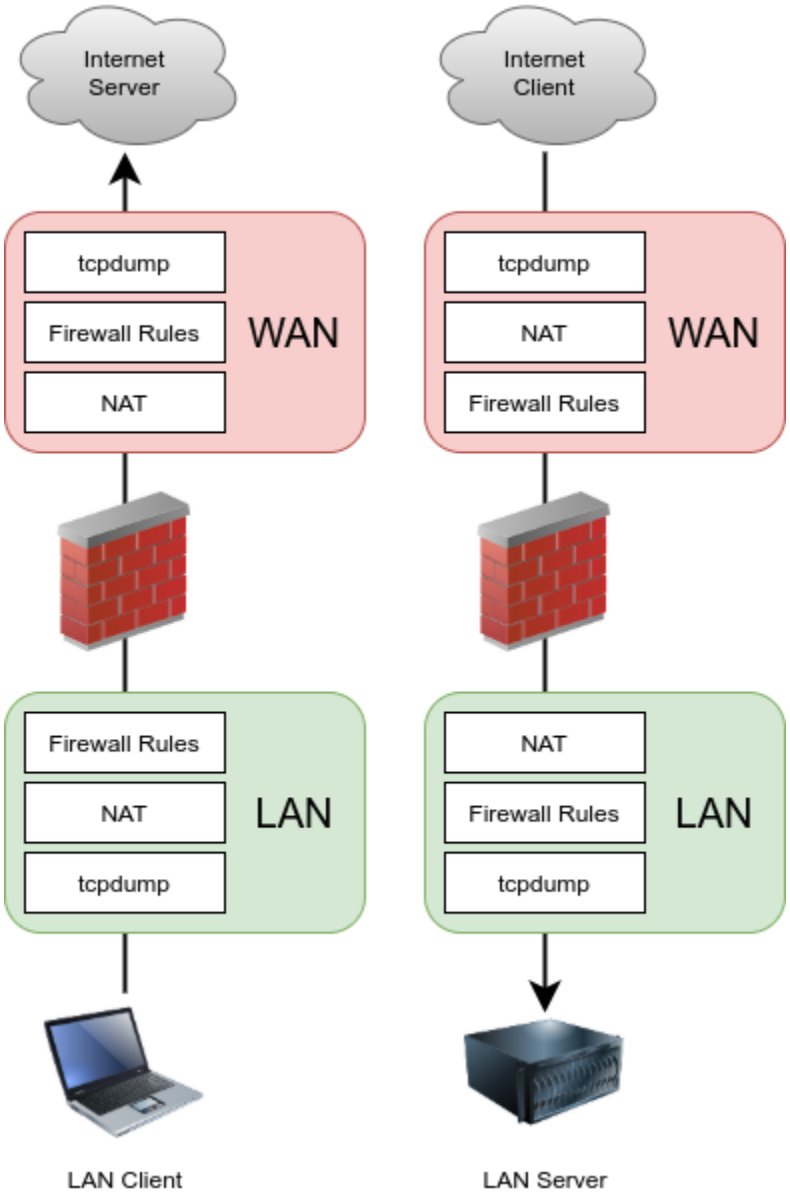


Fig. 7: Ordering of NAT and Firewall Processing

15.3.1 Firewall/NAT Processing Order Example

Traffic from LAN to WAN is processed as described in the following more detailed example. If a type of rules do not exist or do not match, they are skipped.

- Ethernet rules **inbound** on LAN
- Port forward rules then 1:1 NAT rules on the LAN interface (e.g. proxy or DNS redirects)
- Firewall rules for the LAN interface:
  - Floating rules **inbound** on LAN
  - Rules for interface groups including the LAN interface
  - LAN tab rules

  Note: This includes NAT64 rules, at which point the source and destination of matching IPv6 packets are translated to IPv4 as determined by the NAT64 rule parameters.

- 1:1 NAT rules then Outbound NAT rules on WAN
- Floating rules that match **outbound** on WAN
- Ethernet rules **outbound** on WAN

In this case, port forwards on WAN and WAN tab firewall rules do not apply.

For traffic initiated on the WAN, the order is the same, but direction is reversed:

- Ethernet rules **inbound** on WAN
- Port forward rules then 1:1 NAT rules on the WAN interface (e.g. public services)
- Firewall rules for the WAN interface:
  - Floating rules **inbound** on WAN
  - Rules for interface groups including the WAN interface
  - WAN tab rules
- 1:1 NAT rules then Outbound NAT rules on LAN
- Floating rules that match **outbound** on LAN
- Ethernet rules **outbound** on LAN

`tcpdump` is always the first and last thing to see traffic, depending on the direction. First, on the incoming interface before any NAT and firewall processing, and last on the outbound interface. It shows what is on the wire. (See *Packet Capturing*)

See also:

See *Rule Processing Order* for more information about the firewall rule processing order.

Ethernet Rules notes

Ethernet (L2) rules are processed before NAT and traditional firewall rules (Floating, group, or per-interface) in the inbound direction. For outbound traffic, Ethernet rules are processed last after all other rules.

See also:

Ethernet (Layer 2) Rules

Floating Rules notes

Floating rules without `quick` set process as "last match wins" instead of "first match wins". Therefore, if a floating rule is set without `quick` and a packet matches that rule, then it also matches a later rule, the later rule will be used. This is the opposite of the other tab rules (groups, interfaces) and rules with `quick` set which stop processing as soon as a match is made. See *Floating Rules* for more details on how floating rules operate.

15.3.2 Extrapolating to additional interfaces

The previous diagram and lists only illustrate a basic two interface LAN and WAN deployment. When working with additional interfaces, the same rules apply. Traffic between two internal interfaces behaves the same as LAN to WAN traffic, though the default NAT rules will not translate traffic between internal interfaces so the NAT layer does not do anything in those cases. If Outbound NAT rules exist that match traffic between internal interfaces, it will apply as shown.

15.3.3 Rules for NAT

On the way into an interface, NAT applies before firewall rules, so if the destination is translated on the way in (e.g. port forwards rules then 1:1 NAT on WAN), then the firewall rules must match the translated destination. In the typical case of a port forward on WAN, this means the rule must match a destination of the target private IP address on LAN.

For example, with a port forward for TCP port 80 on WAN with an automatically added firewall rule, Figure *Firewall Rule for Port Forward to LAN Host* shows the resulting firewall rule on WAN. The internal IP address on the port forward is 10.3.0.15. Whether using port forward rules or 1:1 NAT rules, firewall rules on all WAN interfaces must use the internal IP address as the destination.

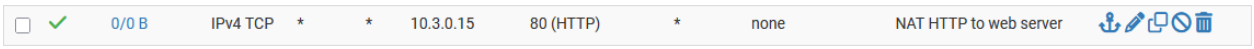


Fig. 8: Firewall Rule for Port Forward to LAN Host

On the way out of an interface, outbound NAT applies before firewall rules, so any floating rules matching outbound on an interface must match the source after it has been translated by 1:1 NAT rules or outbound NAT rules.

15.4 NAT Reflection

NAT reflection refers to the ability to access external services from the internal network using the external (usually public) IP address, the same as if the client were on the Internet. While many commercial and open source firewalls do not implement this functionality, pfSense® software has solid NAT reflection functionality, though some environments require a split DNS infrastructure to fully accommodate this feature.

Split DNS is the best practice to allow local clients to access local resources so that the firewall is not involved in "hairpin" routing of internal services between local hosts.

See also:

Split DNS is covered at the end of this section in *Split DNS*.

15.4.1 Configuring NAT Reflection

To enable NAT Reflection globally:

- Navigate to **System > Advanced** on the **Firewall & NAT**
- Locate the **Network Address Translation** section of the page
- Configure the NAT Reflection options as follows:

  NAT Reflection mode for Port Forwards
  There are three available choices for NAT Reflection mode for port forwards, they are:

  Disable
  NAT Reflection will not be performed, but it may be enabled on a per-rule basis.

  NAT + Proxy
  Enables NAT Reflection using a helper program to send packets to the target of the port forward. This is useful in setups where the interface and/or gateway IP address used for communication with the target cannot be accurately determined at the time the rules are loaded. Reflection rules for use with the proxy are not created for ranges larger than 500 ports and will not be used for more than 1000 ports total between all port forwards. This mode does not work with UDP, only with TCP. Because this is a proxy, the source address of the traffic, as seen by the server, is the firewall IP address closest to the server.

  Pure NAT
  Enables NAT Reflection using only NAT rules in PF to direct packets to the target of the port forward. It has better scalability, but it must be possible to accurately determine the interface and gateway IP address used for communication with the target at the time the rules are loaded. There are no inherent limits to the number of ports other than the limits of the protocols. All protocols available for port forwards are compatible. If servers are on the same subnet as clients, the **Enable automatic outbound NAT for Reflection** option will mask the source of the traffic so it flows properly back through the firewall.

  Reflection Timeout
  This option is only relevant to *NAT + Proxy* mode, and controls how long the NAT proxy daemon will wait before closing a connection. Specify the value in seconds.

  Enable NAT Reflection for 1:1 NAT
  This option allows clients on internal networks to reach locally hosted services by connecting to the external IP address of a 1:1 NAT entry. To fully activate the feature, check both **Enable NAT Reflection for 1:1 NAT** and **Enable automatic outbound NAT for Reflection**. The latter option is only necessary if clients and servers are in the same subnet.

  Enable automatic outbound NAT for Reflection
  When enabled, this option activates additional NAT rules for 1:1 NAT Reflection and Pure NAT mode NAT Reflection for port forwards. These additional rules mask the source address of the client to ensure reply traffic flows back through the firewall. Without this, connections between the client and server will fail as the server will reply directly back to the client using its internal IP address. The client will drop the connection since it expects a reply from the public IP address.

- Click **Save** to activate the new NAT reflection options

NAT Reflection Caveats

NAT reflection is a hack as it loops traffic through the firewall when it is not necessary. Because of the limited options PF allows for accommodating these scenarios, there are some limitations in the pfSense NAT + Proxy reflection implementation. Port ranges larger than 500 ports do not have NAT reflection enabled in NAT + Proxy mode, and that mode is also effectively limited to only working with TCP. The other modes require additional NAT to happen if the clients and servers are connected to the same interface of the firewall. This extra NAT hides the source address of the client, making the traffic appear to originate from the firewall instead, so that the connection can be properly established.

Split DNS is the best means of accommodating large port ranges and 1:1 NAT. Maintaining a split DNS infrastructure is required by many commercial firewalls even, and typically isn't a problem.

15.4.2 Split DNS

The best practice alternative to NAT reflection is deploying a split DNS infrastructure. Split DNS refers to a DNS configuration where, for a given hostname, public Internet DNS resolves to public IP address, and DNS on the internal network resolves to the internal, private IP address. The exact methods to accommodate split DNS vary depending on the specifics of the DNS infrastructure in an organization, but the end result is the same: NAT reflection is not necessary because hostnames resolve to the private IP addresses inside the network and clients contact the servers directly.

Split DNS allows servers to see the true client IP address, and connections between servers and clients in the same subnet go directly, rather than unnecessarily involving the firewall.

The only case that does not work properly with split DNS is when the external and internal port numbers are different. With split DNS, the port number has to be the same in both places.

DNS Resolver/Forwarder Overrides

If pfSense software is acting as the DNS server for internal hosts, then host overrides in the DNS Resolver or DNS Forwarder can provide split DNS functionality.

To add an override to the DNS Resolver:

- Navigate to **Services > DNS Resolver**
- Click  the under **Host Overrides** to reach the Host Override Options page
- Configure the host override as needed, using the internal IP address of the server. See *Host Overrides*. Figure *DNS Resolver Override for example.com* shows an example of a DNS override for `example.com` and `www.example.com`.
- Click **Save**
- Click **Apply Changes**

The DNS Forwarder works identically in this regard. If the DNS Forwarder is enabled instead of the DNS Resolver, add the overrides there.

An override is required for each hostname in use behind the firewall.

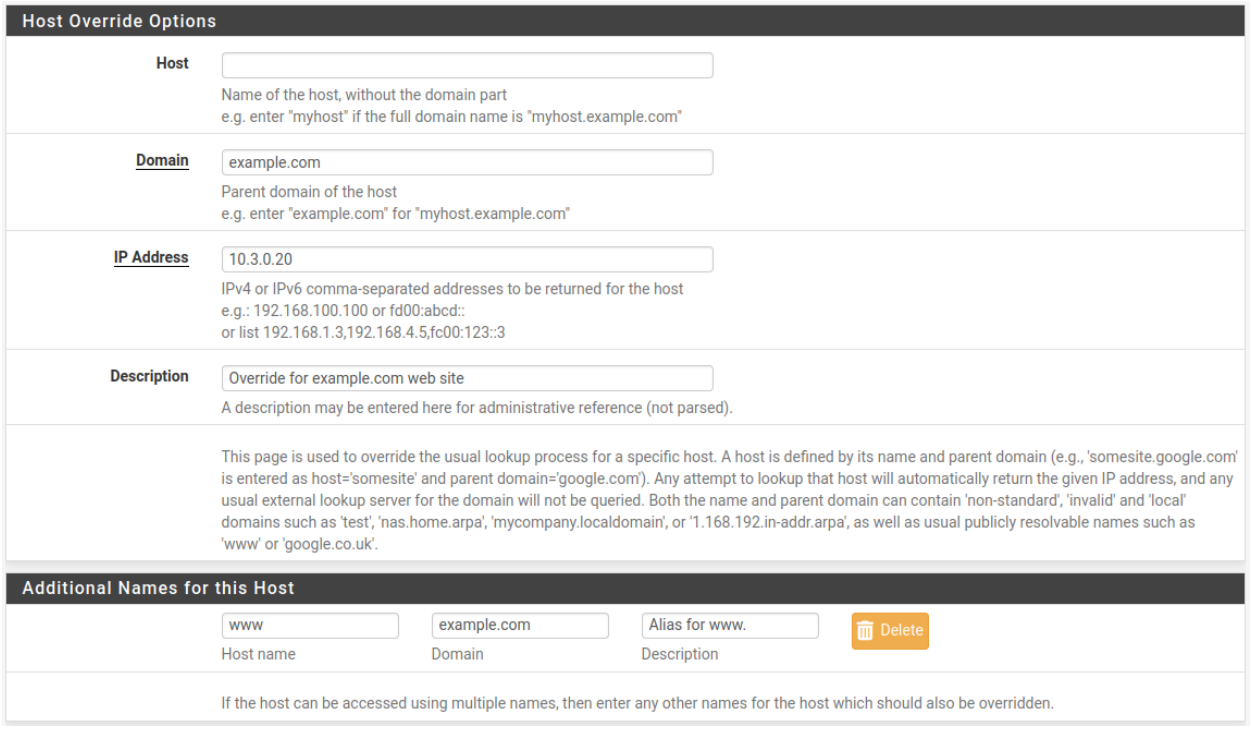


Fig. 9: DNS Resolver Override for example.com

Internal DNS servers

When using a separate DNS server on an internal network, such as Microsoft Active Directory, zones must be created by the DNS server administrator for all domains hosted inside the network, along with all other records for those domains (A, CNAME, MX, etc.).

In environments running the BIND DNS server where the public DNS is hosted on the same server as the private DNS, BIND's views feature is used to resolve DNS differently for internal hosts than external ones. Other DNS servers may offer similar functionality. Check their documentation for information.

15.5 Outbound NAT

Outbound NAT, also known as *Source NAT*, controls how pfSense® software will translate the source address and ports of traffic leaving an interface. To configure Outbound NAT, navigate to **Firewall > NAT**, on the **Outbound** tab.

See also:

- *Ordering of NAT and Firewall Processing*

15.5.1 Outbound NAT Mode

There are four possible **Modes** for Outbound NAT:

Automatic Outbound NAT
: The default option, which automatically performs NAT from internal interfaces, such as LAN, to external interfaces, such as WAN.

Hybrid Outbound NAT
: Utilizes manual rules while also using automatic rules for traffic not matched by manually entered rules. This mode is the most flexible and easy to use for administrators who need a little extra control but do not want to manage the entire list manually.

Manual Outbound NAT
: Only honors the manually entered rules, and nothing more. Offers the most control, but can be tough to manage, and any changes made to internal interfaces or WANs must be accounted for in the rules by hand. If the list is empty when switching from automatic to manual, the list is populated with rules equivalent to the automatically generated set.

Disable Outbound NAT
: Disables all outbound NAT. Useful if the firewall contains only routable addresses (e.g. public IP addresses) on all LANs and WANs.

  See also:

  Disabling Outbound NAT

Note: Even if rules are present in the Outbound NAT screen, the firewall will not honor those rules unless the **Mode** is set to **Hybrid Outbound NAT** or **Manual Outbound NAT**.

When changing the **Mode** value, click the **Save** button to store the new value.

In networks with a single public IP address per WAN, there is usually no reason to enable manual outbound NAT. If some manual control is necessary, hybrid mode is the best choice. In environments with multiple public IP addresses and complex NAT requirements, manual outbound NAT offers more fine-grained control over all aspects of translation.

Tip: For environments using High Availability with CARP, it is important to NAT outbound traffic to a CARP VIP address, as discussed in *High Availability*. This can be accomplished in either hybrid or manual mode.

As with other types of rules on pfSense software, the firewall considers outbound NAT rules from the top of the list down, and it uses the first rule which matches a packet.

Note: Outbound NAT **only** controls what happens to traffic *as it leaves an interface*. It **does not** control the interface a packet uses to exit the firewall. That is handled by the routing table (*Static Routes*) or policy routing (*Policy routing*).

15.5.2 Default Outbound NAT Rules

When set to the default **Automatic Outbound NAT** mode, the firewall maintains a set of NAT rules which translate connections sourced from internal networks to the IP address of a WAN interface through which those connections egress. The firewall also includes static route networks and remote access VPN networks in automatic NAT rules.

When outbound NAT is configured for **Automatic** or **Hybrid** modes, the GUI displays the automatic rules in the lower section of the screen labeled **Automatic Rules**.

If the Outbound NAT rule list is empty, switching to **Manual Outbound NAT** and saving will generate a full set of rules equivalent to the automatic rules.

15.5.3 Outbound NAT Rule Precedence

For outbound packets, 1:1 NAT rules take precedence over outbound NAT rules. This allows 1:1 NAT rules to override default behaviors defined in outbound NAT rules, including automatic outbound NAT.

See also:

- *Ordering of NAT and Firewall Processing*
- *1:1 NAT Rule Precedence*

15.5.4 Outbound NAT Rule Options

Outbound NAT rules are very flexible and are capable of translating traffic in many ways.

Unlike Firewall rules, the GUI displays the NAT rules in a single page with a column indicating which outbound interface is associated with each rule.

Click  from the Outbound NAT page to add a rule to the **top** of the list. Click  to add a rule to the bottom. Place specific rules at the top, and more general rules at the bottom. The rules are processed by the firewall starting at the top of the list and working down, and the firewall uses the first rule to match. Rules may be moved to match in the desired order.

When editing Outbound NAT rules, the options are split into multiple sections.

General Options

These options control general rule behavior and matching parameters.

Disabled
: Toggles whether this rule is active.

Do not NAT
: Prevents the firewall from applying NAT to packets matching the rule as they leave. This is necessary if the traffic would otherwise match a NAT rule, but must not have NAT applied.

  One common use for this is to add a rule exception so that the firewall does not apply NAT to its own IP addresses, especially in the case of CARP where such NAT would break Internet communication from a secondary node while it is in backup mode.

Interface
: The egress interface where this NAT rule will apply.

  Typically, this is WAN or a WAN-type interface, but in some special cases it could be LAN or another internal interface.

Address Family
The address family for which this NAT rule will apply. In nearly all cases this will be *IPv4* or *IPv4+IPv6*.

Note: While it is possible to perform traditional outbound style overload NAT for IPv6 addresses, the best practice is to not apply NAT to IPv6 traffic. See *IPv6 and NAT* for details.

Protocol
The IP protocol this NAT rule will match.

In most cases, Outbound NAT will apply to *any* protocol, but occasionally it is necessary to restrict the protocol upon which the NAT will act. For example, to only perform static port NAT for UDP traffic from a PBX.

Source
The source network which this rule will match and then translate as it leaves the selected **Interface**.

This field can utilize aliases if the **Type** is set to *Network or Alias*.

This is typically a LAN, DMZ, or VPN subnet. The **Source Port** is nearly always left blank to match all ports.

Note: Avoid using a source address of *any* as that will also match traffic from the firewall itself. This will cause problems with gateway monitoring and other firewall-initiated traffic.

Destination
The destination address this rule will match.

This field can utilize aliases if the **Type** is set to *Network or Alias*.

In most cases, the **Destination** remains set to *any* so that the firewall will translate all traffic going to any destination through the selected **Interface**. However, the destination can be restricted when necessary. For example, to translate packets heading to a specific destination differently, such as only performing static port NAT to SIP trunk addresses.

Translation

These options control how NAT translates the source address for packets matching the rule.

Address
Controls what happens to the source address of traffic matching this rule.

The **Address** drop-down contains all defined Virtual IP addresses and subnets, and *Network or Alias* to manually enter a subnet for translation.

Most commonly, this is set to the Interface Address, so the firewall translates the source IP address of connections the IP address of **Interface**, e.g. the WAN IP address.

Note: NAT rules cannot use an alias containing subnets for translation. They can only utilize host aliases or a single manually entered subnet.

Pool Options
Controls how outbound NAT translates source addresses when it has multiple addresses available for translation.

When an outbound NAT rule is configured to use a host alias or manually entered subnet, the rule can translate to a pool of addresses. This can help in large NAT deployments or in areas where several clients require static port to reach the same destination.

Only *Round Robin* types work with host aliases. Any type may be used with a subnet.

Default
Does not define any specific algorithm for selecting a translation address from the pool.

Round Robin
Loops through each potential translation address in the alias or subnet in turn.

Round Robin with Sticky Address
Works the same as *Round Robin* but maintains the same translation address for a given source address as long as states from the source host exist.

Random
Selects a translation address for use from the subnet at random.

Random with Sticky Address
Selects an address at random, but maintains the same translation address for a given source address as long as states from the source host exist.

Source Hash
Uses a hash of the source address to determine the translation address, ensuring that the translated address is always the same for a given source IP address.

Bitmask
Applies the subnet mask and keeps the last portion identical. For example if the source address is `10.10.10.50` and the translation subnet is `192.2.0.0/24`, the rule will change the address to `192.2.0.50`. This works similarly to 1:1 NAT but only in the outbound direction.

Port or Range
Specifies a specific *source* port for translation.

This is almost always left blank, but a client could require this behavior if the client selects a random source port when the server requires a specific source port.

Static Port
Maintains the original source port of the client packet when translating the source IP address. Checking this option disables the **Port** entry box.

Some protocols require this behavior, such as IPsec without NAT-T. Some protocols behave better with this behavior, such as SIP and RTP.

> **Warning:** Do not enable this option with **Endpoint-independent Mapping** checked.

See also:

Static Port

Endpoint-independent Mapping
Enables experimental "Port Restricted Cone" endpoint-independent outbound NAT on this rule. These NAT mappings attempt to preserve port and external address mappings for clients when speaking to multiple remote hosts, but in a dynamic way that does not rely on static port NAT.

This behavior facilitates use cases such as online gaming, peer-to-peer connections, and VoIP.

> **Warning:** Do not enable this option with **Static Port** checked.

See also:

Endpoint-independent Port Restricted Cone NAT

Misc

No XMLRPC Sync
Prevents this rule from synchronizing to other High Availability cluster members via XMLRPC.

See also:

High Availability

> **Warning:** This does not prevent a rule on a secondary node from being overwritten by the primary.

Description
Text describing the rule, such as its intended behavior or name of a service. The best practice is to clearly describe the purpose of the rule in this field.

The description is optional and does not affect functionality of the rule.

Rule Information

This section works similar to the same section in firewall rules to track NAT rule creation and updates.

See also:

Firewall Rule Information

When switching from Automatic Outbound NAT mode to Manual Outbound NAT mode, the firewall marks that change as the source of the rules it creates.

15.5.5 Configuring Outbound NAT Rules

Outbound NAT rules are located at **Firewall > NAT**, on the **Outbound** tab.

See also:

The list of outbound NAT rules works the same as the list of firewall rules for management. See *Introduction to the Firewall Rules screen* for details.

To create a new outbound NAT rule:

- Navigate to **Firewall > NAT**, **Outbound** tab
- Select **Hybrid Outbound NAT** or **Manual Outbound NAT**

  See also:

  Outbound NAT Mode
- Click **Save**

- Click  to create a new outbound NAT rule to the top of the list
- Configure the rule as described in *Configuring Outbound NAT Rules*
- Click **Save**
- Click **Apply Changes**

15.5.6 Outbound NAT Configuration Examples

This section contains example scenarios and configurations for outbound NAT.

See also:

- *Configuring NAT for a VoIP PBX*
- *Configuring NAT for VoIP Phones*
- *Configuring pfSense Software for Online Gaming*

Disabling Outbound NAT

If local interfaces only utilize public IP addresses, and thus NAT is not required to pass traffic through the firewall, disable NAT for routable local subnets. This can be achieved in several ways:

- If NAT is not required for any interface, set the outbound NAT mode to **Disable**.
- Using Hybrid Outbound NAT, create a rule set with **Do not NAT** set to match the routable subnets.
- Using Manual Outbound NAT, delete (or do not create) any NAT rules matching the routable subnets.

See also:

- *Outbound NAT Mode*

In any of the above cases, outbound NAT will no longer be active for those source IP addresses and pfSense software will then route public IP addresses without translation.

Static Port

By default, pfSense software rewrites the source port on all outgoing connections except for UDP port `500` (IKE for IPsec VPN traffic). Some operating systems do a poor job of source port randomization and some do not randomize source ports at all. This makes IP address spoofing easier and makes it possible to fingerprint hosts behind the firewall from their outbound traffic. Rewriting the source port eliminates these potential security vulnerabilities. Outbound NAT rules, including the automatic rules, will show  in the **Static Port** column on rules set to randomize the source port.

Source port randomization breaks some rare applications. The default Automatic Outbound NAT ruleset disables source port randomization for UDP `500` because it will almost always be broken by rewriting the source port. Outbound NAT rules which preserve the original source port are called **Static Port** rules and have  on the rule in the **Static Port** column. All other traffic has the source port rewritten by default.

Other protocols, such as those used by game consoles, may not work properly when NAT rewrites the source port. To disable this functionality, use the **Static Port** option.

To add a rule for a device which requires static source ports:

- Navigate to **Firewall > NAT**, **Outbound** tab
- Select **Hybrid Outbound NAT**

  See also:

  Outbound NAT Mode
- Click **Save**
- Click  to add a new NAT rule to the top of the list
- Configure the rule as described in *Configuring Outbound NAT Rules*

  The rule must match the traffic that requires static port, such as the source address of a PBX or a game console
- Locate the **Translation** section of the page
- Check **Static Port**
- Uncheck **Enable EIM-NAT for UDP connections** (if checked)
- Click **Save**
- Click **Apply Changes**

After making that change, the firewall will preserve the source port on outgoing traffic matching the rule. The best practice is to use strict rules when utilizing static port to avoid any potential conflict if two local hosts use the same source port to talk to the same remote server and port using the same external IP address.

Endpoint-independent Port Restricted Cone NAT

Endpoint-independent NAT Mappings are an experimental feature in pfSense software, currently implemented as "Port Restricted Cone" outbound NAT. "Port Restricted Cone" NAT mappings attempt to preserve port and external address mappings for UDP client connections when speaking to multiple remote hosts, but in a dynamic way that does not rely on *static port NAT*.

These rules enable a client to receive the same external IP address and port on outbound UDP connections when communicating with multiple remote hosts using the same source port. This behavior facilitates use cases such as online gaming, peer-to-peer connections, and VoIP.

Unlike static port NAT, Port Restricted Cone NAT changes the source port when mapping to an external address and port. However, the way it maintains the association of client connections with a specific external source port achieves a similar effect without some of the drawbacks. In particular, this behavior helps avoid issues with multiple local clients using the same source port to reach the same remote host.

Tip: Do not use endpoint-independent Port Restricted Cone NAT if the remote service requires a specific external source port, use *static port NAT* instead which preserves the source port on the client connection.

Inbound communication from a remote host and port is only possible after a local client initiates first contact to that remote host and port. While this is more secure, it is not yet capable of "full cone" NAT which some use cases may require such as certain types of online gaming.

Note: When this is enabled, do not enable **Static Port**. **Static Port** always preserves the original source port, so EIM-NAT has no effect when **Static Port** is also enabled.

To add an outbound NAT rule that performs endpoint-independent Port Restricted Cone NAT:

- Navigate to **Firewall > NAT**, **Outbound** tab
- Select **Hybrid Outbound NAT**

  See also:

  Outbound NAT Mode
- Click **Save**
- Click  to add a new NAT rule to the top of the list
- Configure the rule as described in *Configuring Outbound NAT Rules*

  The rule must match the traffic that should have endpoint-independent Port Restricted Cone NAT mappings, such as the source address of a PBX or game console.
- Locate the **Translation** section of the page
- Uncheck **Static Port** (if checked)
- Check **Enable EIM-NAT for UDP connections**
- Click **Save**
- Click **Apply Changes**

The best practice is to use strict rules when utilizing this option to avoid any potential conflicts.

15.6 Choosing a NAT Configuration

The best NAT configuration for a given deployment depends primarily on the number of public IP addresses available and the number of local services that require inbound access from the Internet.

15.6.1 Single Public IP Address per WAN

When only a single public IP per WAN is available, NAT options are limited.

For outbound connections, Outbound NAT is typically best left on automatic mode or hybrid with some small customizations.

For inbound connections, 1:1 NAT rules can be used with WAN IP addresses, but that can have drawbacks. In this case, the best practice is to only use port forwards.

15.6.2 Multiple Public IP Addresses per WAN

When multiple public IP addresses are available per WAN, numerous options are available for inbound and outbound NAT configuration. Port forwards, 1:1 NAT, and Hybrid or Manual Outbound NAT may all be desirable, depending on the needs of the site.

15.7 NAT and Protocol Compatibility

This section covers a sampling of protocols that have difficulties with NAT in pfSense® software, and how to work around these issues where possible. The majority of network connections from clients have no problems working through connections with NAT. However, some protocols do not work well with NAT and others will not work at all.

Particularly problematic protocols, such as SIP and FTP, embed IP addresses and/or port numbers within packets so when those do not match the post-NAT IP address and port, communication can fail. Protocols which rely on connections using specific source ports, such as SIP from a PBX or IPsec, do not work properly if NAT rewrites the source port. Some protocols are difficult because of limitations of PF itself.

15.7.1 Online Games

Most current game consoles require special configuration to facilitate online play. Some basic games may not have issues with NAT, but it varies greatly by game, platform/console, and online service.

For information on NAT and online gaming, refer to *Configuring pfSense Software for Online Gaming*.

15.7.2 FTP

FTP poses problems with both NAT and firewalls because of the design of the protocol. FTP was initially designed in the 1970s, and the current standard defining the specifications of the protocol was written in 1985. Since FTP was created more than a decade prior to NAT, and long before firewalls were common, it acts in ways that are very unfriendly toward NAT and firewalls.

pfSense software does not include an FTP proxy in the base installation, but there is a client proxy available as an add-on package.

FTP Limitations

Because PF lacks the ability to properly handle FTP traffic without a proxy, and the FTP proxy package implementation is somewhat lacking, there are some restrictions on the usage of FTP.

FTP servers behind NAT

For FTP servers behind NAT, all relevant ports must be manually forwarded in to the server and allowed in firewall rules. Or in the case of 1:1 NAT, only the firewall rules are necessary. Depending on the FTP mode, server software, and client software, some server configuration may also be required.

FTP modes

FTP can act in multiple modes that change the behavior of the client and server, and which side listens for incoming connections. The complications of NAT and firewall rules depend on these modes and whether a remote client is attempting to reach a server behind pfSense, or if a client behind pfSense software is attempting to reach a remote server.

Active Mode

With Active Mode FTP, when a file transfer is requested, the client listens on a local port and then tells the server the client IP address and port. The server will then connect back to that IP address and port in order to transfer the data. This is a problem for firewalls because the port is typically random, though modern clients allow for limiting the range that is used. In the case of a client behind NAT, the IP address given would be a local address, unreachable from the server. Not only that, but a firewall rule would need to be added along with a port forward allowing traffic into this port.

When the FTP proxy package is in use and a client is behind pfSense software connecting to a remote server, the proxy attempts to do three major things: First, it will rewrite the FTP PORT command so that the IP address is the WAN IP address of the firewall, and a randomly chosen port on that IP address. Next, it adds a port forward that connects the translated IP address and port to the original IP address and port specified by the FTP client. Finally, it allows traffic from the FTP server to connect to that "public" port. With Multi-WAN, the proxy will only function on the WAN containing the default gateway.

When everything is working properly, this all happens transparently. The server never knows it is talking to a client behind NAT, and the client never knows that the server isn't connecting directly.

In the case of a server behind NAT, active mode is not usually a problem since the server will only be listening for connections on the standard FTP ports and then making outbound connections back to the clients. The outbound firewall rules must allow the server to make arbitrary outbound connections, and the rules must not policy route those connections out a WAN other than the one that accepted the inbound FTP connection.

Passive Mode

Passive Mode (PASV) acts somewhat in reverse. For clients, it is more NAT and firewall friendly because the server listens on a port when a file transfer is requested, not the client. Typically, PASV mode will work for FTP clients behind NAT without using any proxy or special handling at all.

Similar to the situation in the previous section, when a client requests PASV mode the server will provide the client with its IP address and a random port to which the client can attempt to connect. Since the server is on a private network, that IP address and port will need to be translated and allowed through the firewall. See *FTP Servers and Port Forwards* below for rule requirements. The FTP server must provide the public IP address to which clients connect, but some clients such as FileZilla are smart enough to ignore a given IP address if it is private, and will connect to the original server IP address instead.

Extended Passive Mode

Extended Passive Mode (EPSV) works similar to PASV mode but makes allowances for use on IPv6. When a client requests a transfer, the server will reply with the port to which the client should connect. The same caveats for servers in PASV mode apply here.

FTP Servers and Port Forwards

For FTP servers providing passive mode to clients, the configuration of the FTP server must define a passive port range and must also set the external NAT address, typically the WAN IP address of the firewall. The means of setting these values varies depending on the FTP server software implementation. Consult the FTP server documentation for more information. On the firewall, the passive port range must be forwarded in with port forwards along with TCP port 21.

For FTP servers providing active mode to clients, a port forward is only required for TCP port 21.

FTP Servers and 1:1 NAT

With 1:1 NAT, firewall rules must allow port 21 and the passive port range.

15.7.3 TFTP

Standard TCP and UDP traffic initiates connections to remote hosts using a random source port in the ephemeral port range, which varies by operating system but falls within 1024-65535, and the destination port of the protocol in use. Replies from server to client reverse that: The source port is the client destination port, and the destination port is the client source port. This is how PF associates the reply traffic with connections initiated from inside a network.

TFTP (Trivial File Transfer Protocol) does not follow this convention, however. The standard defining TFTP, RFC 1350, specifies the reply from the TFTP server to client will be sourced from a pseudo-random port number. The TFTP client may choose a source port of 10325 (as an example) and use the destination port for TFTP, port 69. The server for other protocols would then send the reply using source port 69 and destination port 10325. Since TFTP instead uses a pseudo-random source port, the reply traffic will not match the state PF has created for this traffic. Hence, the replies will be blocked because they appear to be unsolicited traffic from the Internet.

TFTP is not a commonly used protocol across the Internet. The only situation that occasionally comes up where this is an issue is with some IP phones that connect to outside VoIP providers on the Internet using TFTP to pull configuration and other information. Most VoIP providers do not require this.

If TFTP traffic must pass through the firewall, a TFTP proxy is available which is configured under **System > Advanced** on the **Firewall & NAT** tab. See *TFTP Proxy* for more information.

15.7.4 GRE / PPTP

The state tracking code in PF for the GRE protocol can only track a single session per public IP address per external server.

The limitations with PPTP in pfSense software are from the same root cause. Because PF can only track a single GRE session per public IP address per external server, if a PPTP VPN connection is in place, only one internal client can connect simultaneously to **the same** a PPTP server on the Internet. A thousand clients can connect simultaneously to a thousand different PPTP servers, but only one simultaneously to a single server. A single client can also connect to an unlimited number of outside PPTP servers.

The only available work around is to use multiple public IP addresses on the firewall, one per client via Outbound or 1:1 NAT. This is not a problem with other types of VPN connections.

Due to the extremely flawed security in PPTP (See *PPTP Warning*), including a complete compromise of the entire protocol, its usage should be discontinued as soon as possible, so this issue is not relevant given the current security standards.

15.8 IPv6 Network Prefix Translation (NPt)

Network Prefix Translation, or NPt for short, works similarly to 1:1 NAT but operates on IPv6 prefixes instead. NPt can be found under **Firewall > NAT** on the **NPt** tab.

NPt takes one prefix and translates it to another. So `2001:db8:1111:2222::/64` becomes `2001:db8:3333:4444::/64` and though the prefix changes, the remainder of the address will be identical for a given host on that subnet.

Warning: NPt on pfSense software **does NOT** function like traditional outbound NAT (overload NAT/PAT). NPt cannot be used to map an internal prefix to a different size prefix or single address in use on a WAN, it must be used with a routed prefix. That type of translation is possible with *outbound NAT rules*.

NPt on pfSense software also **does not** function like NPT66 (RFC 6296), which also changes the host portion, but using specific mathematical rules so it does not change packet checksums. NPt on pfSense software is stateful and maintains the host portion of the address when translating.

There are a few purposes for NPt. With NPt, a LAN can utilize "private" IPv6 space (`fc00::/7`) and it can be translated by NPt to a public, routed, IPv6 prefix as it comes and goes through a WAN. This can help to avoid renumbering the LAN if external providers change, however since anything external that looked for the old prefix must also be adjusted, the usefulness of that can go either way, especially when the configuration must account for avoiding collisions in the `fc00::/7` space for VPN tunnels.

NPt is useful for SOHO IPv6 Multi-WAN deployments. The likelihood that a home or small business end user will have their own provider-independent IPv6 space and a BGP feed is very small. In these cases, the firewall can utilize a routed prefix from multiple WANs to function similarly to Multi-WAN on IPv4. As traffic leaves the second WAN sourced from the LAN subnet, NPt will translate it to the equivalent IP address in the routed subnet for that WAN. The LAN can either use one of the routed prefixes and do NPt on the other WANs, or use addresses in `fc00::/7` and do NPt on all WANs. The best practice is to avoid using the `fc00::/7` space for this task. For more information on Multi-WAN with IPv6, see *Configuring Multi-WAN for IPv6*.

When adding an NPt entry, there are few options to consider as NPt is fairly basic:

Disabled
: Toggles whether this rule is actively used.

Interface
: Selects the Interface where this NPt rule takes effect as the traffic exits.

Source IPv6 Prefix
: The local (e.g. LAN) IPv6 subnet and prefix length, typically the `/64` on LAN or other internal network.

Destination IPv6 Prefix
: The routed external IPv6 subnet and prefix length to which the Internal IPv6 Prefix will be translated. This is **NOT** the prefix of the WAN itself. It must be a network routed to this firewall via **Interface**

Description
: A brief description of the purpose for this entry.

Figure *NPt Example* shows an NPt rule where the LAN IPv6 subnet `2001:db8:1111:2222::/64` will be translated to `2001:db8:3333:4444::/64` as it leaves the `HE_CABLE` interface.

See also:

- *Configuring NAT for a VoIP PBX*
- *Configuring NAT for VoIP Phones*
- *Accessing Port Forwards from Local Networks*
- *Using NAT and FTP without a Proxy*
- *Troubleshooting NAT*
- *Troubleshooting NAT Port Forwards*
- *Troubleshooting 1:1 NAT*
- *Troubleshooting NAT Reflection*

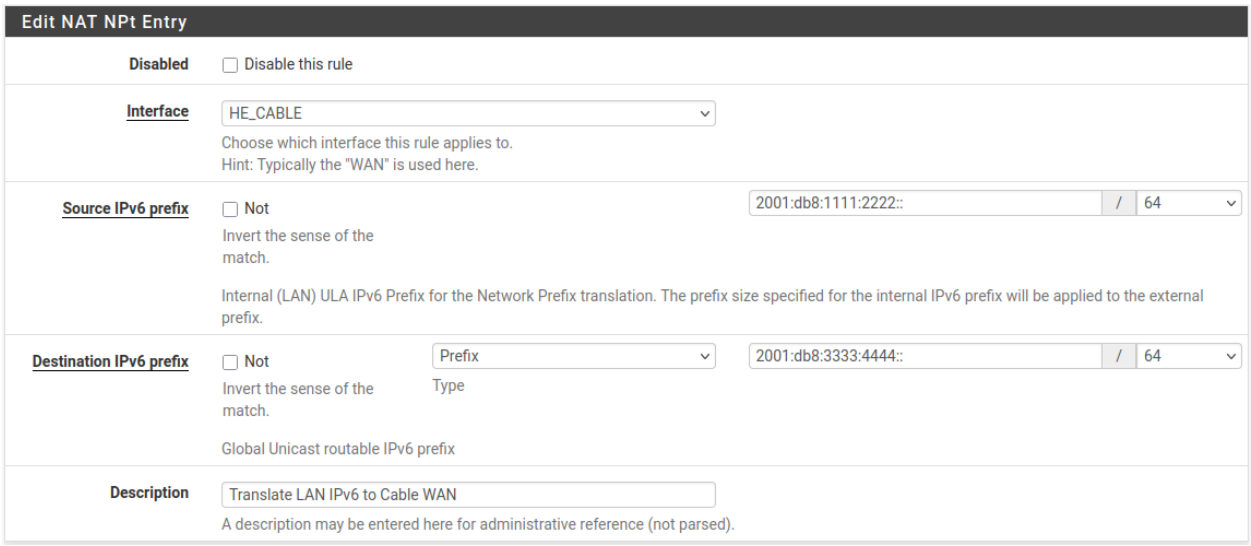


Fig. 10: NPt Example

- *NAT64*

In its most common usage, Network Address Translation (NAT) allows multiple computers using IPv4 to be connected to the Internet using a single public IPv4 address. pfSense® software enables these simple deployments, but also accommodates much more advanced and complex NAT configurations required in networks with multiple public IP addresses.

NAT is configured in two directions: inbound and outbound. Outbound NAT defines how traffic leaving a local network destined for a remote network, such as the Internet is translated. Inbound NAT refers to traffic entering a network from a remote network. The most common type of inbound NAT is *port forwards*, which is also the type many administrators are most familiar with.

See also:

Hangouts Archive to view the May 2016 Hangout on NAT with pfSense software version 2.3 and the earlier August 2014 Hangout on Network Address Translation.

15.9 Default NAT Configuration

This section describes the default NAT configuration present on pfSense software. The most appropriate NAT configuration that can be determined is generated automatically. In some environments, this configuration may not be suitable, and pfSense software fully enables changing it from the web interface. This is a contrast from many other open source firewall distributions, which do not allow the capabilities commonly required in all but small, simple networks.

15.9.1 Default Outbound NAT Configuration

In a typical two-interface setup with LAN and WAN, the default NAT configuration automatically translates Internet-bound traffic to the WAN IP address. When multiple WAN interfaces are configured, traffic leaving any WAN interface is automatically translated to the address of the WAN interface a connection exits.

Static port is automatically configured for IKE (part of IPsec). Static port is covered in more detail in *Static Port*.

For detecting WAN-type interfaces for use with NAT, pfSense software looks for the presence of a gateway selected on the interface configuration if it has a static IP address, or pfSense software assumes the interface is a WAN if it is a dynamic type such as PPPoE or DHCP.

15.9.2 Default Inbound NAT Configuration

By default, nothing is allowed in from the Internet on the WAN interface. If traffic initiated on the Internet must be allowed to reach a host on the internal network, port forwards or 1:1 NAT are required. This is covered in the coming sections.

CHAPTER

SIXTEEN

ROUTING

One of the primary functions of a firewall is routing traffic. This chapter covers several topics related to routing including gateways, static routes, routing protocols, routing of public IP addresses, and displaying routing information.

16.1 Gateways

Gateways are the key to routing; They are routers on directly connected networks through which a host can reach other networks. The kind of gateway most people are familiar with is a *default* gateway, which is the router through which a host will communicate to the Internet or any other networks it doesn't have a more specific route to reach. Gateways are also used for static routing, where certain hosts or networks must be reached via specific routers. On most networks a gateway resides in the same subnet as one of the interfaces on a host. For example, if a firewall has an IP address of `192.168.22.5/24`, then a gateway to another network would have to be somewhere inside `192.168.22.x` if the other network is reachable through that interface.

Note: One notable exception to this is point-to-point interfaces like those used in PPP-based protocols, which often have gateway IP addresses in another subnet because they are not used in the same way.

16.1.1 Gateway Address Families (IPv4 and IPv6)

When working with routing and gateways the functionality and procedures are the same for both IPv4 and IPv6 addresses. However, all the addresses for a given route must involve addresses of the same family. For example, an IPv6 network must be routed through an IPv6 gateway. A route cannot be created for an IPv6 network using an IPv4 gateway address. When working with gateway groups the same restriction applies: All gateways in a gateway group must be of the same address family.

16.1.2 Managing Gateways

Before a gateway can be utilized for any purpose, it must be added to the firewall configuration.

If a gateway will be used for a WAN-type interface, it can be added on the configuration page for that interface (See *Interface Configuration Basics*), or it may be added first manually and then selected from the drop-down list on the interface configuration.

Dynamic interface types such as DHCP, PPPoE, and some assigned tunnel interfaces receive an automatic gateway that is noted as **Dynamic** in the gateway list. The parameters for such gateways can be adjusted the same as the parameters for a static gateway.

Note: Deleting a dynamic gateway will clear its custom settings but the dynamic gateway itself cannot be removed.

To add or manage gateways, navigate to **System > Routing**, **Gateways** tab.

On the screen there are a variety of options to manage gateway entries:

-  **Add** at the bottom of the list creates a new gateway
-  edits an existing gateway
-  creates a copy of an existing gateway
-  disables an active gateway
-  enables a disabled gateway
-  deletes a gateway

See also:

The individual options for gateways are discussed in detail in *Gateway Settings*.

Managing the Default Gateway

The **Default Gateway** section at the bottom of **System > Routing**, **Gateways** tab controls which gateway(s) are used by default when the firewall routes traffic. Traffic from the firewall itself will follow the default gateway, as will traffic passing through the firewall when it does not match policy routing rules or other more specific routes.

There are two controls in the section which set the default gateway for IPv4 and IPv6 respectively.

The default gateway can have one of the following values:

Automatic
: The firewall will automatically use gateways from this list (from the top down) for the default gateway, switching to the next item in the list if gateways fail or are marked down.

  For more control over this behavior, use a gateway group instead.

  Warning: This function can automatically select gateways from VPNs (e.g. IPsec, WireGuard, OpenVPN) and other sources, which may not be what the user intends. These gateways may not allow the firewall to reach the Internet, which may prevent regular traffic flow.

  The best practice for failover is to create a custom gateway group with viable Internet gateways in the intended order by tier rather than relying on the automatic behavior.

Gateway
: The selected single gateway is always used for the default gateway.

Gateway Group
The firewall uses the selected *gateway group* to select a default gateway. It will change from one gateway to another if the preferred default fails.

> **Warning:** This function is not compatible with load balancing, only failover. When using a gateway group for the default gateway, the group must only have one gateway in each tier.

None
No default gateway for the address family will be added to the routing table.

Note: Though default gateway switching is handy for handling traffic from the firewall itself, it is not always the best fit for user traffic. When using gateway switching instead of policy routing the firewall states are not able to track gateway information which allows the firewall to selectively kill states for specific gateways. See *State Killing on Gateway Failure*.

16.2 Gateway Settings

When adding or editing a gateway, the GUI presents a page with the options for controlling gateway behavior.

The only required settings are the **Interface**, **Address Family**, **Name**, and the **Gateway** (IP address).

Interface
The interface through which the gateway is reached. For example, if this is a local gateway on the LAN subnet, choose the *LAN* interface here.

Address Family
Either *IPv4* or *IPv6*, depending on the type of address for this gateway.

Name
The **Name** for the gateway, as referenced in the gateway list, and various drop-down and other selectors for gateways. It can only contain alphanumeric characters, or an underscore, but no spaces. For example: `WANGW`, `GW_WAN`, and `WANGATE` are valid but `WAN GW` is not allowed.

Gateway
The IP address of the gateway. As mentioned previously, this must reside in a subnet directly configured on the selected **Interface**.

Note: Rare cases which require a gateway not in the interface subnet can still function, but require an additional setting. See the **Use Non-Local Gateway** option in *Advanced Gateway Settings* for details.

Disable Gateway Monitoring
By default, the gateway monitoring daemon will ping each gateway periodically to monitor latency and packet loss for traffic to the monitored IP address. This data is used for gateway status information and also to draw the Quality RRD graph. If this monitoring is undesirable for any reason, it may be disabled by checking **Disable Gateway Monitoring**. Note that if the gateway status is not monitored, then Multi-WAN will not work properly as it cannot detect failures.

Disable Gateway Monitoring Action
When set, the gateway monitoring daemon will take no action if the status of the gateway changes.

By default, the gateway monitoring daemon will trigger actions on events when a gateway status changes, such as when a gateway becomes unresponsive, suffers from high latency, or when it recovers and returns to an online status. This allows the firewall to react to changes and perform tasks such as enacting gateway failover and recovery policies, changing gateways in load balancing gateway group tiers, and restarting services for various daemons utilizing the interface on which the gateway resides.

Disabling this behavior can be useful in certain situations if the administrator wants to monitor a gateway without that monitoring causing additional disruptions.

Tip: On firewalls with a single WAN or outbound path, gateway monitoring actions can typically be disabled since no actions are necessary if a gateway fails.

Monitor IP
The **Monitor IP** address option configures the IP address used by the gateway monitoring daemon to determine the gateway status using ICMP echo requests ("pings").

By default, the gateway monitoring daemon will ping the gateway IP address. This is not always desirable, especially in the case where the gateway IP address is local, such as on a cable modem or fiber CPE. In those cases it makes more sense to ping something farther upstream, such as an ISP DNS server or a server on the Internet. Another case is when an ISP is prone to upstream failures, so pinging a host on the Internet is a more accurate test to determine if a WAN is usable rather than testing the link itself. Some popular choices include Google public DNS servers, or popular websites such as Google or Yahoo. If the IP address specified in this box is not directly connected, a static route is added to ensure that traffic to the **Monitor IP** address leaves via the expected gateway. Each gateway must have a unique Monitor IP address.

The status of a gateway as perceived by the firewall can be checked by visiting **Status > Gateways** or by using the **Gateways** widget on the dashboard. If the gateway shows **Online**, then the monitor IP address is successfully responding to pings.

Static Route
By default the firewall adds static routes for gateway monitor IP addresses to ensure traffic to the monitor IP address leaves via the correct interface. Enabling this checkbox overrides that behavior, so the user can manually manage routes or paths to monitor IP addresses.

Force State
When **Mark Gateway as Down** is checked, the gateway will always be considered down, even when pings are returned from the monitor IP address. This is useful for cases when a WAN is behaving inconsistently and the gateway transitions are causing disruption. The gateway can be forced into a *down* state so that other gateways may be preferred until it stabilizes.

State Killing on Gateway Failure
Optionally overrides the default behavior of the global State Killing on Gateway Failure option. See *State Killing on Gateway Failure*.

Use Global Behavior
Uses the default global behavior configured on **System > Advanced**, Miscellaneous tab.

Do not kill states on gateway failure
Excludes this gateway from the default state killing behavior.

Kill states using this gateway when it is down
Kills states for this gateway when it is down even if the global default setting would not do so.

Description
An optional **Description** of the gateway entry for reference. A short note about the gateway or interface it's used for may be helpful, or it may be left blank.

16.2.1 Advanced Gateway Settings

Several parameters can be changed to control how a gateway is monitored or treated in a Multi-WAN scenario. Most users will not need to alter these values. To access the advanced options, click the  **Display Advanced** button. If any of the advanced options are set, this section is automatically expanded. For more information on using multiple WAN connections, see *Multiple WAN Connections*.

Weight
When using Multi-WAN load balancing, if two gateways have different amounts of bandwidth, the **Weight** parameter adjusts the ratio at which connections are sent through each gateway.

For example if WAN1 has 50Mbit/s and WAN2 has 100Mbit/s, weight WAN1 as *1* and WAN2 as *2*. Then for every three connections that go out, one will use WAN1 and two will use WAN2. Using this method, connections are distributed in a way that is more likely to better utilize the available bandwidth. Weight from *1* to *30* may be chosen.

Data Payload
To conserve bandwidth, the `dpinger` daemon sends a ping with a payload size of `1` by default so that minimal data is contained within the ICMP echo request. ICMP data is padded in echo requests under `60` bytes, so any size below `60` will be equivalent on the wire.

Note: The default is not `0` because in some circumstances a CPE, ISP router, or intermediate hop may drop or reject ICMP packets without a payload. Since data sizes under `60` are padded, there is no advantage to using a size of `0`, especially when considering the potential for problems.

Latency Thresholds
The **Latency Thresholds** fields control the amount of latency that is considered normal for this gateway. This value is expressed in milliseconds (ms). The default values are **From** `200` and **To** `500`.

The value in the **From** field is the lower boundary at which the gateway would be considered in a warning state, but not down. If the latency exceeds the value in the **To** field, it is considered down and removed from service. The proper values in these fields can vary depending on what type of connection is in use, and what ISP or equipment is between the firewall and the monitor IP address.

Some common situations may require adjusting these values. For instance some DSL lines operate acceptably even at higher latency, so increasing the **To** parameter to `700` or more would lower the number of times the gateway would be considered down when, in fact, it was passing traffic sufficiently. Another example is a GIF tunnel to a provider such as he.net for IPv6. Due to the nature of GIF tunnels and load on the tunnel servers, the tunnel could be working acceptably even with latency as high as `900` ms as reported by ICMP ping responses.

Packet Loss Thresholds
Similar to **Latency Thresholds**, the **Packet Loss Thresholds** control the amount of packet loss to a monitor IP address before it would be considered unusable. This value is expressed as a percentage, `0` being no loss and `100` being total loss. The default values are **From** `10` and **To** `20`.

The value in the **From** field is the lower boundary at which the gateway would be considered in a warning state, but not down. If the amount of packet loss exceeds the value in the **To** field, it is considered down and removed from service. The proper values in these fields can vary depending on

what type of connection is in use, and what ISP or equipment is between the firewall and the monitor IP address.

As with latency, connections can be prone to different amounts of packet loss and still function in a usable way, especially if the path to a monitor IP address drops or delays ICMP in favor of other traffic. We have observed unusable connections with minor amounts of loss, and some that are usable even when showing 45% loss. If loss alarms occur on a normally functioning WAN gateway, enter higher values in the **From** and **To** fields until a good balance for the circuit is achieved.

Probe Interval
The value in the **Probe Interval** field controls how often a ping is sent to the monitor IP address, in *milliseconds*. The default is to ping twice per second (`500` ms).

In some situations, such as links that need monitored but have high data charges, even a small ping every second can add up. This value can be safely increased so long as it less than or equal to the **Alert Interval** and also does not violate the constraint on the **Time Period** listed below. Lower values will ping more often and be more accurate, but consume more resources. Higher values will be less sensitive to erratic behavior and consume less resources, at the cost of accuracy.

Note: The quality graph is averaged over seconds, not intervals, so as the **Probe Interval** is increased the accuracy of the quality graph is decreased.

Loss Interval
Time in milliseconds before packets are treated as lost. The default is `2000` ms (2 seconds). Must be greater than or equal to the **High Latency Threshold**.

If a circuit is known to have high latency while operating normally, this can be increased to compensate.

Time Period
The amount of time, in milliseconds, over which ping results are averaged. The default is `60000` (60 seconds, one minute). A longer **Time Period** will take more time for latency or loss to trigger an alarm, but it is less prone to be affected by erratic behavior in ping results.

The **Time Period** must be greater than **twice** the sum of the **Probe Interval** and **Loss Interval**, otherwise there may not be at least one completed probe.

Alert Interval
The time interval, in milliseconds, at which the daemon checks for an alert condition. The default value is `1000` (1 second). This value must be greater than or equal to the **Probe Interval**, because an alert could not possibly occur between probes.

Use Non-Local Gateway
The **Use non-local gateway through interface specific route** option allows a non-standard configuration where a gateway IP address exists outside an interface subnet. Some providers attempting to scrape the bottom of the IPv4 barrel have resorted to this in order to not put a gateway into each customer subnet. Do not activate this option unless required to do so by the upstream provider.

16.3 Gateway Groups

Gateway groups are a set of gateways, but are treated as one entity in gateway fields of the GUI. Groups will appear in the gateway drop-downs available on, for example, firewall rule editing.

Gateway groups are managed from the **Gateway Groups** tab on **System > Routing**.

16.3.1 Gateway Group Options

When creating a gateway group, the following options are available:

Group Name
The name of this gateway group. The name must be less than 32 characters in length, and may only contain letters a-z, digits 0-9, and an underscore. This will be the name used to refer to this gateway group in the **Gateway** field in firewall rules. This field is required.

Gateway Priority
This list contains every gateway on the firewall to select which gateways will be a part of this group. The GUI will filter the list address family after the first selection.

Tier
The priority level for this gateway. The value may be from 1-5 or *Never* to exclude the gateway from this group.

Lower values are higher priority. For example, gateways on *Tier 1* are used before gateways on *Tier 2*, and so on.

Gateways on the same tier are used by the firewall for load balancing when possible. Load balancing naturally performs failover as failed gateways are removed from the pool available for load balancing.

Gateways on different tiers result in failover from gateways on lower tiers to those higher tiers. For example, if Tier 1 contains only one gateway, and it fails, then the next tier (Tier 2) is checked for available gateways and the firewall uses those instead, and so on.

> **Warning:** Some firewall features that can use gateway groups are only compatible with failover, not load balancing. For example, when using a gateway group for the default gateway or as a VPN endpoint, each gateway must be on a separate tier.

Virtual IP
When using a gateway group for failover in certain contexts which require binding a specific address, such as IPsec, this option controls which address on an interface is used for that purpose. For example, in an HA pair this could be a CARP VIP used as an endpoint for IPsec tunnels.

Leave it set to the default *Interface Address* when a specific address is not required by any use of the gateway group.

Keep Failover States
Controls the state-killing behavior for the gateway group when configured for failover. This behavior takes effect when a higher-priority gateway returns to an online state. Only affects states created by policy routing rules. This option overrides the global behavior (see *Gateway Monitoring*).

Keep states on gateway recovery
Policy routing states are unaffected when a higher-priority gateway returns to an online

state. Connections established on failover gateways will remain on those gateways until reconnected.

Kill states on gateway recovery
States created by policy routing rules using this gateway group are killed when a higher-priority gateway returns to an online state. This option does not affect traffic from the firewall itself.

Trigger Level

Configures how the firewall manages the gateway group entries when certain types of gateway events occur.

Member Down
Marks the gateway as down only when it is completely down, past one or both of the higher thresholds configured for the gateway. This catches the worst sort of failures, when the gateway is completely unresponsive, but may miss more subtle issues with the circuit that can make it unusable long before the gateway reaches that level.

Packet Loss
Marks the gateway as down when packet loss crosses the lower alert threshold (See *Advanced Gateway Settings*).

High Latency
Marks the gateway as down when latency crosses the lower alert threshold (See *Advanced Gateway Settings*).

Packet Loss or High Latency
Marks the gateway as down for either type of alert.

Description
Text describing the purpose of this gateway group.

16.3.2 Tier Priority Example

Example:

- WANGW: Tier1
- OPT1GW: Tier2
- OPT3GW: Tier3

In the example above OPT1GW would be used if WANGW fails, OPT3GW will be used if both WANGW and OPT1GW fail.

16.3.3 Connection-Based Round-Robin Load Balancing Example

Example:

- WANGW: Tier1
- OPT1GW: Tier1
- OPT3GW: Tier1

In the example above all gateways have the same **Tier** value. When this group is used by a firewall rule, connections matching that rule will perform connection-based round-robin load balancing between all the gateways.

Note: If any of the gateways fail, they are automatically removed from active usage in the group, effectively resulting in failover in addition to load balancing.

16.3.4 See Also

Multi-WAN Guide

16.4 Static Routes

Static routes are used when hosts or networks are reachable through a router other than the default gateway. The firewall knows about the networks directly attached to it, and it reaches all other networks as directed by the routing table. In networks where an internal router connects additional internal subnets, a static route must be defined for those networks to be reachable.

Note: The routers through which these other networks are reached must first be added as gateways. See *Gateways* for information on adding gateways.

Static routes are managed at **System > Routing** on the **Routes** tab.

See also:

- *Accessing Firewall Services over IPsec*
- *Policy Routing Configuration*

16.4.1 Static Route Configuration

When adding or editing a static route, the following options are available:

Destination Network
The network and subnet mask reachable using this route. This may be an IPv4 address (subnet ID), IPv6 prefix, or an *alias*.

> **Warning:** When editing an alias which is used by a static route, the firewall makes changes to the routing table to reflect the new alias content as soon as the alias is saved. The firewall does not wait for the user to click **Apply Changes** in this circumstance.

Gateway
The router through which this network is reached.

Disabled
Check if the static route should not be used, only defined.

Description
Text to describe the route, its purpose, etc.

16.4.2 Managing Static Routes

To add a route:

- Navigate to **System > Routing** on the **Routes** tab
- Click  **Add** to create a new static route
- Fill in the configuration as described in *Static Route Configuration*
- Click **Save**
- Click **Apply Changes**

To manage existing routes, navigate to **System > Routing** on the **Routes** tab. On the screen there are a variety of options to manage routes:

-  edits an existing route
-  creates a copy of an existing gateway
-  deletes a route
-  disables an active route
-  enables a disabled route

16.4.3 Example Static Route

Figure *Static Routes* illustrates a scenario where a static route is required.

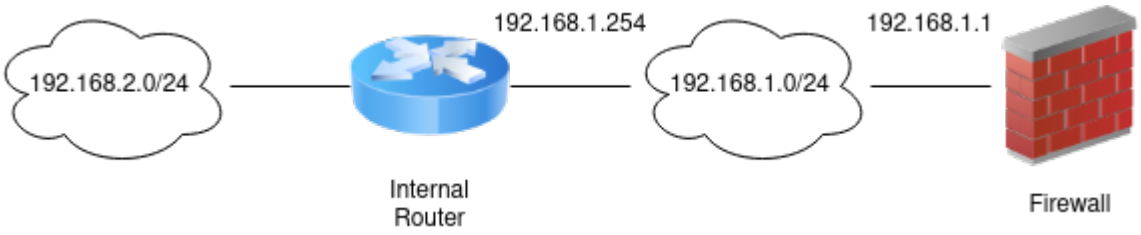


Fig. 1: Static Routes

Because the `192.168.2.0/24` network in Figure *Static Routes* is not on an interface directly connected to the firewall, a static route is required for the firewall to know how to reach that network. Figure *Static Route Configuration* shows the appropriate static route for the above diagram. As mentioned earlier, before a static route may be added a gateway must first be defined.

LAN firewall rules must allow traffic to pass from a source of the networks reachable via static routes on LAN, and outbound NAT must also accommodate these networks.

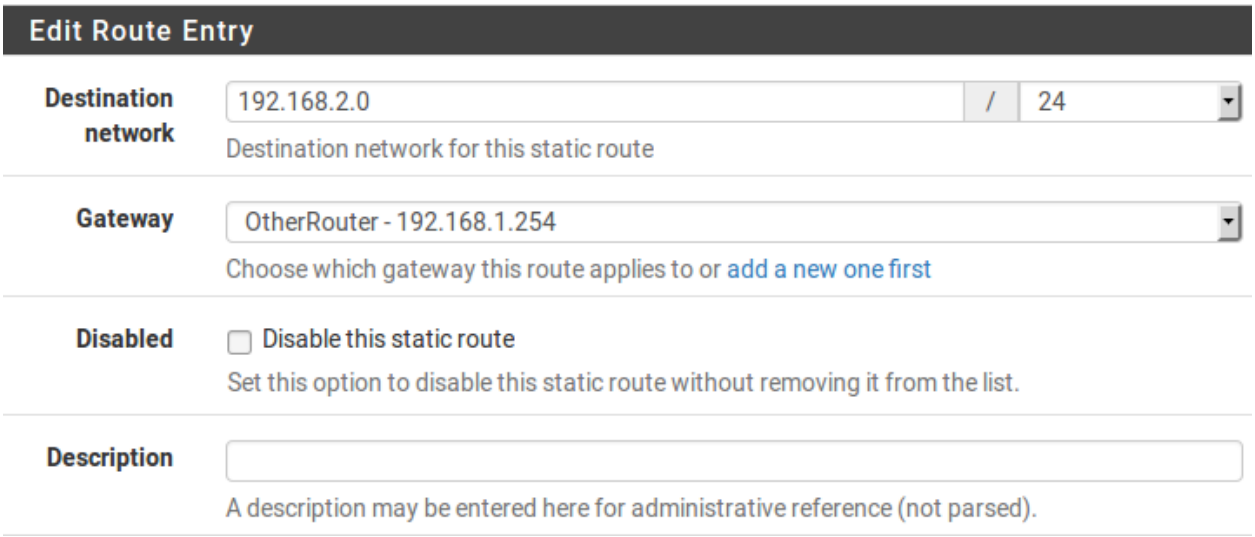


Fig. 2: Static Route Configuration

16.4.4 Bypass Firewall Rules for Traffic on Same Interface

In some environments using static routes, traffic ends up routing asymmetrically. This means the traffic will follow a different path in one direction than the traffic flowing in the opposite direction. Take Figure *Asymmetric Routing* for example.

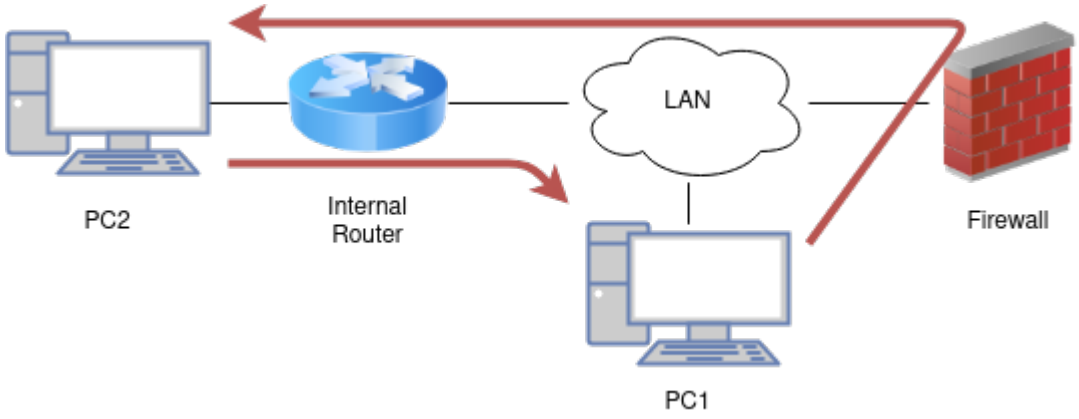


Fig. 3: Asymmetric Routing

Traffic from PC1 to PC2 will go through the firewall since it is the default gateway for PC1, but traffic in the opposite direction will go directly from the router to PC1.

Since this is a stateful firewall, it must see traffic for the entire connection to be able to filter traffic properly. With asymmetric routing such as in this example, any stateful firewall will drop legitimate traffic because it cannot properly keep state without seeing traffic in both directions. This generally only affects TCP, since other protocols do not have a formal connection handshake the firewall can recognize for use in state tracking.

Note: A connection may appear to work for a short time and then stop. This can be from the firewall removing a state which doesn't transition to a fully open state, or it may be from clients expiring an ICMP redirect.

In asymmetric routing scenarios, there is an option in the firewall GUI which can be used to prevent legitimate traffic from being dropped. The option adds firewall rules which allow all traffic between networks defined in static routes using a more permissive set of rule options and state handling. To activate this option:

- Click **System > Advanced**
- Click the **Firewall/NAT** tab
- Check **Bypass firewall rules for traffic on the same interface**
- Click **Save**

Alternatively, firewall rules may be added manually to allow similar traffic. Two rules are needed, one on the interface tab where the traffic enters (e.g. LAN) and another on the **Floating** tab:

- Navigate to **Firewall > Rules**
- Click the tab for the interface where the traffic will enter (e.g. **LAN**)
- Click  **Add** to add a new rule to the top of the list
- Use the following settings:

  Protocol
  TCP

  Source
  The local systems utilizing the static route (e.g. *LAN Net*)

  Destination
  The network on the other end of the route

  TCP Flags
  Check **Any flags** (Under **Advanced Features**)

  State Type
  Sloppy State (Under **Advanced Features**)

- Click **Save**
- Click the **Floating** tab
- Click  **Add** to add a new rule to the top of the list
- Use the following settings:

  Interface
  The interface where the traffic originated (e.g. **LAN**)

  Direction
  Out

  Protocol
  TCP

  Source
  The local systems utilizing the static route (e.g. *LAN Net*)

  Destination
  The network on the other end of the route

  TCP Flags
  Check **Any flags** (Under **Advanced Features**)

State Type
Sloppy State (Under **Advanced Features**)

- Click **Save**

If additional traffic from other sources or destinations is shown as blocked in the firewall logs with TCP flags such as "TCP:SA" or "TCP:PA", the rules may be adjusted or copied to match that traffic as well.

Note: If filtering of traffic between statically routed subnets is required, it must be done on the router and not the firewall since the firewall is not in a position on the network where it can effectively control that traffic.

See also:

- *Troubleshooting Asymmetric Routing*

16.4.5 ICMP Redirects

When a device sends a packet to its default gateway, and the gateway knows the sender can reach the destination network via a more direct route, it will send an ICMP redirect message in response and forward the packet as configured. The ICMP redirect causes a route for that destination to be temporarily added to the routing table of the sending device, and the device will subsequently use that more direct route to reach that destination.

This will only work if the client OS is configured to permit ICMP redirects, which is typically the case by default.

ICMP redirects are common when static routes are present which point to a router on the same interface as client PCs and other network devices. The asymmetric routing diagram from the previous section is an example of this.

ICMP redirects have a mostly undeserved bad reputation from some in the security community because they allow modification of a client routing table. However, they are not the risk that some imply, as to be accepted, the ICMP redirect message must include the first 8 bytes of data from the original datagram. A host in a position to see that data and hence be able to successfully forge illicit ICMP redirects is in a position to accomplish the same end result in multiple other ways.

16.5 Multi-Path Routing

Current versions of pfSense® software include kernels built with the option `ROUTE_MPATH` which enables multi-path routing.

This allows the routing table to contain multiple entries to the same destination, which allows for weight-based balancing of traffic including Equal-cost multi-path routing (ECMP) if all gateways for a destination are weighted the same.

Warning: Multi-path routing can only be utilized by the dynamic routing daemon package *FRR*. The base system GUI cannot manage multiple routes to the same destination, but this capability is planned for future releases.

16.5.1 Multi-Path Routing Behavior

Packets may only take alternate paths when they are different in some way. When there are multiple paths to a destination the operating system computes outbound flow hashing for connections to decide which path to use for a packet. This concept is similar to the *LAGG Hash Algorithm*. The hashing takes the 5-tuple connection property set into account: "(protocol, source address, destination address, source port, destination port)".

For example, testing with ICMP only from one host to another with no variation may never see flows take a second path. In contrast, testing repeated TCP connections could take multiple paths if the source port is randomized. The best way to test is by using multiple sources and destinations passing through the firewall and not from the firewall itself.

16.5.2 View Nexthop Data

The first thing to check is the routing table to ensure that there are multiple routes to the same destination:

```
: netstat -rWn4
Routing tables

Internet:
Destination        Gateway            Flags    Nhop#    Mtu      Netif Expire
[...]
10.5.0.0/16        198.51.100.5       UGS          0   1500        ix3
10.5.0.0/16        203.0.113.5        UGS          0   1500        ix2
```

Note: The route table in the status output of the FRR package will also show multiple entries for the same destination with different `via` addresses.

Next, there are two items to check to verify that multi-path routing is taking effect: Nexthop data and Nexthop group data.

Check the nexthop data with `netstat -4onW` and/or `netstat -6onW`:

```
: netstat -4onW
Nexthop data

Internet:
Idx   Type         IFA                Gateway             Flags      Use Mtu       ␣
↪Netif     Addrif Refcnt Prepend
[...]
32           v4/gw 198.51.100.17      198.51.100.5       GS           0   1500      ␣
↪ix3                1
33           v4/gw 203.0.113.17       203.0.113.5        GS           0   1500      ␣
↪ix2                1
```

Check nexthop group data with `netstat -4OnW` and/or `netstat -6OnW`:

```
: netstat -4OnW
Nexthop groups data

Internet:
GrpIdx  NhIdx     Weight   Slots           Gateway     Netif  Refcnt
34        ------- ------- ------- ----------------- ---------       2
```

(continues on next page)

(continued from previous page)

```
                32      1       1       198.51.100.5        ix3
                33      1       1        203.0.113.5        ix2
```

The outputs of those commands should show both gateways and indicate that they belong to the same "group".

See also:

- *Route Table Contents*
- *Multiple WAN Connections*
- *IPv6 Router Advertisements*
- *FRR Package*
- *Routing Public IP Addresses*
- *Dynamic Routing Protocol Basics*
- *Troubleshooting Gateway Monitoring*
- *Troubleshooting "No buffer space available" Errors*
- *Troubleshooting Network Connectivity*
- *Troubleshooting Traceroute Output*
- *Troubleshooting Website Access*
- *Troubleshooting Routes*

CHAPTER

SEVENTEEN

BRIDGING

17.1 Creating a Bridge

In pfSense® software, bridges are added and removed at **Interfaces > Assignments** on the **Bridges** tab. Using bridges, any number of ports may be bound together easily. Each bridge created in the GUI will also create a new bridge interface in the operating system, named `bridgeX` where `X` starts at 0 and increases by one for each new bridge. These interfaces may be assigned and used like most other interfaces, which is discussed later in this chapter.

To create a bridge:

- Navigate to **Interfaces > Assignments** on the **Bridges** tab.
- Click **Add** to create a new bridge.
- Select at least one entry from **Member Interfaces**. Select as many as needed using `Ctrl`-click.
- Add a **Description** if desired.
- Click **Show Advanced Options** to review the remaining configuration parameters as needed. For most cases they are unnecessary.
- Click **Save** to complete the bridge.

Note: A bridge may consist of a single member interface, which can help with migrating to a configuration with an assigned bridge, or for making a simple span/mirror port.

17.2 Advanced Bridge Options

There are numerous advanced options for a bridge and its members. Some of these settings are quite involved, so they are discussed individually in this section.

17.2.1 (Rapid) Spanning Tree Options

Spanning Tree is a protocol that helps switches and devices determine if there is a loop and cut it off as needed to prevent the loop from harming the network. There are quite a few options that control how spanning tree behaves which allow for certain assumptions to be made about specific ports or to ensure that certain bridges get priority in the case of a loop or redundant links. More information about STP may be found in the FreeBSD ifconfig(8) man page, and on Wikipedia.

Protocol

The **Protocol** setting controls whether the bridge will use IEEE 802.1D Spanning Tree Protocol (*STP*) or IEEE 802.1w Rapid Spanning Tree Protocol (*RSTP*). RSTP is a newer protocol, and as the name suggests it operates much faster than STP, but is backward compatible. The newer IEEE 802.1D-2004 standard is based on RSTP and makes STP obsolete.

Select STP only when older switch gear is in use that does not behave well with RSTP.

STP Interfaces

The **STP Interfaces** list reflects the bridge members upon which STP is enabled. `Ctrl`-click to select bridge members for use with STP.

Valid Time

Set the **Valid Time** for a Spanning Tree Protocol configuration. The default is `20` seconds. The minimum is `6` seconds and the maximum is `40` seconds.

Forward Time

The **Forward Time** option sets the time that must pass before an interface begins forwarding packets when Spanning Tree is enabled. The default is `15` seconds. The minimum is `4` seconds and the maximum is `30` seconds.

Note: A longer delay will be noticed by directly connected clients as they will not be able to pass traffic, even to obtain an IP address via DHCP, until their interface enters forwarding mode.

Hello Time

The **Hello Time** option sets the time between broadcasting of Spanning Tree Protocol configuration messages. The **Hello Time** may only be changed when operating in legacy STP mode. The default is `2` seconds. The minimum is `1` second and the maximum is `2` seconds.

Bridge Priority

The **Bridge Priority** for Spanning Tree controls whether the OS will select this bridge first for blocking if it detects a loop. The default is `32768`. The minimum is `0` and the maximum is `61440`. Values must be a multiple of `4096`. Lower priorities are given precedence, and values lower than `32768` indicate eligibility for becoming a root bridge.

Hold Count

The **Hold Count** for Spanning Tree is the number of packets transmitted before being rate limited. The default is `6`. The minimum is `1` and the maximum is `10`.

Port Priorities

The **Priority** fields set the Spanning Tree priority for each bridge member interface. Lower priorities are given preference when deciding which ports to block and which remain forwarding. Default priority is `128`, and must be between `0` and `240`.

Path Costs

The **Path Cost** fields control the Spanning Tree path cost for each bridge member. The default is calculated from the link speed. To change a previously selected path cost back to automatic, set the cost to `0`. The minimum is `1` and the maximum is `200000000`. Lower cost paths are preferred when making a decision about which ports to block and which remain forwarding.

17.2.2 Cache Settings

Cache Size sets the maximum size of the bridge address cache, similar to the MAC or CAM table on a switch. The default is `100` entries. If there will be numerous devices communicating across the bridge, set this higher.

Cache entry expire time controls the timeout of address cache entries in seconds. If set to `0`, then address cache entries will not be expired. The default is `240` seconds (Four minutes).

17.2.3 Span Port

Selecting an interface as the **Span port** on the bridge will transmit a copy of every frame received by the bridge to the selected interface. This is most useful for snooping a bridged network passively on another host connected to the span ports of the bridge with something such as Snort, tcpdump, etc. The selected span port may not be a member port on the bridge.

17.2.4 Edge Ports / Automatic Edge Ports

If an interface is set as an **Edge port**, it is always assumed to be connected to an end device, and *never* to a switch; It assumes that the port can never create a layer 2 loop. Only set this on a port when it will never be connected to another switch. By default, ports automatically detect edge status, and they can be selected under **Auto Edge** ports to *disable* this automatic edge detection behavior.

17.2.5 PTP Ports / Automatic PTP Ports

If an interface is set as a **PTP port**, it is always assumed to be connected to a switch, and not to an end user device; It assumes that the port can potentially create a layer 2 loop. It should only be enabled on ports that are connected to other switches which have RSTP enabled. By default, ports automatically detect PTP status, and they can be selected under Auto **PTP ports** to *disable* this automatic PTP detection behavior.

17.2.6 Sticky Ports

An interface selected in **Sticky Ports** will have its dynamically learned addresses cached as though they were static once they enter the cache. Sticky entries are never removed from the address cache, even if they appear on a different interface. This could be used a security measure to ensure that devices cannot move between ports arbitrarily.

17.2.7 Private Ports

An interface marked as a **Private Port** will not communicate with any other port marked as a **Private Port**. This can be used to isolate end users or sections of a network from each other if they are connected to separate bridge ports marked in this way. It works similar to "Private VLANs" or client isolation on a wireless access point.

17.3 Bridging and Interfaces

A bridge interface (e.g. *bridge0*) itself may be assigned as interface. This allows the bridge to act as a normal interface and have an IP address placed upon it rather than a member interface.

Configuring the IP address on the bridge itself is best in nearly all cases. The main reason for this is due to the fact that bridges are dependent on the state of the interface upon which the IP address is assigned. If the IP address for the bridge is configured on a member interface and that interface is down, the whole bridge will be down and no longer passing traffic. The most common case for this is a wireless interface bridged to an Ethernet LAN NIC. If the LAN NIC is unplugged, the wireless would be dead unless the IP address was configured on the bridge interface and not LAN. Another reason is that if limiters must be used for controlling traffic, then there must be an IP address on the bridge interface for them to work properly. Likewise, in order for Captive Portal or a transparent proxy to function on an internal bridge the IP address must be configured on the assigned bridge and not a member interface.

17.3.1 Swapping Interface Assignments

Before getting too far into talking about moving around bridge interface assignments, it must be noted that these changes should be made from a port that is **not** involved in the bridge. For example, if bridging WLAN to LAN, make the change from WAN or another OPT port. Alternately, download a backup of `config.xml` and manually make the changes. Attempting to make changes to a port while managing the firewall from that port will most likely result loss of access to the GUI, leaving the firewall unreachable.

Note: It is tempting to create the bridge and then merely swap the interface assignments, but that won't work because it would end up with the bridge added to itself. For example, with LAN and WLAN, create a bridge LANBRIDGE, and then try to swap LAN and LANBRIDGE, it wouldn't work because LAN is specified in the bridge configuration.

Easy Method: Move settings to the new interface

The easiest, though not the quickest, path in the GUI is to remove the settings from the LAN interface individually (IP address, DHCP, etc.) and then activate them on the newly assigned bridge interface.

Quick but Tricky: Reassign the Bridge as LAN

Though this method is a bit trickier than moving the settings, it can be much faster especially in cases where there are lots of firewall rules on LAN or a complex DHCP configuration. In this method, some hoop-jumping is required, but ultimately the bridge ends up as the LAN interface, and it retains the LAN IP address, all the former firewall rules, DHCP, and other interface configuration.

- Assign and configure the bridge members that have not yet been handled. Review the steps below to ensure the interface settings are correct even if the interfaces have already been assigned and configured.
  - Navigate to **Interfaces > Assignments**
  - Choose the interface from the **Available network ports** list
  - Click **Add**
  - Navigate to the new interface configuration page, e.g. **Interfaces > OPT2**

    Enable
    Checked

    Description
    `WiredLAN2` or similar

    IPv4 Configuration Type
    None

    IPv6 Configuration Type
    None

    Block private networks
    Unchecked

    Block bogon networks
    Unchecked
  - Click **Save**
  - Click **Apply Changes**
  - Repeat for additional unassigned future bridge members
- Create the new bridge
  - Navigate to **Interfaces > Assignments** on the **Bridges** tab
  - Click **Add** to create a new bridge
  - Enter a **Description**, such as `LAN Bridge`
  - Select all the new bridge members **EXCEPT** the *LAN* interface in the **Member interfaces** list
  - Click Save
- Change the bridge filtering **System Tunable** to disable member interface filtering
  - Navigate to **System > Advanced**, **System Tunables** tab
  - Locate the entry for **net.link.bridge.pfil_member** or create a new entry if one does not exist, using that name for the **Tunable**

  - Click  to edit an existing entry
  - Enter `0` in the **Value** field
  - Click **Save**
- Navigate to **Interfaces > Assignments**
- Change the assignment of **LAN** to `bridge0`
- Click **Save**
- Assign and configure the old LAN interface as described previously, setting its IP configuration types to *None* and naming it `WiredLAN`
- Edit the bridge and select the newly assigned `WiredLAN` as a bridge member
- Change the bridge filtering **System Tunable** to enable bridge interface filtering
  - Use the procedure described previously, but set **net.link.bridge.pfil_bridge** to 1

Now the former LAN interface, along with the new bridge members, are all on a common layer 2 with the bridge assigned as LAN along with the other configuration.

Quickest but Most Difficult: Hand Edit config.xml

Hand editing `config.xml` can be very fast for those familiar with the configuration format in XML. This method is easy to get wrong, however, so be sure to have backups and install media nearby in case a mistake is made.

When hand editing `config.xml` to accomplish this task, do as follows:

- Assign the additional bridge members and set their IP configuration types to *None*
- Create the bridge, including *LAN* and *LAN2* and other bridge members
- Assign the bridge (e.g. as OPT2) and enable it, also with an IP configuration type of *None*
- Download a backup of `config.xml` from **Diagnostics > Backup/Restore**
- Open `config.xml` in a text editor that understands UNIX line endings
- Change the *LAN* assignment to `bridge0`
- Change the former *LAN* assignment to what used to be the bridge (e.g. *OPT2*)
- Edit the bridge definition to refer to *OPT2* and not *LAN*
- Save the changes
- Restore the edited `config.xml` from **Diagnostics > Backup/Restore**

The firewall will reboot with the desired setup. Monitor the console to ensure the settings were applied correctly, and no errors are encountered during the boot sequence.

17.3.2 Assigned Bridge MAC Addresses and Windows

The MAC address for a bridge is determined randomly when the bridge is created, either at boot time or when a new bridge is created. That means that on each reboot, the MAC address can change. In many cases this does not matter, but Windows Vista, 7, 8, and 10 use the MAC address of the gateway to determine if they are on a specific network. If the MAC changes, the network identity will change and its status as public, private, etc. may need to be corrected. To work around this, enter a MAC address on the assigned bridge interface to spoof it. Then clients will always see the same MAC for the gateway IP address.

17.4 Bridging and Firewall Rules

Filtering with bridged interfaces functions similar to routed interfaces, but there are some configuration choices to alter exactly how the filtering behaves.

17.4.1 Apply Firewall Rules on Bridges or Interfaces

By default, firewall rules are applied on each member interface of the bridge on an inbound basis, like any other routed interface.

It is possible to decide whether the filtering happens on the bridge member interfaces, or on the bridge interface itself. This is controlled by two *System Tunables* values on **System > Advanced** on the **System Tunables** tab, as seen in Figure *Bridge Filtering Tunables*. The `net.link.bridge.pfil_member` tunable controls whether the firewall will honor rules on the bridge member interfaces. By default, this is on (`1`). The `net.link.bridge.pfil_bridge` tunable controls whether the firewall will honor rules on the bridge interface itself. By default, this is off (`0`). At least one of these must be set to `1`.

net.link.bridge.pfil_member	Packet filter on the member interface	1	
net.link.bridge.pfil_bridge	Packet filter on the bridge interface	0	

Fig. 1: Bridge Filtering Tunables

When filtering on the bridge interface itself, traffic will hit the rules as it enters from any member interface. The rules are still considered "inbound" like any other interface rules, but they work more like an interface group since the same rules apply to each member interface.

17.4.2 Firewall Rule Macros

Only one interface of a bridge will have an IP address set, the others will have none. For these interfaces, their firewall macros such as *OPT1 address* and *OPT1 net* are undefined because the interface has no address and thus no subnet.

If filtering is performed on bridge members, keep this fact in mind when crafting rules and explicitly list the subnet or use the macros for the interface where the IP address resides.

17.4.3 Ethernet Rules on Bridge Interfaces

Applying *Ethernet (Layer 2) Rules* on a bridge requires changing a *System Tunables* value at **System > Advanced** on the **System Tunables** tab.

The `net.link.bridge.ipfw` tunable controls whether the firewall will honor Ethernet rules on a bridge interface itself. Though the tunable name mentions IPFW, it controls all link-level packet filtering hooks on bridges.

By default, this tunable is off (`0`). To enable Ethernet rule filtering on bridge interfaces, this must be set to `1`.

17.5 Bridging Two Internal Networks

When bridging two internal networks as described in *Internal Bridges* there are some special considerations to take for certain services on the firewall.

Note: There are additional requirements and restrictions when bridging wireless interfaces because of the way 802.11 functions. See *Bridging and wireless* for more information.

17.5.1 DHCP and Internal Bridges

When bridging one internal network to another, two things need to be done. First, ensure that DHCP is only running on the interface containing the IP address and not the bridge members without an address. Second, an additional firewall rule may be necessary at the top of the rules on the member interfaces to allow DHCP traffic.

Note: This only applies to filtering being performed on member interfaces, not filtering performed on the bridge.

When creating a rule to allow traffic on an interface, normally the source is specified similar to *OPT1 Subnet* so that only traffic from that subnet is allowed out of that segment. With DHCP, that is not enough. Because a client does not yet have an IP address, a DHCP request is performed as a broadcast. To accommodate these requests, create a rule on the bridge member interfaces with the following settings:

- Navigate to **Firewall > Rules** on the tab for the bridge member
- Click  **Add** to add a new rule to the top of the list

  Protocol
  UDP

  Source
  `0.0.0.0`

  Source Port
  `68`

  Destination
  `255.255.255.255`

  Destination port
  `67`

  Description
  `Allow DHCP`

- Click **Save** and **Apply Changes**

The rule will look like Figure *Firewall rule to allow DHCP*.

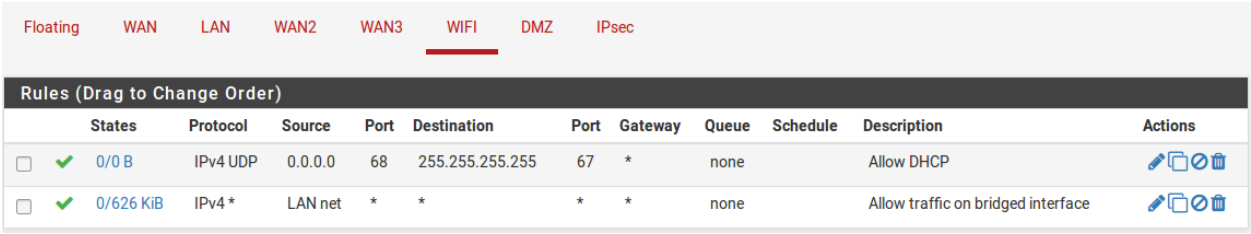


Fig. 2: Firewall rule to allow DHCP

After adding the rule, clients in the bridged segment will be able to successfully make requests to the DHCP daemon listening on the interface to which it is bridged.

DHCPv6 is a bit more complicated to allow since it communicates to and from both link-local and multicast IPv6 addresses. See Figure *Firewall Rule to Allow both DHCP and DHCPv6* for the list of required rules. These can be simplified with aliases into one or two rules containing the proper source network, destination network, and ports.

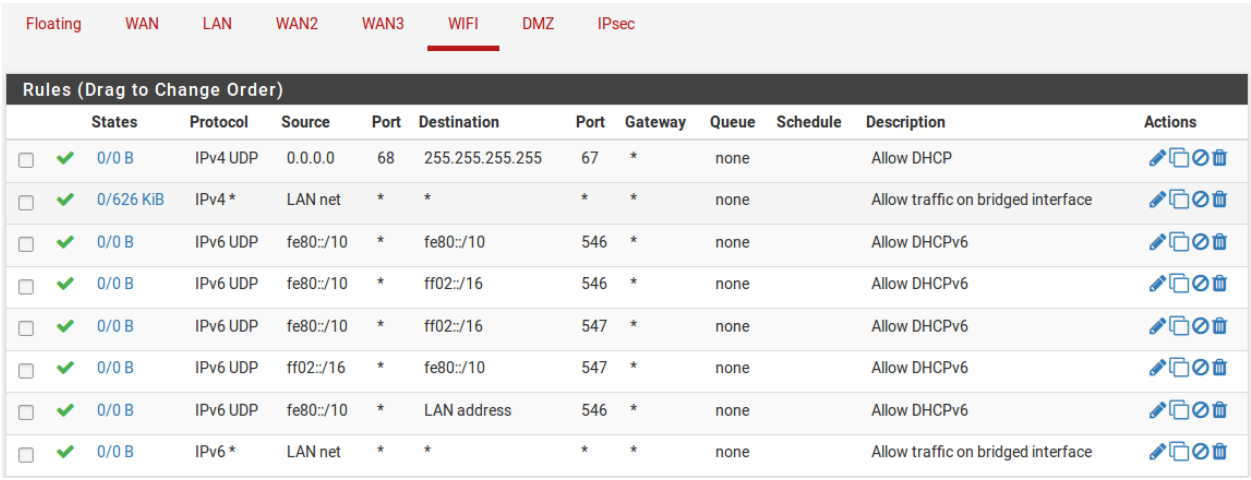


Fig. 3: Firewall Rule to Allow both DHCP and DHCPv6

17.6 Bridging interoperability

Bridged interfaces are different from normal interfaces, thus there are a few features that are incompatible with bridging and others where additional considerations are necessary to accommodate bridging. This section covers features that work differently with bridging than with non-bridged interfaces.

17.6.1 Captive portal

Captive portal (*Captive Portal*) is not compatible with transparent bridging because it requires an IP address on the interface being bridged, used to serve the portal contents, and that IP address must be the gateway for clients. This means that it is not possible, for example, to bridge LAN and WAN and hope to capture clients with the portal.

This can work when bridging multiple *local* interfaces to all route through pfSense® (e.g. LAN1, LAN2, LAN3, etc.). It will work if the bridge interface is assigned, the bridge interface has an IP address, and that IP address is used as the gateway by clients on the bridge. See *Swapping Interface Assignments* for a procedure to place the IP address on an assigned bridge interface.

17.6.2 High Availability

High availability (*High Availability*) is not recommended with bridging. Some users have had mixed success with combining the two in the past, but great care must be taken to handle layer 2 loops, which are unavoidable in a HA+bridge scenario. When two network segments are bridged, they are in effect merged into one larger network, as discussed earlier in this chapter. When HA is added into the mix, that means there will be two paths between the switches for each respective interface, creating a loop.

Managed switches can handle this with Spanning Tree Protocol (STP) but unmanaged switches have no defenses against looping. Left unchecked, a loop will bring a network to its knees and make it impossible to pass any traffic. STP may be configured on bridges to help, though there may still be unexpected results.

Consult switch documentation for information on STP configuration. Configure the switch to give preference to the port(s) on the primary node.

17.6.3 Multi-WAN

Transparent bridging by its nature is incompatible with multi-WAN in many of its uses. When using bridging between a WAN and LAN/OPT interface, commonly something other than pfSense will be the default gateway for the hosts on the bridged interface, and that router is the only device that can direct traffic from those hosts. This doesn't prevent multi-WAN from being used with other interfaces on the same firewall that are not bridged, it only impacts the hosts on bridged interfaces where they use something other than pfSense as their default gateway. If multiple internal interfaces are bridged together and pfSense is the default gateway for the hosts on the bridged interfaces, then multi-WAN can be used the same as with non-bridged interfaces.

17.6.4 Limiters

For limiters to function with bridging, the bridge itself must be assigned, and the bridge interface must have the IP address and not a member interface.

17.6.5 LAN NAT

For port forwards on LAN to function in a bridge scenario, the situation is the same as Captive Portal: It will only function for LAN bridges and not WAN/LAN bridges, the IP address must be on the assigned bridge interface, and that IP address must be used as the gateway for local clients.

17.6.6 Mixing Bridged and NAT Segments

For hosts behind the NAT/routed segment, NAT must occur as traffic exits toward the bridged systems so that the return traffic will come back to the firewall.

For hosts on the bridged segment to reach hosts behind the NAT segment directly, a static route could be used on the bridged hosts or upstream gateway to send the "private" subnet traffic to the IP address of the firewall in the bridged network.

17.6.7 VLANs

A VLAN cannot use a bridge as a parent interface.

Normally each interface on the pfSense® firewall represents its own broadcast domain with a unique IP subnet. In some circumstances it is desirable or necessary to combine multiple interfaces onto a single broadcast domain, where two ports on the firewall will act as if they are on the same switch, except traffic between the interfaces can be controlled with firewall rules. Typically, this is done so multiple interfaces will act as though they are on the same flat network using the same IP subnet and so that clients all share broadcast and multicast traffic.

Certain applications and devices rely on broadcasts to function, but these are found more commonly in home environments than corporate environments. For a practical discussion, see *Bridging and wireless*.

For services running on the firewall, bridging can be problematic. Features such as limiters, Captive Portal, and transparent proxies require special configuration and handling to work on bridged networks. Specifically, the bridge itself must be assigned and the only interface on the bridge with an IP address must be the assigned bridge. Also, in order for these functions to work, the IP address on the bridge must be the address used by clients as their gateway. These issues are discussed more in-depth in *Bridging interoperability*.

17.7 Types of Bridges

There are two distinct types of bridges: Internal bridges and Internal/external bridges. Internal bridges connect two local interfaces such as two LAN interfaces or a LAN interface and a wireless interface. Internal/external bridges connect a LAN to a WAN resulting in what is commonly called a "transparent firewall".

17.7.1 Internal Bridges

With an internal type bridge, ports on the firewall are linked such that they behave similar to switch ports, though with the ability to filter traffic on the ports or bridge and with much lower performance than a switch. The firewall itself is still visible to the local connected clients and acts as their gateway, and perhaps DNS and DHCP server. Clients on the bridged segments may not even know there is a firewall between them.

This type of configuration is commonly chosen by administrators to isolate and control a portion of the network, such as a wireless segment, or to make use of additional ports on the firewall in lieu of a proper switch where installing a switch would be impractical. Though it is not recommended, this type of bridge can also be used to join two remote networks over certain types of VPN connections.

See also:

The Hangouts Archive contains a video which includes practical examples of internal bridges: May 2015 Hangout on Wireless Access Points.

17.7.2 Internal/External Bridges

An Internal/External type bridge, also known as a "transparent firewall", is used to insert a firewall between two segments without altering the other devices. Most commonly this is used to bridge a WAN to an internal network so that the WAN subnet may be used "inside" the firewall, or internally between local segments as an in-line filter. Another common use is for devices behind the firewall to obtain IP addresses via DHCP from an upstream server on the WAN.

In a transparent firewall configuration the firewall does not receive the traffic directly or act as a gateway, it merely inspects the traffic as it passes through the firewall.

Note: Devices on the internal side of this bridge **must** continue to use the upstream gateway as their own gateway. Do not set any IP address on the firewall as a gateway for devices on a transparent bridge.

NAT is not possible with this style of bridge because NAT requires the traffic to be addressed to the firewall's MAC address directly in order to take effect. Since the firewall is not the gateway, this does not happen. As such, rules to capture traffic such as those used by a transparent proxy do not function.

17.8 Bridging and Layer 2 Loops

When bridging, care must be taken to avoid layer 2 loops, or a switch configuration must be in place that handles loops. A layer 2 loop is when, either directly or indirectly, the switch has a connection back to itself. If a firewall running pfSense has interfaces bridged together, and two interfaces are plugged into the same switch on the same VLAN, a layer 2 loop has been created. Connecting two patch cables between two switches also does this.

Managed switches employ Spanning Tree Protocol (STP) to handle situations like this, because it is often desirable to have multiple links between switches, and the network shouldn't be exposed to complete meltdown by someone plugging one network port into another network port. STP is not enabled by default on all managed switches, and is almost never available with unmanaged switches. Without STP, the result of a layer 2 loop is frames on the network will circle endlessly, and the network will completely cease to function until the loop is removed. Check the switch configuration to ensure the feature is enabled and properly configured.

pfSense enables STP on bridge interfaces to help with loops, but it can still lead to unexpected situations. For instance, one of the bridge ports would shut itself down to stop the loop, which could cause traffic to stop flowing unexpectedly or bypass the firewall entirely.

In a nutshell, bridging has the potential to completely melt down the network unless anyone that plugs devices into the switch is careful.

CHAPTER

EIGHTEEN

VIRTUAL LANS (VLANS)

18.1 Terminology

This section defines the terminology required to successfully deploy VLANs.

Trunking
Trunking refers to a means of carrying multiple VLANs on the same physical switch port. The frames leaving a trunk port are marked with an 802.1Q tag in the header, enabling the connected device to differentiate between multiple VLANs. Trunk ports are used to connect multiple switches, and for connecting any devices that are capable of 802.1Q tagging and require access to multiple VLANs. This is commonly limited to the firewall or router providing connectivity between VLANs, in this case, the firewall running pfSense® software, as well as any connections to other switches containing multiple VLANs.

VLAN Identifier (VLAN ID, VID)
Each VLAN has an identifier number (ID) for distinguishing tagged traffic. This is a number between `1` and `4094`.

> **Warning:** The default VLAN on switches is VLAN `1`, and this VLAN should not be used when deploying VLAN trunking. This is discussed further in *VLANs and Security*.

Aside from avoiding the use of VLAN `1`, VLAN numbers may be chosen at will. Some designs start with VLAN `2` and increment by one until the required number of VLANs is reached. Another common design is to use the third octet in the subnet of the VLAN as the VLAN ID. For example, if the environment contains networks `10.0.10.0/24`, `10.0.20.0/24` and `10.0.30.0/24`, it is logical to use VLANs `10`, `20`, and `30` respectively. Choose a VLAN ID assignment scheme that makes sense for a given network design.

Parent interface
The underlying interface where a VLAN resides is known as its **Parent Interface**. For example, `igb0` or `igc0`. When VLANs are configured on pfSense software, each is assigned a virtual interface. The virtual interface name is crafted by combining the parent interface name plus the VLAN ID, separated by a period (`.`). For example, for VLAN `20` on `igb0`, the interface name is `igb0.20`.

Note: The sole function of the parent interface is, ideally, to be the parent for the defined VLANs and not used directly. In some situations this will work, but can cause difficulties with switch configuration, and it requires use of the default VLAN on the trunk port, which is best to avoid as discussed further in *VLANs and Security*.

Access Port
An access port refers to a switch port providing access to a single VLAN, where the frames are not tagged with an 802.1Q header. Normal client-type devices are connected to access ports, which will comprise the majority of switch ports. Devices on access ports do not need knowledge of VLANs or tagging. They see the network on their port the same as they would a switch without VLANs.

Double tagging (QinQ)
QinQ refers to the double tagging of traffic, using both an outer (802.1ad) and inner (802.1Q) tag. This can be useful in large ISP environments, other very large networks, or networks that must carry multiple VLANs across a link that can only carry a single VLAN tag. Triple tagging is also possible.

pfSense software offers QinQ functionality, though it is not a very commonly used feature. These types of environments generally need the kind of routing power that only a TNSR or a high end ASIC-based router can provide, and QinQ adds a level of complexity that is unnecessary in most environments.

See also:

QinQ Configuration

Private VLAN (PVLAN)
PVLAN, sometimes called **Port Isolation**, refers to capabilities of some switches to segment hosts within a single VLAN. Normally hosts within a single VLAN function the same as hosts on a single switch without VLANs configured. PVLAN provides a means of preventing hosts on a VLAN from talking to any other host on that VLAN, only permitting communication between that host and its default gateway. This isn't directly relevant to pfSense software, but is a common question. Switch functionality such as this is the only way to prevent communication between hosts in the same subnet. Without a function like PVLAN, no network firewall can control traffic within a subnet because it never touches the default gateway.

18.2 VLANs and Security

VLANs are a great way to segment a network and isolate subnetworks, but there are security issues which need to be taken into account when designing and implementing a solution involving VLANs. VLANs are not inherently insecure, but misconfiguration can leave a network vulnerable. There have also been past security problems in switch vendor implementations of VLANs.

18.2.1 Segregating Trust Zones

Because of the possibility of misconfiguration, networks of considerably different trust levels should be on separate physical switches. For example, while the same switch could technically be used with VLANs for all internal networks as well as the network outside the firewalls, that should be avoided as a simple misconfiguration of the switch could lead to unfiltered Internet traffic entering the internal network. At a minimum, use two switches in such scenarios: One for outside the firewall and one inside the firewall. In many environments, DMZ segments are also treated separately, on a third switch in addition to the WAN and LAN switches. In others, the WAN side is on its own switch, while all the networks behind the firewall are on the same switches using VLANs. Which scenario is most appropriate for a given network depends on its specific circumstances, level of risk, and security concerns.

18.2.2 Using the default VLAN 1

Because VLAN 1 is the default ("native") VLAN, it may be used in unexpected ways by the switch. It is similar to using a default-allow policy on firewall rules instead of default deny and selecting what is needed. Using a different VLAN is always better, and ensure that only the ports are selected that must be on that VLAN, to better limit access. Switches will send internal protocols such as STP (Spanning Tree Protocol), VTP (VLAN Trunking Protocol), and CDP (Cisco Discover Protocol) untagged over the native VLAN, where the switches use these protocols. It is generally the best practice to keep that internal traffic isolated from data traffic.

If VLAN 1 must be used, take great care to assign every single port on every switch to a different VLAN except those that must be in VLAN 1, and do not create a management interface for the switch on VLAN 1. The native VLAN of the switch group should also be changed to a different, unused, VLAN. Some switches may not implement any of these workarounds, and so it is typically easier to move data to a different VLAN instead of fussing with making VLAN 1 available. With VLAN ID 2 through 4094 to choose from, it is undoubtedly better to ignore VLAN 1 when designing a new VLAN scheme.

18.2.3 Using a trunk port default VLAN

When VLAN tagged traffic is sent over a trunk on the native VLAN, tags in the packets that match the native VLAN may be stripped by the switch to preserve compatibility with older networks. Worse yet, packets that are double tagged with the native VLAN and a different VLAN will only have the native VLAN tag removed when trunking in this way and when processed later, that traffic can end up on a different VLAN. This is also called "VLAN hopping".

As mentioned in the previous section, any untagged traffic on a trunk port will be assumed to be the native VLAN, which could also overlap with an assigned VLAN interface. Depending on how the switch handles such traffic and how it is seen by pfSense® software, using the interface directly could lead to two interfaces being on the same VLAN.

18.2.4 Limiting access to trunk ports

Because a trunk port can talk to any VLAN in a group of trunked switches, possibly even ones not present on the current switch depending on the switch configurations, it is important to physically secure trunk ports. Also make sure there are no ports configured for trunking that are left unplugged and enabled where someone could hook into one, accidentally or otherwise. Depending on the switch, it may offer dynamic negotiation of trunking. Ensure this functionality is disabled or properly restricted.

18.2.5 Other Issues with Switches

Over the years there have been reports of rare cases where VLAN-based switches have leaked traffic across VLANs while under heavy loads, or if a MAC address of a PC on one VLAN is seen on another VLAN. These issues tend to be in older switches with outdated firmware, or extremely low-quality managed switches. These types of issues were largely resolved many years ago, when such security problems were common. No matter what switch from what brand is used for a network, research to see if it has undergone any kind of security testing, and ensure the latest firmware is loaded on the switch. While these issues are a problem with the switch, and not pfSense, they are part of a network's overall security.

Many of the items here are specific to particular makes and models of switches. Security considerations differ based on the switch being used on a network. Refer to its documentation for recommendations on VLAN security.

18.3 VLAN Configuration

This section covers how to configure VLANs in pfSense® software.

18.3.1 VLAN Configuration Options

The following options are available when creating or editing VLANs in the GUI:

Parent Interface
The physical interface which will carry traffic for this VLAN.

VLAN Tag Type
The VLAN tag protocol identifier (TPID) which specifies the type of VLAN tag in a packet. This can be one of two values:

C-Tag
"Customer" tag, value `0x8100`. This is the default and most common value for typical 802.1Q VLAN interfaces. If there is only one layer of VLAN tagging, C-Tag will nearly always be the correct choice. When a packet has multiple tags, C-Tags are the inner tag.

S-Tag
"Service" tag, value `0x88a8`. S-tags are used for the outer tag when double tagging, such as in *802.1ad QinQ*.

VLAN Tag
The VLAN identifier (VID), which is an integer value within the range `1-4094`.

VLAN Priority
802.1p Priority Code Point (PCP) value to use for all packets on this VLAN, as described in *VLAN Priority*. This value is an integer from `0-7`.

See also:

- *VLAN Priority*
- *VLAN Priority (Match and Set)*

Description
Some text to identify the purpose of the VLAN, such as `DMZ` or `Storage`.

18.3.2 GUI VLAN Configuration

To add a VLAN in the GUI:

- Navigate to **Interfaces > Assignments**
- Click the **VLANs** tab
- Click  **Add** to add a new VLAN
- Configure the VLAN as described in *VLAN Configuration Options*
- Click **Save**

Afterward, assign the new VLAN as an interface and configure as any other interface.

See also:

Refer to *Interface Configuration* for details on interface assignment and configuration.

18.3.3 GUI VLAN Configuration Example

In this example, WAN and LAN are assigned as `igc3` and `igc2` respectively. There is also an `igc1` interface that will be used as the VLAN parent interface.

To configure these VLANs in the firewall GUI:

- Navigate to **Interfaces > Assignments** to view the interface list
- Click the **VLANs** tab
- Click  **Add** to add a new VLAN
- Configure the VLAN as shown in Figure *Edit VLAN*:

  Parent Interface
  `igc1`

  VLAN Tag Type
  C-Tag (0x8100)

  VLAN tag
  `10`

  VLAN Priority
  Blank

  Description
  `DMZ`

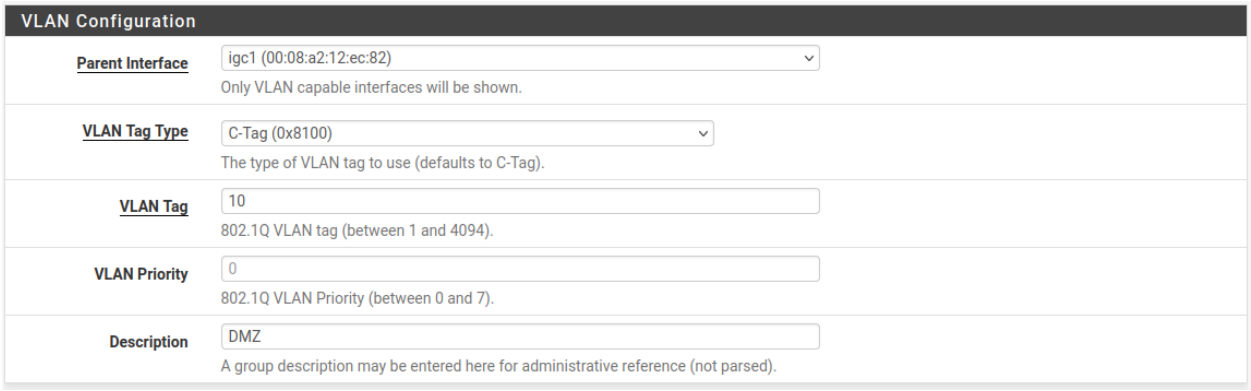


Fig. 1: Edit VLAN

- Click **Save** to return to the VLAN list, which now includes the newly added VLAN `10`
- Repeat the process to add additional VLANs, such as VLAN `20`

  These can be seen in Figure *VLAN list*:

Now assign the VLANs as interfaces:

- Navigate to **Interfaces > Assignments**
- Click the **Interface Assignments** tab

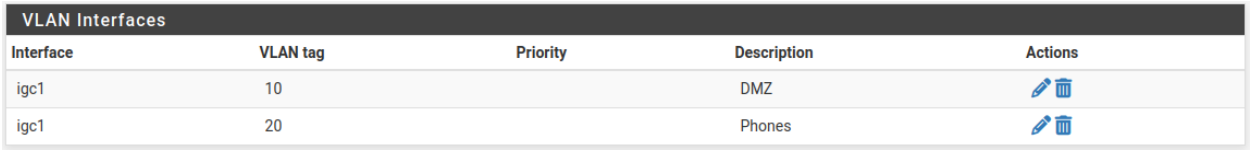


Fig. 2: VLAN list

- Select the VLAN to add from the **Available Network Ports** list, such as *VLAN 10 on igc1 (DMZ)*
- Click **Add** to assign the network port
- Repeat the last two steps to assign *VLAN 20 on igc1 (Phones)*

When finished, the interfaces will look like Figure *Interfaces list with VLANs*

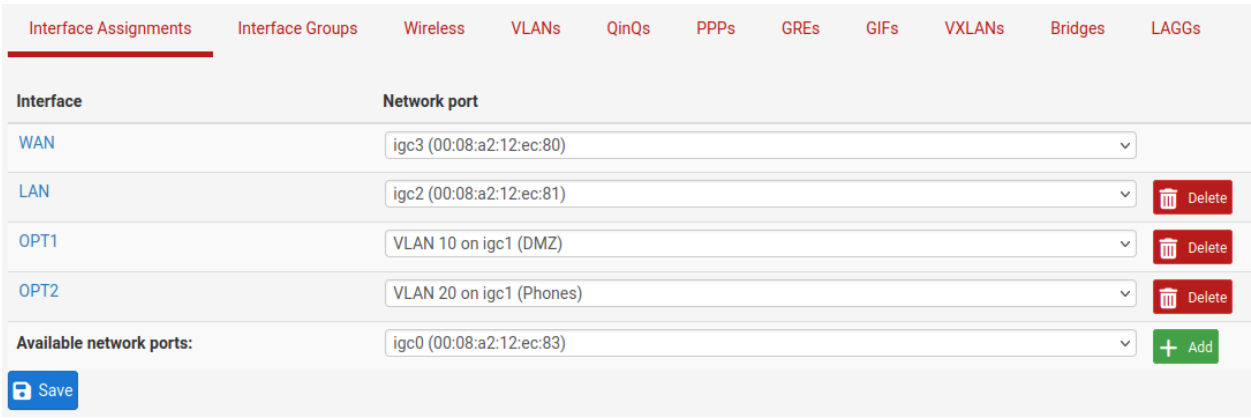


Fig. 3: Interfaces list with VLANs

The VLAN-based OPT interfaces behave as any other OPT interfaces do, which means they must be enabled, configured, have firewall rules added, and services like the DHCP Server will need to be configured if needed. See *Interface Configuration Basics* for more information on configuring optional interfaces.

18.3.4 Console VLAN Configuration Example

VLANs can be configured at the console using the **Assign Interfaces** function, but this method only offers limited functionality compared to the VLAN options available in the GUI.

Tip: Visit the GUI later and edit the VLAN entries to fill in any options not possible during the console VLAN creation process.

See also:

VLAN Configuration Options

The following example shows how to configure two VLANs, ID `10` and `20`, with `igc1` as the parent interface. The VLAN interfaces are assigned as **OPT1** and **OPT2**:

```
 0) Logout (SSH only)                  9) pfTop
 1) Assign Interfaces                 10) Filter Logs
 2) Set interface(s) IP address       11) Restart GUI
 3) Reset admin account and password  12) PHP shell + pfSense tools
 4) Reset to factory defaults         13) Update from console
 5) Reboot system                     14) Disable Secure Shell (sshd)
 6) Halt system                       15) Restore recent configuration
 7) Ping host                         16) Restart PHP-FPM
 8) Shell

Enter an option: 1

Valid interfaces are:

igc0    00:08:a2:xx:xx:83 (down) Intel(R) Ethernet Controller I226-V
igc1    00:08:a2:xx:xx:82 (down) Intel(R) Ethernet Controller I226-V
igc2    00:08:a2:xx:xx:81   (up)
igc3    00:08:a2:xx:xx:80   (up)

Do VLANs need to be set up first?
If VLANs will not be used, or only for optional interfaces, it is typical to
say no here and use the webConfigurator to configure VLANs later, if required.

Should VLANs be set up now [y|n]? y

WARNING: all existing VLANs will be cleared if you proceed!

Do you want to proceed [y|n]? y

VLAN Capable interfaces:

igc0    00:08:a2:xx:xx:83 (down)
igc1    00:08:a2:xx:xx:82 (down)
igc2    00:08:a2:xx:xx:81   (up)
igc3    00:08:a2:xx:xx:80   (up)

Enter the parent interface name for the new VLAN (or nothing if finished): igc1
Enter the VLAN tag (1-4094): 10

VLAN Capable interfaces:

igc0    00:08:a2:xx:xx:83 (down)
igc1    00:08:a2:xx:xx:82 (down)
igc2    00:08:a2:xx:xx:81   (up)
igc3    00:08:a2:xx:xx:80   (up)

Enter the parent interface name for the new VLAN (or nothing if finished): igc1
Enter the VLAN tag (1-4094): 20

VLAN Capable interfaces:

igc0    00:08:a2:xx:xx:83 (down)
igc1    00:08:a2:xx:xx:82 (down)
```

(continues on next page)

(continued from previous page)

```
igc2    00:08:a2:xx:xx:81   (up)
igc3    00:08:a2:xx:xx:80   (up)

Enter the parent interface name for the new VLAN (or nothing if finished): <enter>

VLAN interfaces:

igc1.10     VLAN tag 10, parent interface igc1
igc1.20     VLAN tag 20, parent interface igc1

If the names of the interfaces are not known, auto-detection can
be used instead. To use auto-detection, please disconnect all
interfaces before pressing 'a' to begin the process.

Enter the WAN interface name or 'a' for auto-detection
(igc0 igc1 igc2 igc3 igc1.10 igc1.20 or a): igc3

Enter the LAN interface name or 'a' for auto-detection
NOTE: this enables full Firewall/NAT mode.
(igc0 igc1 igc2 igc1.10 igc1.20 a or nothing if finished): igc2

Enter the Optional 1 interface name or 'a' for auto-detection
(igc0 igc1 igc1.10 igc1.20 a or nothing if finished): igc1.10

Enter the Optional 2 interface name or 'a' for auto-detection
(igc0 igc1 igc1.20 a or nothing if finished): igc1.20

Enter the Optional 3 interface name or 'a' for auto-detection
(igc0 igc1 a or nothing if finished):<enter>

The interfaces will be assigned as follows:

WAN  -> igc3
LAN  -> igc2
OPT1 -> igc1.10
OPT2 -> igc1.20

Do you want to proceed [y|n]? y

Writing configuration...done.
One moment while the settings are reloading... done!
```

After a few seconds, the firewall settings will reload and the console menu will reload.

18.4 VLAN Priority

802.1p, also known as IEEE P802.1p or Priority Code Point, is a method of tagging packets at layer 2 with a specific quality of service priority.

pfSense software can set VLAN Priorities in VLAN configuration (*VLAN Configuration Options*) and set or match them in firewall rules (*VLAN Priority (Match and Set)*).

There are eight levels of priority for 802.1p, and each has an associated two-letter code. In order from the lowest priority to the highest, they are:

PCP 0, BK
Background

PCP 1, BE
Best Effort – This is the default priority when no PCP is specified.

PCP 2, EE
Excellent Effort

PCP 3, CA
Critical Applications

PCP 4, VI
Video

PCP 5, VO
Voice

PCP 6, IC
Internetwork Control

PCP 7, NC
Network Control

See also:

- *Configuring Switches with VLANs*
- *QinQ Configuration*

VLANs enable a switch to carry multiple discrete broadcast domains, allowing a single switch to function as if it were multiple switches. VLANs are commonly used for network segmentation in the same way that multiple switches can be used: To place hosts on a specific segment, isolated from other segments. Where trunking is employed between switches, devices on the same segment need not reside on the same switch. Devices that implement trunking can also communicate on multiple VLANs through a single physical port.

This chapter covers VLAN concepts, terminology and configuration in pfSense® software.

18.5 Requirements

There are two requirements, both of which must be met to deploy VLANs.

1. **802.1Q VLAN capable switch**

   Every decent managed switch manufactured in the last 25 years is capable of 802.1Q VLAN trunking.

   Warning: VLANs **cannot** be used with an unmanaged switch.

2. **Network adapter capable of VLAN tagging**

   A NIC that implements hardware VLAN tagging or long frames is helpful, but not required. Interfaces without appropriate hardware can still use VLANs, but not as fast or as efficiently.

   Each VLAN frame has a 4 byte 802.1Q tag added in the header, so the frame size can be up to 1522 bytes. A NIC with hardware features for VLAN tagging or long frames is helpful because other adapters will not be able to send frames larger than the normal 1518 byte maximum with 1500 MTU Ethernet. This means interfaces without those hardware features must use a lower MTU, and thus are limited in how much data they can send in a packet, reducing efficiency.

   Note: If an adapter is listed as having long frame functionality, that does not guarantee the specific implementation of that NIC chipset properly implements long frames. Realtek `rl(4)` NICs are the biggest offenders. Many will work fine, but some do not properly implement long frames, and some will not accept 802.1Q tagged frames at all. If problems are encountered using one of the NICs listed as having long frame functionality, the best practice is to try an interface with VLAN hardware tagging instead. There are no known similar problems with NICs listed as having VLAN hardware tagging.

Ethernet interfaces with VLAN hardware functionality:

`ae(4)`, `age(4)`, `alc(4)`, `ale(4)`, `bce(4)`, `bge(4)`, `bxe(4)`, `cxgb(4)`, `cxgbe(4)`, `em(4)`, `igb(4)`, `ix(4)`, `jme(4)`, `liquidio(4)`, `msk(4)`, `mxge(4)`, `nge(4)`, `re(4)`, `sge(4)`, `stge(4)`, `ti(4)`, `vge(4)`.

Ethernet interfaces with long frame functionality :

`axe(4)`, `bfe(4)`, `cas(4)`, `dc(4)`, `et(4)`, `fwe(4)`, `fxp(4)`, `gem(4)`, `le(4)`, `nfe(4)`, `rl(4)`, `sis(4)`, `sk(4)`, `ste(4)`, `vr(4)`, `vte(4)`, `xl(4)`.

CHAPTER

NINETEEN

MULTIPLE WAN CONNECTIONS

19.1 Multi-WAN Terminology and Concepts

This section covers terminology and concepts necessary to understand the deployment of multi-WAN functionality with pfSense® software.

See also:

WAN vs LAN Interfaces

19.1.1 Policy routing

Policy routing refers to a means of routing traffic by matching a policy, typically involving firewall rules or access control lists. This type of routing can consider more factors than the destination IP address of the traffic, as is done with the routing table in most operating systems and routers.

In pfSense software the **Gateway** field available when editing or adding firewall rules enables the use of policy routing. The **Gateway** field contains all gateways defined on the firewall under **System > Routing**, plus any gateway groups.

Policy routing provides a powerful means of directing traffic to use an appropriate path, since anything a firewall rule can match may be used as criteria for making policy routing decisions. Specific hosts, subnets, protocols, ports, and more can be used to direct traffic.

Note: Remember on per-interface rule tabs that all firewall rules, including policy routing rules, are processed in top-down order, and the first match wins.

19.1.2 Gateway Groups

Gateway groups define how a chosen set of gateways provide failover and/or load balancing functionality. They are configured under **System > Routing**, on the **Gateway Groups** tab.

See also:

See *Gateway Groups* for more.

19.1.3 Failover

Failover refers to the ability to switch from one or more WANs to an alternate set of WANs if the preferred connections fail. This is useful for situations where traffic should utilize a specific WAN connection unless it is unavailable.

See also:

To fail from one *firewall* to another, rather than from one *WAN* to another, see *High Availability*.

19.1.4 Load Balancing

The **Load Balancing** functionality in pfSense software distributes **connections** over multiple WAN connections in a round-robin fashion. This feature operates on a **per-connection** basis, not a per-packet basis. If a gateway that is part of a load balancing group fails, the interface is marked as down and removed from all groups until it recovers, thus a load balanced configuration effectively also includes failover functionality.

19.1.5 Monitor IP Addresses

When configuring failover or load balancing, each gateway is associated with a monitor IP address (*Gateway Settings*). In a typical configuration, the firewall will ping this IP address and if it stops responding, the gateway is marked as down. Options on the gateway group can select different failure triggers besides packet loss. The other triggers are high latency, a combination of either packet loss or high latency, or when the circuit is down.

What constitutes failure?

The topic is a little more complex than "if pings to the monitor IP address fail, the gateway is marked as down." The actual criteria for a failure depend on the options chosen when creating the gateway group and the individual settings on a gateway.

The settings for each gateway that control when it is considered up and down are all discussed in *Advanced Gateway Settings*. The thresholds for packet loss, latency, downtime, and even the probing interval of the gateway are all individually configurable.

19.1.6 State Killing/Forced Switch

When a gateway has failed, the firewall can optionally flush states to force clients to reconnect, and in doing so they will use a gateway that is online instead of a gateway that is down. This can be done for the entire state table or selectively for only gateways that are down. When clearing states for a specific gateway, it can only clear states created by policy routing rules.

This currently only works one-way, meaning that it can move connections off of a failing gateway, but it cannot force them back if the original gateway comes back online.

This is an optional behavior, and it is **not** enabled by default as it can be disruptive. For information on changing this setting, see *State Killing on Gateway Failure*.

19.1.7 Default Gateway Switching

Traffic exiting the firewall itself will use the default gateway unless a static route sends the packet along a different path. If the default gateway is on a WAN that is down, daemons on the firewall will be unable to make outbound connections, depending on the capabilities of the daemon and its configuration.

The default gateway for the firewall can be set to a gateway group or set to an automatic mode, which will switch the default to the next available gateway if the normal default gateway fails, and then switched back when that WAN recovers. See *Managing the Default Gateway* for details.

Note: Daemons bound to non-default WANs which have no static routes influencing their outbound traffic may also fail in certain cases even when all WANs are online. See *Configuring the Firewall Default State Policy* for details.

19.2 Policy Routing, Load Balancing and Failover Strategies

This section provides guidance on common multi-WAN goals and how they can be achieved with pfSense® software.

19.2.1 Bandwidth Aggregation

One of the primary desires with multi-WAN is bandwidth aggregation. Load balancing can help accomplish this goal. There is, however, one caveat: If the firewall has two 50 Mbit/s WAN circuits, it cannot get 100 Mbit/s of throughput with a *single* client connection. Each individual connection must be tied to only one specific WAN. This is true of any multi-WAN solution other than MLPPP. The bandwidth of two different Internet connections cannot be aggregated into a single large "pipe" without involvement from the ISP. With load balancing, since individual connections are balanced in a round-robin fashion, 100 Mbit/s of throughput can only be achieved using two 50 Mbit/s circuits when multiple connections are involved. Applications that utilize multiple connections, such as download accelerators, will be able to achieve the combined throughput capacity of the two or more connections.

Note: Multi-Link PPPoE (MLPPP) is the only WAN type which can achieve full aggregate bandwidth of all circuits in a bundle, but MLPPP requires special implementation from the ISP. For more on MLPPP, see *Multi-Link PPPoE (MLPPP)*.

In networks with numerous internal machines accessing the Internet, load balancing will reach speeds near the aggregate throughput by balancing the many internal connections out all the WAN interfaces.

19.2.2 Segregation of Priority Services

Consider a site which has a reliable, high quality Internet connection that offers low bandwidth, or high costs for excessive transfers, and another connection that is fast but of lesser quality (higher latency, more jitter, or less reliable). In these situations, services can be segregated between the two Internet connections by their priority. High priority services may include VoIP, traffic destined to a specific network such as an outsourced application provider, or specific protocols used by critical applications, amongst other options. Low priority traffic commonly includes any permitted traffic that doesn't match the list of high priority traffic. Policy routing rules can be setup to direct the high priority traffic out the high quality Internet connection, and the lower priority traffic out the lesser quality connection.

Another example of a similar scenario is getting a dedicated Internet connection for quality critical services such as VoIP, and only using that connection for those services.

19.2.3 Failover Only

There are scenarios where the best practice is to only use failover. For example, users who have a secondary backup Internet connection with a low bandwidth cap such as a 4G/LTE modem, and only want to use that connection if their primary connection fails. Gateway groups configured for failover can achieve this goal.

Another usage for failover is to ensure a certain protocol or destination always uses only one WAN unless it goes down.

19.2.4 Unequal Cost Load Balancing

pfSense software can achieve unequal cost load balancing by setting appropriate weights on gateways as discussed in *Advanced Gateway Settings*. By setting a weight on a gateway, it will be used more often in a gateway group. Weights can be set from *1* to *30*, allowing

Table 1: Unequal Cost Load Balancing

WAN_GW weight	WAN2_GW weight	WAN load	WAN2 load
3	2	60%	40%
2	1	67%	33%
3	1	75%	25%
4	1	80%	20%
5	1	83%	17%
30	1	97%	3%

Note that this distribution is strictly balancing the number of *connections*, it does not take interface throughput or existing load into account. This means bandwidth usage may not necessarily be distributed equally, though in most environments it works out to be roughly distributed as configured over time. This also means if an interface is loaded to its capacity with a single high throughput connection, the firewall will still direct additional connections to that interface.

19.3 Multi-WAN Caveats and Considerations

This section contains the caveats and considerations specific to multi-WAN in pfSense® software.

19.3.1 Multiple WANs sharing a single gateway IP

Due to the way PF handles multi-WAN connections, traffic can only be directed using the gateway IP address of a circuit, which is fine for most scenarios. If the firewall has multiple connections on the same ISP using the same subnet and gateway IP address, as is common when using multiple cable modems, an intermediate NAT device must be used on all but one of them so that the firewall sees each WAN gateway as a unique IP address.

When using the NAT device it can be configured to forward all traffic back to the firewall which can help with using that WAN for other services. However, some protocols, such as VoIP, will have problems if they use a WAN with NAT in such a configuration.

If at all possible, contact the ISP and have them configure the WAN circuits such that they are in different subnets with different gateways.

One exception to this is a PPP type WAN such as PPPoE. PPP type WANs are capable of having the same gateway on multiple interfaces, but each gateway entry must be configured to use a different monitor IP address (See *Gateway Settings*).

19.3.2 Multiple PPPoE WANs

When multiple PPPoE lines from the same ISP are present and the ISP offers Multi-Link PPPoE (MLPPP), it may be possible to bond the lines into a single aggregate link. This bonded link has total bandwidth of all lines together in a single WAN as seen by the firewall. Configuration of MLPPP is covered in *Multi-Link PPPoE (MLPPP)*.

19.3.3 Local Services and Multi-WAN

There are additional considerations with local services and multi-WAN, since any traffic initiated from the firewall itself will not be affected by policy routing configured on internal interface rules. Traffic from the firewall itself always follows the routing table. Hence, static routes are required under some circumstances when using additional WAN interfaces, otherwise only the WAN interface with the default gateway would be used.

The firewall can be configured to change the default gateway if the preferred default fails. See *Managing the Default Gateway* for details.

In the case of traffic initiated on the Internet destined for any WAN interface, pfSense software automatically uses the `reply-to` directive in all WAN-type interface rules, which ensures the reply traffic is routed back out the correct WAN interface.

Note: Daemons bound to non-default WANs which have no static routes influencing their outbound traffic may also fail in certain cases even when all WANs are online. See *Configuring the Firewall Default State Policy* for details.

DNS Resolver

The default settings for the DNS Resolver require using failover for the default gateway to work properly with Multi-WAN. See *Managing the Default Gateway* for details. As an alternative to using default gateway switching, a few changes can be made to make the DNS Resolver more accommodating to Multi-WAN, including enabling forwarding mode. The details are described later in this chapter.

DNS Forwarder

The DNS servers used by the DNS forwarder must have gateways defined if they use a non-default WAN interface, as described later in this chapter. That is the only requirement for using the DNS forwarder in multi-WAN environments.

Dynamic DNS

Dynamic DNS entries can be set using a gateway group for their interface. This will move a Dynamic DNS entry between WANs in failover mode, allowing a public hostname to shift from one WAN to another in case of failure.

IPsec

IPsec is fully compatible with multi-WAN. A static route is automatically added for the remote tunnel peer address pointing to the specified WAN gateway to ensure the firewall sends traffic out the correct path when it initiates a connection. For mobile connections, the client always initiates the connection, and the reply traffic is correctly routed by the state table.

An IPsec tunnel may also be set using a gateway group as its interface for failover. This is discussed further in *Multi-WAN Environments*.

OpenVPN

OpenVPN multi-WAN capabilities are described in *OpenVPN and Multi-WAN*. Like IPsec, it can use any WAN or a gateway group.

CARP and multi-WAN

CARP is multi-WAN capable so long as all WAN interfaces use static IP addresses and there are at least three public IP addresses available per WAN. This is covered in *High Availability Configuration Example with Multi-WAN*.

19.3.4 IPv6 and Multi-WAN

IPv6 is also capable of performing in a multi-WAN capacity, but will usually require Network Prefix Translation (NPt) on one or more WANs. This is covered in more detail in *Configuring Multi-WAN for IPv6*.

19.4 Summary of Multi-WAN Requirements

This is a brief summary of configuration changes necessary for a fully implemented multi-WAN setup:

- Create a gateway group under **System > Routing** on the **Gateway Groups** tab
- Set a failover gateway group for the default gateway as described in *Managing the Default Gateway* (Technically optional but a best practice)
- Configure the DNS Resolver or Forwarder for Multi-WAN, starting by:
  - Use a failover gateway group for the default gateway (DNS Resolver in default resolver mode)
  - Set at least one unique DNS server for each WAN gateway under **System > General Setup** with a gateway set (DNS Resolver in forwarding mode or DNS Forwarder)
- Use the gateway group on LAN firewall rules

These topics are covered in detail by the following sections.

19.5 Load Balancing and Failover with Gateway Groups

Gateway Groups are necessary components of a Load Balancing or Failover configuration. The group itself does not cause any action to be taken, but when the group is used later, such as in policy routing firewall rules, the group settings define how the items utilizing the group will behave.

The same gateway may be included in multiple groups so that several different scenarios can be configured at the same time. For example, some traffic can be load balanced, and other traffic can use failover, and the same WAN can be used in both capacities by using different gateway groups.

A common example setup for a two WAN firewall contains three groups:

LoadBalance
: Gateways for WAN1 and WAN2 both on Tier 1

PreferWAN1
: Gateway for WAN1 on Tier 1, and WAN2 on Tier 2

PreferWAN2
: Gateway for WAN1 on Tier 2, and WAN2 on Tier 1

The best practice for any strategy is to have at least one failover group and to set that failover group to be used as the default gateway on the firewall. This ensures that the firewall always has a viable default gateway, and using a gateway group ensures that the correct gateways are used for this function and in the intended order. See *Managing the Default Gateway* for details.

19.5.1 Configuring a Gateway Group for Load Balancing or Failover

To create a gateway group for Load Balancing or Failover:

- Navigate to **System > Routing**, **Gateway Groups** tab
- Click  **Add** to create a new gateway group
- Fill in the options on the page as described in *Gateway Group Options*
- Click **Save**

Load Balancing

Any two gateways on the same tier are load balanced. For example, if *Gateway A*, *Gateway B*, and *Gateway C* are all Tier 1, connections would be balanced between them. Gateways that are load balanced will automatically fail over between each other. When a gateway fails it is removed from the group, so in this case if any one of A, B, or C went down, the firewall would load balance between the remaining online gateways.

Weighted Balancing

If two WANs must be balanced in a weighted fashion due to differing amounts of bandwidth between them, that can be accommodated by adjusting the **Weight** parameter on the gateway as described in *Unequal Cost Load Balancing* and *Advanced Gateway Settings*.

Failover

The firewall prefers gateways on a **lower** number tier. If the gateways on the lowest number tier are down then it looks for gateways on a higher numbered tier.

For example, if *Gateway A* is on Tier 1, *Gateway B* is on Tier 2, and *Gateway C* is on Tier 3, then the firewall uses *Gateway A* first. If *Gateway A* goes down, then the firewall uses *Gateway B*. If both *Gateway A* and *Gateway B* are down, then the firewall uses *Gateway C*.

Complex/Combined Scenarios

By extending the concepts above for Load Balancing and Failover, complicated scenarios are possible that combine both load balancing and failover. For example, if *Gateway A* is on Tier 1, and *Gateway B* and *Gateway C* are on Tier 2, then *Gateway D* on Tier 3, the following behavior occurs: *Gateway A* is preferred on its own. If *Gateway A* is down, then traffic would be load balanced between *Gateway B* and *Gateway C*. Should either *Gateway B* or *Gateway C* go down, the remaining online gateway in that tier would still be used. If *Gateway A*, *Gateway B*, and *Gateway C* are all down, traffic would fail over to *Gateway D*.

Any other combination of the above can be used, so long as it can be arranged within the limit of **5** tiers.

19.5.2 Problems with Load Balancing

Some websites store session information including the client IP address, and if a subsequent connection to that site is routed out a different WAN interface using a different public IP address, the website will not function properly. This has been more common with banks and other security-minded sites. One method of working around this issue is to create a failover group and direct traffic destined to these sites to the failover group rather than a load balancing group. Alternately, perform failover for all HTTPS traffic.

The sticky connections feature of PF is intended to resolve this problem. It is safe to use, and should alleviate this scenario, but there is also a downside to using the sticky option. When using sticky connections, the firewall remembers an association between the *client IP address* and a given *gateway*, it is **not** based off of the destination. When the sticky connections option is enabled, a single client would not load balance its connections between multiple WANs, but it would be associated with whichever gateway it happened to use for its first connection. Once all the client states have expired, the client may exit a different WAN for its next connection, resulting in a new gateway pairing. As such, it works best in environments with many clients where one client utilizing a single WAN does not have a large impact.

19.6 Interface and DNS Configuration

The first two items to configure for Multi-WAN are interfaces and DNS.

19.6.1 Interface Configuration

Set up the primary WAN as previously described in *Setup Wizard*. Then for the additional WAN interfaces, perform the following tasks:

- Assign the interfaces if they do not yet exist
- Visit the **Interfaces** menu entry for each additional WAN (e.g. **Interfaces > OPT1**)
- Enable the interface
- Enter a suitable name, such as `WAN2`
- Select the desired type of IP address configuration depending on the Internet connection type.
- Enter the remaining details for the type of WAN. For example, on static IP connections, fill in the IP address, subnet mask, and add or select a gateway.

19.6.2 DNS Configuration

DNS is critical for Internet connectivity. For multi-WAN to function correctly the firewall must always be able to resolve DNS for itself and on behalf of local clients utilizing the DNS Resolver or DNS Forwarder.

When the firewall configuration only includes DNS servers from a single WAN, an outage of that WAN results in a complete Internet outage since DNS will no longer function.

To ensure DNS resolution uses multiple WANs, consult one of the following sections:

Resolver Mode
: If the DNS Resolver is in resolver mode, see *DNS Resolver and Multi-WAN*.

Forwarding Mode
: If the DNS Resolver is set for **forwarding** mode, or if the DNS Forwarder is in use, then the firewall must be configured with DNS servers for each WAN as described in *DNS Forwarding and Static Routes*.

DNS Resolver and Multi-WAN

The DNS Resolver can work with multi-WAN, but the exact configuration depends on the desired behavior and current settings, especially the chosen *DNS Resolver mode*.

If the DNS Resolver is using its default resolver mode, such as for environments which require DNSSEC, then it can still function with multi-WAN but requires using failover for the default gateway. See *Managing the Default Gateway*.

Even in resolver mode if the firewall is set to use or fall back to remote DNS servers under **System > General**, **DNS Resolution Behavior**, then it is still useful to configure gateways for individual DNS servers as described in *DNS Forwarding and Static Routes*

DNS Forwarding and Static Routes

When using the DNS Resolver in forwarding mode or the DNS Forwarder, the firewall uses its routing table to reach the configured DNS servers. This means without any static routes configured, it will only use the WAN with the default gateway to reach DNS servers.

Gateways must be selected for each DNS server defined on the firewall. This forces the firewall to use a specific WAN interface to reach a given DNS server. At least one gateway from each WAN should be selected where possible.

Note: Most ISPs prohibit recursive queries from hosts outside their network, hence the firewall must use the correct WAN interface when accessing DNS servers for a specific ISP.

DNS servers obtained from a dynamic WAN are automatically routed back out the appropriate dynamic WAN.

To configure DNS server gateways:

- Navigate to **System > General Setup**
- Define at least one *unique* DNS server for each WAN
- Select an appropriate gateway for each DNS server so it uses a specific WAN

  Note: If the gateway entries for these WANs use DNS servers for their monitor IP addresses, ensure there is no conflict between those values and the selected gateways in this list.

Note: Each entry must be unique; the same DNS server cannot be entered more than once.

If using the DNS Resolver, ensure it is set for forwarding mode:

- Navigate to **Services > DNS Resolver**
- Check **Enable Forwarding Mode**
- Uncheck **Enable DNSSEC Support**
- Click **Save**, then **Apply Changes**

19.7 Multi-WAN and NAT

The default NAT rules generated by pfSense® software will translate any traffic leaving a WAN-type interface to the IP address of that interface. In a default two interface LAN and WAN configuration, pfSense software will NAT all traffic from the LAN subnet leaving the WAN interface to the WAN IP address. Adding more WAN-type interfaces extends this to NAT any traffic leaving a WAN-type interface to that interface IP address. This is all handled automatically unless Manual Outbound NAT is enabled.

Warning: NAT does not influence the path taken by connections, only how addresses on packets traversing an interface are translated by the firewall.

Policy routing firewall rules direct connections to specific WAN interfaces, and the Outbound and 1:1 NAT rules specify how the addresses on packets for those connections will be translated by the firewall as it leaves that WAN.

19.7.1 Multi-WAN and Manual Outbound NAT

If **Manual Outbound NAT** must be used with multi-WAN, ensure manual outbound NAT rules are present for all WAN-type interfaces.

19.7.2 Multi-WAN and Port Forwarding

Each port forward applies to a single WAN interface. A given port can be opened on multiple WAN interfaces by using multiple port forward entries, one per WAN interface. The easiest way to accomplish this is:

- Add a port forward on the first WAN connection as usual
- Click  to the right of that entry to add another port forward based on the selected one
- Change the **Interface** to the desired WAN
- Click **Save**

The `reply-to` keyword in PF, which the firewall automatically places on WAN-type interface rules by default, ensures that when traffic comes in over a specific WAN interface, the return traffic will go back out the way it came into the firewall. Thus, port forwards can be actively used on all WAN interfaces at any time, regardless of any failover configuration that may be present. This is especially useful for mail servers as an address on a secondary WAN can be used as a backup MX, allowing the site to receive mail even when the primary line is down.

See also:

This `reply-to` behavior is configurable, for information on this setting, see *Disable Reply-To*.

19.7.3 Multi-WAN and 1:1 NAT

1:1 NAT entries are specific to a single WAN interface and, like outbound NAT, they only control what happens to addresses on packets as they pass through an interface. Internal systems can be configured with a 1:1 NAT entry on each WAN interface, or a 1:1 entry on one or more WAN interfaces and use the default outbound NAT on others. Where 1:1 entries are configured, they always override any other Outbound NAT configuration for that specific interface.

If a local device must always use a 1:1 NAT entry on a specific WAN, then traffic from that device must be forced to use that specific WAN gateway with policy routing firewall rules.

19.8 Policy Routing Configuration

At this point the firewall is prepared for Multi-WAN but not fully configured. With default gateway switching the firewall will have basic failover, but it cannot yet use more advanced failover or load balancing behaviors without policy routing firewall rules in place.

See also:

For information on default gateway switching, see *Managing the Default Gateway*.

19.8.1 Configuring the Firewall Default State Policy

The default State Policy (*Firewall State Policy*) is not directly related to policy routing but can affect how it functions for traffic originating on the firewall itself.

Daemons bound to WANs that are not default, and which have no static route configured to control their outbound behavior, may fail to pass outbound traffic when the default policy is set to "Interface Bound States". This is common for certain types of VPN clients, such as OpenVPN. If these are necessary, consider changing the default policy to "Floating States".

See *Interface Bound States* for additional information on this failure case.

19.8.2 Configuring Firewall Rules for Policy Routing

Setting a **Gateway** on a firewall rule will cause traffic matching the rule to use the chosen gateway or group, following the configured behavior of the group.

The easiest way to configure a firewall for policy routing is to edit the existing default pass rule for the LAN and select the gateway group there. With that set, any traffic matching the default pass rule on the LAN will use the chosen gateway or group.

To make that edit:

- Navigate to **Firewall > Rules**, **LAN** tab
- Click  on the row with the default pass rule
- Click  **Display Advanced** under **Extra Options**
- Select the desired gateway group from the **Gateway** drop-down list
- Click **Save**
- Click **Apply Changes**

Only the most basic of deployments will be satisfied with that configuration, most configurations are more complex. Continue reading for more factors that can require additional configuration.

19.8.3 Bypassing Policy Routing

If there are other local interfaces, VPNs, MPLS interfaces, or traffic that must otherwise obey the routing table, then that traffic must be configured to bypass policy routing. This is simple to do by making a rule to match the traffic in question and then placing that rule **above** any rules that have a gateway configured, because the first rule to match is the one that is used.

This can be generalized by making an alias for any *RFC1918* traffic which would cover all private networks, and then using that in a rule. This alias would contain at least `192.168.0.0/16`, `172.16.0.0/12`, and `10.0.0.0/8`.

In Figure *Bypass Policy Routing Example Rules*, local and VPN traffic bypasses policy routing, HTTPS traffic prefers WAN2, and all other traffic is load balanced:

Rules (Drag to Change Order)

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
✓	0/0 B	*	*	*	LAN Address	443 80 22	*	*		Anti-Lockout Rule	⚙
☐ ✓	0/0 B	IPv4 *	LAN net	*	RFC1918	*	*	none		Bypass policy routing for local/VPN traffic	⚓ ✎ ⧉ ⊘ 🗑
☐ ✓ ⚙	0/0 B	IPv4 TCP	LAN net	*	*	443 (HTTPS)	PreferWAN2	none		Force HTTPS out WAN2, fail to WAN1	⚓ ✎ ⧉ ⊘ 🗑
☐ ✓ ⚙	0/10 KiB	IPv4 *	LAN net	*	*	*	LoadBalance	none		Load Balance LAN Traffic	⚓ ✎ ⧉ ⊘ 🗑

Fig. 1: Bypass Policy Routing Example Rules

19.8.4 Mixing Failover and Load Balancing

As shown in Figure *Bypass Policy Routing Example Rules*, failover and load balancing can be used at the same time by carefully ordering the rules on an interface. Rules are processed from the top down and the first match wins. By placing more specific rules near the top of the list, and the general "match all" style rules at the bottom, any number of different combinations are possible with rules using different gateways or groups.

19.8.5 Enforcing Gateway Use

There are situations where traffic should only ever use one gateway and never load balance or failover. In this example, a device must only exit via a specific WAN and lose all connectivity when that WAN fails.

First, set the **Gateway** on a firewall rule matching traffic from this device to a specific WAN Gateway. If that gateway is down, the rule will act as if the gateway was not set at all, so it needs to be taken a couple steps further.

Add a rule immediately below the rule matching the traffic, but set to reject or block instead. This rule must not have a gateway set.

Next, configure the firewall to omit rules for gateways that are down (*Gateway Monitoring*):

- Navigate to **System > Advanced** on the **Miscellaneous** tab
- Check **Do not create rules when gateway is down**
- Click **Save**

With that option enabled, the first rule will be omitted entirely, falling through to the next matching rule. This way, when the first rule is omitted automatically, traffic will be stopped by the block rule.

19.9 Verifying Functionality

After completing the multi-WAN configuration the best practice is to test the functionality to verify it functions as expected. The following sections describe how to test each portion of a multi-WAN configuration.

19.9.1 Testing Failover

Testing Multi-WAN in a controlled manner immediately after configuration is a key step in the process. Do not make the mistake of waiting until an Internet connection fails naturally for the first test, only to discover problems when they are much more difficult and stressful to fix.

First, navigate to **Status > Gateways** and ensure all WAN gateways show as **Online** under **Status**, as well as on the **Gateway Groups** tab. If they do not, verify that a proper monitor IP address is used as discussed in *Gateway Settings*.

Creating a WAN Failure

There are a number of ways to simulate a WAN failure. For any type, first try unplugging the target WAN interface Ethernet cable from the firewall.

For cable and DSL connections, try powering off the modem/CPE, and in a separate test, unplug the coax or phone line from the modem. For fiber, wireless, and other types of connections with a router outside pfSense® software, try unplugging the Internet connection from the router, and also turning off the router itself.

All the described testing scenarios will likely end with the same result. However, there are some circumstances where trying all these things individually will find a fault that would not have otherwise been noticed until an actual failure. One of the most common is unknowingly using a monitor IP address assigned to the DSL or cable modem. Hence, when the coax or phone line is disconnected, simulating a provider failure rather than an Ethernet or modem failure, the monitor ping still succeeds since it is pinging the modem. From what the firewall was told to monitor, the connection is still up, so it will not fail over even if the upstream connection is actually down. There are other types of failure that can similarly only be detected by testing all the individual possibilities for failure. The monitor IP address can be edited on the gateway entry as covered in *Gateway Settings*.

Verifying Interface Status

After creating a WAN failure, refresh **Status > Gateways** to check the current status. As the gateway monitoring daemon notices the loss, the loss will eventually move past the configured alarm thresholds, and it will mark the gateway as down.

19.9.2 Verifying Load Balancing Functionality

This section describes how to verify the functionality of a load balancing configuration.

Verifying HTTP Load Balancing

The easiest way to verify HTTP load balancing is to visit a website that displays the public IP address the client used to access the site. A page on the Netgate site is available for this purpose, and countless other sites offer the same functionality. Search for "what is my IP address" and numerous websites are returned that will show the public IP address making the HTTP request. Many of those sites tend to be full of advertisements, so Netgate provides a handful of sites which report only the client IP address:

- https://files00.netgate.com/ip
- https://files01.netgate.com/ip
- https://www.pfsense.org/ip

Browsers have a habit of keeping open server connections and caching results, so the best browser-based test is to either load multiple sites, or to close the browser window between attempts to load a site. During each connection attempt, a different IP address should be shown if load balancing is working correctly. If other traffic is present on the network, the IP address may not appear to change on every page load. Repeat the test several times and the IP address should change at least a few times. As an alternative to using a browser, command line utilities such as `curl` do not keep persistent sessions and are more accurate for testing this behavior.

If the IP address never changes, try several different sites, and make sure the browser is requesting the page again, and not returning something from its cache or using a persistent connection to the server. Other good steps include manually deleting the cache, closing and reopening the browser, and trying multiple web browsers. If all of those fail to return the expected result, then start troubleshooting the load balancer configuration further.

Testing load balancing with traceroute

The `traceroute` utility (or `tracert` in Windows) shows the network path taken to a given destination. See *Using traceroute* for details on using `traceroute`. With load balancing, running a `traceroute` from a client host behind the firewall should show a different path being taken for each attempt. Due to the way `traceroute` functions, wait at least one minute after stopping a traceroute before starting another test so that the states will expire, or try different destinations on each attempt.

Using Traffic Graphs

The real time traffic graphs under **Status > Traffic Graph** and on the Traffic Graphs dashboard widget are useful for showing the real time throughput on WAN interfaces. Only one graph at a time can be shown per browser window when using **Status > Traffic Graph**, but additional windows or tabs can be opened in the browser to see all WAN interfaces simultaneously. The traffic graphs widget for the Dashboard enables the simultaneous display of multiple traffic graphs on a single page to simplify this process. If load balancing is working correctly, activity will be observed on all WAN interfaces.

The RRD traffic graphs under **Status > Monitoring** are useful for longer-term and historical evaluation of WAN utilization.

Note: Bandwidth usage may not be exactly equally distributed, since connections are directed on a round-robin basis without regard for bandwidth usage.

19.10 IPsec in Multi-WAN Environments

IPsec on pfSense® software can work well with multiple WAN connections.

19.10.1 Alternate / Non-Default WAN

When using Multi-WAN with IPsec, pick the appropriate **Interface** choice for the WAN-type interface to which the tunnel will connect. If the connection will enter via WAN, pick WAN. If the tunnel will use a different WAN, choose whichever OPT WAN interface is needed. The firewall will automatically add a static route to ensure that the traffic to the **Remote Gateway** uses the appropriate WAN.

19.10.2 Failover with Gateway Groups and Dynamic DNS

IPsec can fail between multiple WANs, but it requires some coordination and relies upon *gateway groups* and *dynamic DNS*. If the first gateway goes down the tunnel will move to the next available WAN in the group. When the first WAN comes back up, the tunnel will be rebuilt there again.

Note: Due to its reliance on DNS, this type of failover can take several minutes to establish a tunnel after failover or recovery.

First, set up a failover type gateway group with only *one* gateway per tier.

Next, choose the failover gateway group from the **Interface** list on the IPsec phase 1 configuration.

Next, set up a new dynamic DNS entry for a hostname using the same gateway group as its interface. There are numerous dynamic DNS providers available for this purpose. The firewall will update the Dynamic DNS entry with the active WAN IP address when a WAN fails or recovers.

On the remote side of the tunnel, set the peer address to be the new dynamic DNS hostname. This peer will track updates to the hostname so that it will know to accept traffic from the newly activated WAN.

Note: If a peer allows multiple remote gateway addresses for a tunnel, and all WAN interfaces on the pfSense software side are static, then those static IP addresses can be used instead of relying on DNS.

19.10.3 Failover with Routed IPsec and Dynamic Routing

In some environments it is possible to use routed IPsec (VTI) to achieve faster Multi-WAN failover.

This method uses one VTI IPsec tunnel per WAN connecting to the same number of WANs at the remote peer. These VTI tunnels are kept up at all times.

Dynamic routing is then setup on all the tunnels using the *FRR Package* to select an active path to the remote endpoint. Depending on the protocols used (e.g. OSPF vs BGP) and settings, failover can happen in seconds instead of minutes.

19.11 Using OpenVPN with Multi-WAN

OpenVPN servers can be used with any WAN, or multiple WANs, as can OpenVPN clients. This document covers only a remote access OpenVPN server, but a similar process could be applied for site to site VPNs. For OpenVPN client instances on pfSense software, in most cases it's as simple as picking the gateway group for the interface.

There are many different ways to configure multiple WANs with OpenVPN on pfSense® software for remote access or site to site VPNs.

See also:

Many of these scenarios were covered during the September 2014. "Advanced OpenVPN Concepts" presentation available through Hangouts Archive.

19.11.1 OpenVPN Configuration

First, get OpenVPN working as desired on the primary WAN interface. Once it is properly functioning, make a backup.

19.11.2 Bind to Localhost and Setup Port Forwards

The OpenVPN configuration needs to be adjusted so it can be reached from either WAN. The simplest way to do this is by changing the **Interface** on the VPN connection to be *Localhost*, and then adding a port forward on each WAN to redirect the OpenVPN port to *Localhost* (127.0.0.1).

For example: If there are two WANs and the OpenVPN server is running on port *1194*, set the **Interface** to *Localhost*, then add two port forwards:

WAN1
: *UDP*, Source any, Destination *WAN1 Address* port `1194`, redirect target `127.0.0.1` port `1194`

WAN2
: *UDP*, Source any, Destination *WAN2 Address* port `1194`, redirect target `127.0.0.1` port `1194`

19.11.3 Configure Clients

Clients may be configured to use the second WAN by adding a second *remote* statement to their configuration, such as:

```
remote x.x.x.x 1194 udp
```

Where `x.x.x.x` is the second WAN IP address or host name.

This process can be automated by using the OpenVPN Client Export package. When exporting a client, in **Host Name Resolution** choose one of:

Automagic Multi-WAN IPs (port forward targets)
: Adds a remote statement for each port forward found targeting the interface binding and port used by this VPN, uses the IP address of each WAN as-is.

Automagic Multi-WAN DDNS Hostnames (port forward targets)
: Like above, but uses the first located Dynamic DNS hostname for a given WAN. If the WAN is a private IP address, this may be the better choice.

19.11.4 More than two WAN connections

The same steps can be repeated to add more WAN connections. Add a port forward to any additional WAN. Clients will need an updated configuration file if another WAN is added later.

19.12 Multi-WAN on a Stick

In the router world, Cisco and others refer to a VLAN router as a "router on a stick" since it can be a functioning router with only one physical network connection. pfSense® software can be configured in this manner as well, using VLANs and a managed switch capable of 802.1q trunking. Most of the deployments running more than 5 WANs use this methodology to limit the number of physical interfaces required on the firewall. In such a deployment, the WAN interfaces all reside on one physical interface on the firewall, with the internal network(s) on additional physical interfaces. Figure *Multi-WAN on a stick* illustrates this type of deployment.

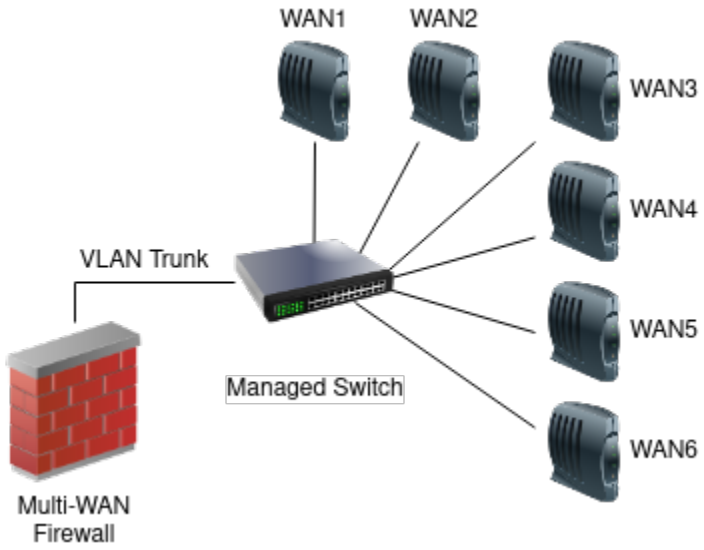


Fig. 2: Multi-WAN on a stick

19.13 Multi-Link PPPoE (MLPPP)

Multi-Link PPPoE (MLPPP) is a unique WAN option that bonds together multiple PPPoE lines from the same ISP to form one larger virtual circuit. This means a firewall can get the true aggregate bandwidth of all circuits in the bundle. For example, if a firewall has three 10 Mbit/s DSL lines in a bundle, it could potentially get 30Mbit/s from a single connection.

19.13.1 Requirements

The largest hurdle for MLPPP is that the ISP must offer the feature on the circuits connected to the firewall. Few ISPs are willing to implement MLPPP, so if an ISP is available that does, it would be worth taking advantage of that fact. Additionally, each line must be on a separate interface connected to the firewall running pfSense® software.

19.13.2 Setup

Setup for MLPPP is simple:

- Configure a WAN for a single line with the correct credentials
- Navigate to **Interfaces > Assign**, **PPPs** tab
- Click  to edit the entry for the PPPoE WAN
- Ctrl-click to select the other physical interfaces that belong to the same MLPPP bundle
- Click **Save**

The firewall will then attempt to bond the lines using MLPPP.

19.13.3 Caveats

One downside to using MLPPP is that troubleshooting is much more difficult. Statistics and status are not available for the individual lines. To determine the status, read through the PPP log, as there is not yet a way to query the lines separately. In some cases it's obvious if a line is down, as there may be a noticeable problem at the modem (out of sync) or that the maximum attainable bandwidth is reduced.

See also:

- *Troubleshooting Multi-WAN*
- *Configuring Multi-WAN for IPv6*

The multiple WAN (multi-WAN) capabilities in pfSense® software allow a firewall to utilize multiple Internet connections to achieve more reliable connectivity and greater throughput capacity.

> **Warning:** Before proceeding with a multi-WAN configuration, the firewall must have a functional two interface (LAN and WAN) configuration.

pfSense software is capable of handling numerous WAN interfaces, with multiple deployments using over 10 WANs in production.

All WAN-type interfaces are treated identically in the GUI. Anything that can be done with the primary WAN can also be done with an additional OPT WAN interface. There are no significant differences between the primary WAN and additional WANs.

This section starts by covering items to consider when implementing *any* multi-WAN solution, then covers multi-WAN configuration with pfSense software.

See also:

For a brief run-down of what to configure when setting up Multi-WAN on pfSense software, see *Summary of Multi-WAN Requirements*.

19.14 Choosing Internet Connectivity

The ideal choice of Internet connectivity depends largely upon the options available at a given location, but there are some additional factors to take into consideration.

19.14.1 Cable Paths

Speaking from the experience of those who have seen first hand the effects of multiple cable-seeking backhoes, as well as nefarious copper thieves, it is highly desirable to obtain connectivity choices for a multi-WAN deployment which utilize disparate cabling paths. In many locations, DSL connections as well as any others utilizing copper pairs are carried on a single cable subject to the same cable cut, and others from the same company such as fiber circuits may run along the same poles or conduits.

If one connection comes in over copper pair (DSL), choose a secondary connection utilizing a different type and path of cabling. Cable connections are typically the most widely available option not subject to the same outage as copper services. Other options include fiber service or fixed wireless coming in on a different path from copper services.

Two connections of the same type cannot be relied upon to provide redundancy in most cases. An ISP outage or cable cut will commonly take down all connections of the same type. Some people use multiple DSL lines or multiple cable modems, though the only redundancy that typically offers is isolating a site from modem or other CPE (Customer Premise Equipment) failure. Consider multiple connections from the same provider as a solution only for additional bandwidth, as the redundancy such a deployment offers is minimal.

19.14.2 Paths to the Internet

Another consideration when selecting Internet connectivity for a site is the path from the connection itself to the Internet. For redundancy purposes, multiple Internet connections from the same provider, especially of the same type, should not be relied upon as they could all fail concurrently.

With larger providers, two different types of connections such as a Fiber line and DSL will usually traverse significantly different networks until reaching core parts of the network. These core network components are generally designed with high redundancy and any problems are addressed quickly since they have widespread effects. Hence, such connectivity is isolated from most ISP issues, but since they commonly utilize the same cable path, it still leaves a site vulnerable to extended outages from cable cuts.

19.14.3 Better Redundancy, More Bandwidth, Less Money

In the past, high-grade services such as DS1 or DS3 circuits were the choice for environments with high availability requirements. Generally the Service Level Agreements (SLA) offered on DS1 and DS3 connections were better than other types of connectivity, and those circuits were generally seen as more reliable. End users have largely left such circuits behind, however, because they are too slow or too costly by today's standards. With the multi-WAN capabilities on pfSense software a site can have more bandwidth and better redundancy for less money in many cases. Fiber services are rapidly becoming more widespread, shaking up this concept by providing extremely large amounts of bandwidth for relatively low cost, though such services may still have a less-than-desirable SLA for outage response.

Most organizations which require high availability Internet connections do not want to rely upon DSL, cable or other "lesser class" broadband Internet connections. While they are usually significantly faster and cheaper, the unfavorable SLA is enough to make many companies think twice. In areas where multiple lower cost broadband options are available, such as fiber and cable, the combination of pfSense software and two low cost Internet connections provides more bandwidth and better redundancy at a lower cost. The chance of two different broadband connections going down simultaneously is significantly less than the chance of any single service outage. Adding a backup Cable or DSL line to supplement a much faster fiber line ensures connectivity will continue when an outage occurs on the fiber line, even if it is a rare occurrence.

CHAPTER

TWENTY

VIRTUAL PRIVATE NETWORKS

20.1 Common deployments

There are four common uses of the VPN capabilities of pfSense, each covered in this section.

20.1.1 Site-to-site connectivity

Site-to-site connectivity is primarily used to connect networks in multiple physical locations where a dedicated, always-on, connection between the locations is required. This is frequently used to connect branch offices to a main office, connect the networks of business partners, or connect a network to another location such as a data center environment.

Before the proliferation of VPN technology, private WAN circuits were the only solution to connect multiple locations. These technologies include point-to-point dedicated circuits, packet switching technologies such as frame relay and ATM, and more recently, MPLS (Multiprotocol Label Switching) and fiber and copper based metropolitan Ethernet services. While these types of private WAN connectivity provide reliable, low latency connections, they are also very costly with recurring monthly fees. VPN technology has grown in popularity because it provides the same secure site to site connectivity using Internet connections that are generally much less costly.

20.1.2 Limitations of VPN connectivity

Performance is an important consideration when planning a VPN solution. In some networks, only a private WAN circuit can meet the requirements for bandwidth or latency. Latency is usually the biggest factor. A point to point DS1 circuit has end to end latency of about 3-5 ms, while the latency to the first hop on an ISP network will generally be at least that much if not higher. Metro Ethernet services or fiber circuits have end to end latency of about 0-3 ms, usually less than the latency to the first hop of an ISP network. That will vary some based on geographical distance between the sites. The stated numbers are typical for sites within a few hundred miles of each other. VPNs usually see latency of around 30-60 ms depending on the Internet connections in use and the geographical distance between the locations. Latency can be minimized, and VPN performance maximized by using the same ISP for all VPN locations, but this isn't always feasible.

Certain protocols perform very poorly with the latency inherent in connections over the Internet. Microsoft file sharing (SMB) is a common example. At sub-10 ms latency, it performs well. At 30 ms or higher, it's sluggish, and at more than 50 ms it's painfully slow, causing frequent hangs when browsing folders, saving files, etc. Getting a simple directory listing requires numerous round trip connections between the client and server, which significantly exacerbates the increased delay of the connection. In Windows Vista and Server 2008, Microsoft introduced SMB 2.0 which includes new capabilities to address the issue described here. SMB 2.0 enables the sending of multiple actions in a single request, as well as the ability to pipeline requests, meaning the client can send additional requests without waiting for the response from prior requests. If a network uses exclusively Vista and Server 2008 or newer operating systems this won't be a concern, but given the rarity of such environments, this will usually be a consideration. SMB 3.0 further improves in this area by allowing multiple streams.

Two more examples of latency sensitive protocols are Microsoft Remote Desktop Protocol (RDP) and Citrix ICA. There is a clear performance and responsiveness difference with these protocols between sub-20 ms response times typically found in a private WAN, and the 50-60+ ms response times common to VPN connections. If remote users work on published desktops using thin client devices, there will be a notable performance difference between a private WAN and VPN. Whether that performance difference is significant enough to justify the expense of a private WAN will vary from one environment to another.

There may be other network applications in an environment that are latency sensitive, where the reduced performance of a VPN is unacceptable. Or all locations may be within a relatively small geographical area using the same ISP, where the performance of a VPN rivals that of private WAN connections.

20.1.3 Remote access

Remote access VPNs enable users to securely connect into a network from any location where an Internet connection is available. This is most frequently used for mobile workers (often referred to as "Road Warriors") whose job requires frequent travel and little time in the office, and to give employees the ability to work from home. It can also allow contractors or vendors temporary access to a network. With the proliferation of smartphones, users have a need to securely access internal services from their phones using a remote access VPN.

20.1.4 Protection for wireless networks

A VPN can provide an additional layer of protection for wireless networks. This protection is two-fold: It provides an additional layer of encryption for traffic traversing the wireless network, and it can be deployed in such a way that it requires additional authentication before access to network resources is permitted. These VPNs are deployed mostly the same as remote access VPNs. Configuring this type of VPN is covered in *Additional protection for a wireless network*.

20.1.5 Secure relay

Remote access VPNs can be configured in a way that passes all traffic from the client system over the VPN. This is nice to have when using untrusted networks, such as wireless hotspots as it lets a client push all its Internet traffic over the VPN and out to the Internet from the VPN server. This protects the user from a number of attacks that are possible on untrusted networks, though it does have a performance impact since it adds additional hops and latency to all connections. That impact is usually minimal with high speed connectivity when the client and VPN server are relatively close geographically.

20.2 Choosing a VPN solution

Each VPN solution has pros and cons. This section will cover the primary considerations in choosing a VPN solution, providing the information necessary to choose the best solution for a given environment.

20.2.1 Interoperability

To interoperate with a firewall or router product from another vendor, the choices are limited by items compatible with both endpoints.

Note: Interoperability in this sense isn't applicable with VPN types not listed here since they are not intended for site-to-site applications.

IPsec

IPsec is usually the best choice since it is included with nearly every VPN-capable device. It also prevents being locked into any particular firewall or VPN solution. For interoperable site-to-site connectivity, IPsec is usually the only choice.

OpenVPN

OpenVPN is interoperable with a few other packaged firewall/VPN solutions, but not many.

WireGuard

WireGuard, like OpenVPN, is compatible with only a few other packaged firewall and VPN solutions.

20.2.2 Authentication considerations

All VPN types on the firewall are capable of user authentication except WireGuard. IPsec can also work with shared keys, and both IPsec and OpenVPN can utilize certificates. Using OpenVPN with certificates, TLS authentication, and User Authentication is the most secure method. OpenVPN certificates can also be password protected, in which case a compromised certificate alone isn't adequate for connecting to a VPN if it is set to only use certificates. The lack of additional authentication can be a security risk in that a lost, stolen, or compromised system containing a key or certificate means whoever has access to the device can connect to a VPN until that loss is discovered and the certificate revoked.

While not ideal, a lack of username and password authentication on a VPN isn't as great a risk as it may seem. A compromised system can easily have a keylogger installed to capture the username and password information and easily defeat that protection. In the case of lost or stolen systems containing keys, if the hard drive isn't encrypted, the keys can be used to connect. However, adding password authentication isn't of great help there either, as usually the same username and password will be used to log into the computer, and most passwords are crackable within minutes using modern hardware when an attacker has access to an unencrypted drive. Password security is also frequently compromised by users with notes on their laptop or in their laptop case with their password written down. As with any security implementation, the more layers utilized, the better, but it's always a good idea to keep these layers in perspective.

20.2.3 Ease of configuration

None of the available VPN options are extremely difficult to configure, but there are differences between the options.

IPsec

Has numerous configuration options and can be difficult for the uninitiated.

OpenVPN

OpenVPN requires the use of certificates for remote access in most environments, which comes with its own learning curve and can be a bit arduous to manage. There is a wizard to handle the most common OpenVPN remote access configurations and the OpenVPN client export packages eases the process of getting the clients up and running.

WireGuard

Has few options, thus configuration is simple. Lacks facilities for automated configuration and deployment leading to increased manual management.

20.2.4 Multi-WAN capable

If users require the ability to connect to multiple WANs, IPsec, OpenVPN, and WireGuard are capable of handling such configurations.

20.2.5 Client availability

VPN Client software is a program that handles connecting to the VPN and handling any other related tasks like authentication, encrypting, routing, etc. For remote access VPNs, the availability of VPN client software is a primary consideration. All options are cross-platform compatible with many different operating systems, but some require installing third-party clients. IPsec in EAP-MSCHAPv2 mode, IPsec in EAP-TLS mode, and IPsec in Xauth mode are the only options with clients built into some popular desktop and mobile operating systems. Other operating systems vary and may include more or less IPsec modes or may even include OpenVPN or WireGuard, as is the case with many Linux distributions. If using built-in clients is a must, consult the operating system documentation for all required client platforms to see if a common option is available and then check to see if that mode is possible.

In some cases multiple remote access VPNs may be required to accommodate all clients. For example, IPsec could be used for some and OpenVPN for others. Some organizations prefer to keep things consistent, so there is a trade-off to be made but for the sake of compatibility it may be worth offering multiple options.

IPsec

IPsec clients are available for Windows, macOS, BSD, Linux, and others. Though the native clients may only implement certain specific modes and configurations. General-use IPsec clients are not included in the OS except for some Linux and BSD distributions. macOS includes both IKEv2 and Cisco (xauth) IPsec clients. There are free and commercial options available with a user-friendly GUI.

macOS 10.11, along with Windows 7 and later include clients for IPsec in specific modes using IKEv2: EAP-TLS and EAP-MSCHAPv2. Both options are compatible with pfSense software and are covered in *IPsec Mobile Clients Tab*.

The Cisco-style IPsec client included with macOS and iOS devices is fully compatible with IPsec using xauth. Configuration for the iOS client is covered in *Configuring IPsec IKEv2 Remote Access VPN Clients on iOS*.

Many Android phones also include a compatible IPsec client, which is discussed in *Configuring IPsec IKEv2 Remote Access VPN Clients on Android*.

OpenVPN

OpenVPN has clients available for Windows, macOS, all the BSDs, Linux, Solaris, and Windows Mobile, but the client does not come pre-installed in any of these operating systems.

Android devices can use a freely available OpenVPN client that works well and doesn't require rooting the device. That client is covered in *Installing the OpenVPN Client on Android*. There are other options available if the device is rooted, but that is beyond the scope of this documentation.

iOS also has a native OpenVPN client. For more information, see *Installing the OpenVPN Client on iOS*.

WireGuard

WireGuard clients are available for a variety of operating systems including Windows, macOS, FreeBSD, Linux, Android, iOS, and more. Some Linux installations include WireGuard but in most cases it requires installation of a third party client.

20.2.6 Firewall friendliness

VPN protocols can cause difficulties for many firewalls and NAT devices. This is primarily relevant to remote access connectivity, where users will be behind a myriad of firewalls mostly controlled by third parties with varying configurations and capabilities.

IPsec

IPsec uses both UDP port 500 and the ESP protocol to function. Some firewalls don't handle ESP traffic well where NAT is involved, because the protocol does not have port numbers like TCP and UDP that make it easily trackable by NAT devices. IPsec clients behind NAT may require NAT Traversal to function, which encapsulates the ESP traffic over UDP port 4500.

OpenVPN

OpenVPN is very firewall-friendly. Since it uses a single UDP or TCP port and is not affected by common NAT functions such as rewriting of source ports, it is rare to find a firewall which will not work with OpenVPN. The only possible difficulty is if the protocol and port in use is blocked. Some administrators use a common port like UDP 53 (usually DNS), or TCP 80 (usually HTTP) or TCP 443 (usually HTTPS) or to evade most egress filtering.

WireGuard

Similar to OpenVPN in this regard, WireGuard uses a single UDP port and thus is not affected by firewall and NAT issues which may affect other protocols.

20.2.7 Cryptographically secure

One of the critical functions of a VPN is to ensure the confidentiality of the data transmitted.

IPsec

Tunnels using pre-shared keys can be broken if a weak key is used. Use a strong key, at least 10 characters in length containing a mix of upper and lowercase letters, numbers and symbols. Use of certificates is preferred, though somewhat more complicated to implement.

OpenVPN

Encryption is compromised if the private keys of the PKI structure are compromised, though the use of multiple factors such as TLS authentication on top of PKI can mitigate some of the danger.

WireGuard

WireGuard encryption relies on pre-generated public/private key pairs and an optional pre-shared key. The peers only need the public keys of other peers, and the optional pre-shared key. The private keys and pre-shared key (if present) must be protected.

20.2.8 NAT Inside Tunnels

While any use of NAT is undesirable, there are some occasions which can benefit from its use inside tunnels. Primarily, it can be useful for working around subnet conflicts or for setting up "outbound" style NAT when the remote endpoint only expects a single address (e.g. a VPN provider with no LAN-to-LAN routing)

IPsec

NAT with IPsec depends on the mode, either tunnel or VTI.

Tunnel

Phase 2 entries in *tunnel* mode can use BINAT (1:1) and Overload/PAT style NAT. See *NAT with IPsec Phase 2 Networks* for details.

VTI

Phase 2 entries in *VTI* mode can utilize NAT when using a special **IPsec Filter Mode** setting which is not compatible with tunnel mode. See *Filtered on Assigned IPsec Interfaces* for details.

OpenVPN

OpenVPN is compatible with inbound NAT (e.g. port forwards) and outbound NAT using the group OpenVPN tab and also on assigned interfaces. Depending on the environment and configuration there may be some special considerations, such as ensuring proper return routing for post-NAT subnets.

WireGuard

WireGuard is compatible with inbound NAT (e.g. port forwards) and outbound NAT using the group WireGuard tab and also on assigned interfaces. Some cases may require using single peer tunnels or carefully crafted **Allowed IPs** lists to ensure correct return routing. See *Design Considerations* and *WireGuard and Rules / NAT*.

20.2.9 Per-tunnel Firewall Rules

Each VPN type has a common group tab for rules, and some are also capable of using separate rulesets for individual tunnels.

Warning: Rules on group tabs are considered before per-interface rules. For per-interface rules to match, rules on the group tab must not match the same packets.

IPsec

IPsec in tunnel mode is not currently compatible with per-tunnel rules, its traffic can only be filtered by rules on the IPsec tab.

Phase 2 entries in *VTI* mode can utilize per-interface rules when using a special **IPsec Filter Mode** setting which is not compatible with tunnel mode. See *Filtered on Assigned IPsec Interfaces* for details.

OpenVPN

When assigned as an interface, OpenVPN instances can utilize per-tunnel rules. See *Assigning OpenVPN Interfaces*.

WireGuard

When assigned as an interface, WireGuard instances can utilize per-tunnel rules. See *Assign a WireGuard Interface* and *WireGuard and Rules / NAT*.

20.2.10 Automatic Mobile Client Configuration

Depending on the deployment, mobile (Remote Access) clients can receive automatic configuration in certain cases.

IPsec

In IKEv2 mode, clients can automatically receive an IP address allocated from a pool, along with DNS configuration.

In IKEv1 mode with Xauth, in addition to the above, clients can also receive a list of networks to route across the VPN.

OpenVPN

OpenVPN clients can automatically receive an IP address allocated from a pool, and numerous additional options can be pushed to clients to control their behavior from the server side (routing, DNS, and many others).

WireGuard

WireGuard mobile clients must be configured statically. On the server side, a client tunnel address must be setup in the **Allowed IPs** for a peer. The same address must be configured on the client. This varies by OS/Platform, some read it from the configuration and others require it to be configured on interfaces via CLI commands. Networks to route must likewise be manually added on the client configuration **Allowed IPs** list and, depending on the client, may also need to be added to its operating system routing table.

20.2.11 Routing

IPsec

IPsec in *Tunnel* mode uses policies, not routes, and thus does not respect the operating system routing table.

IPsec in *VTI* mode is compatible with static and dynamic routing (e.g. BGP, OSPF) and works with the operating system routing table.

OpenVPN

In SSL/TLS tun mode with multiple clients, OpenVPN uses its internal routing on client-specific configurations to determine which clients receive traffic for specific subnets. In this type of configuration, dynamic routing is not possible.

In SSL/TLS tun mode with a /30 size subnet (one client per server), dynamic routing is possible using OSPF or BGP. In that configuration OpenVPN does not need to track internal routing and can rely on the operating system routing table alone.

In tap mode, dynamic routing is possible as packets can be handed off using L2/ARP information rather than relying on internal routing in OpenVPN.

WireGuard

WireGuard routes specific subnets to peers based on the **Allowed IPs** list, but also requires operating system routing table entries for traffic to enter a WireGuard tunnel.

When a WireGuard tunnel has more than one peer, the **Allowed IPs** list lets WireGuard determine internally which clients receive traffic for specific subnets. Due to this internal routing, dynamic routing is not possible in a configuration where WireGuard has multiple peers per tunnel.

For a WireGuard tunnel with a single peer, WireGuard can forward arbitrary networks to the peer without having them all listed in **Allowed IPs**. Thus, in this situation, it can take advantage of dynamic routing using BGP. OSPF is also possible but requires additional configuration.

See *WireGuard Routing* for more information.

20.2.12 Recap

Table *Features and Characteristics by VPN Type* shows an overview of the considerations provided in this section.

Table 1: Features and Characteristics by VPN Type

VPN Feature	IPsec	OpenVPN	WireGuard
User Authentication	Yes	Yes	No
Client included in most OSes	Varies by mode	No	No
Widely interoperable	Yes	No	No
Multi-WAN	Yes	Yes	Yes
Cryptographically secure	Yes	Yes	Yes
Firewall friendly	Yes (NAT-T or IKEv2)	Yes	Yes
In-tunnel NAT	Limited	Yes	Yes
Per-tunnel Firewall Rules	Limited	Yes	Yes
Automatic Mobile Config	Yes	Yes	No
Static Routing	Yes (VTI Only)	Internal	Internal
Dynamic Routing	Yes (VTI Only)	Varies	Varies

20.3 Remote Access Mobile VPN Client Compatibility

A variety of remote access ("mobile") VPN configuration styles are available to accommodate nearly any potential client. The table below shows which operating systems have compatible clients with some of the most common remote access VPN configurations.

Table 2: Mobile/Remote Access VPN Client Availability

Protocol/Operating System	Windows 10	Android	iOS	macOS
OpenVPN	3PA 1 2	3PA 4	3PA 6	3PA 2
WireGuard	3PA 3	3PA 3	3PA 3	3PA 3
IPsec PSK	?	Varies	?	?
IPsec RSA	?	Varies	?	?
IPsec IKEv2 EAP-MSCHAPv2/RADIUS	Yes	Yes (11+), 3PA (4.x+) 5	Yes	Yes
IPsec IKEv2 EAP-TLS	Yes	Yes (11+), 3PA (4.x+) 5	Yes	Yes

Legend:

- Yes = OS Native Client Available
- 3PA = Third Party Client Required
- Varies = Varies by device model and vendor options

Unless otherwise stated, UNIX clients (*BSD, Linux, etc.) can utilize any style with manual configurations, but the availability of GUI configuration tools varies by distribution.

20.4 VPNs and Firewall Rules

VPNs and firewall rules are handled somewhat inconsistently in pfSense® software. This section describes how firewall rules are handled for each of the individual VPN options. For the automatically added rules discussed here, the addition of those rules may be disabled by checking **Disable all auto-added VPN rules** under **System > Advanced** on the **Firewall/NAT** tab.

20.4.1 IPsec

Traffic necessary to establish configured and enabled IPsec tunnels is automatically allowed into the firewall as described in *Outer IPsec Traffic*.

Traffic encapsulated within an active tunnel mode IPsec connection is controlled via user-defined rules on the **IPsec** tab under **Firewall > Rules**. Traffic for VTI mode works the same way by default but can operate on a per-interface basis in certain conditions. See *Tunneled IPsec Traffic from Remote to Local* for details.

20.4.2 OpenVPN

OpenVPN does not automatically add rules to WAN interfaces. The OpenVPN remote access VPN Wizard offers to optionally create rules to pass WAN traffic and traffic on the OpenVPN interface.

Traffic encapsulated within an active OpenVPN connection is controlled via user-defined rules on the **OpenVPN** tab under **Firewall > Rules**.

OpenVPN interfaces may also be assigned similar to other interfaces. In such cases the **OpenVPN** tab firewall rules still apply, but there is a separate tab specific to the assigned VPN instance that controls traffic only for that one VPN.

20.4.3 WireGuard

WireGuard does not automatically add rules to WAN interfaces. Rules must be added to the appropriate WAN interface(s) to allow traffic to reach the ports for WireGuard instances.

Traffic encapsulated within WireGuard is controlled via user-defined rules on the **WireGuard** tab under **Firewall > Rules**.

WireGuard interfaces may also be assigned similar to other interfaces. In such cases the **WireGuard** tab firewall rules still apply, but there is a separate tab specific to the assigned VPN instance that controls traffic only for that one VPN.

20.5 IPv6 VPN and Firewall Rules

As mentioned briefly in *Firewall and VPN Concerns*, special care must be taken when routing IPv6 traffic across a VPN and using publicly routable subnets. The same advice also applies to IPv4, but it's much less common to have clients on both sides of an IPv4 VPN using publicly routable addresses.

The main issue is that because it's possible to route all the way from one LAN to the other LAN across the Internet, then traffic could be flowing unencrypted between the two networks if the VPN is down (or not present at all!). This is far from ideal because although connectivity is available, if any traffic were intercepted in between the two networks and that traffic was using an unencrypted protocol like HTTP, then it could compromise the network.

One way to prevent this is to not allow traffic from the remote IPv6 LAN in on the opposing side's WAN rules. Only allow traffic from the remote side's subnet on the firewall rules for whichever VPN type is being used to protect the traffic. An explicit block rule could also be added to the top of the WAN rules to ensure that this traffic cannot enter from the WAN directly. A better method is to use a floating rule to reject outbound traffic on WAN destined for VPN hosts/remote local networks. This way the insecure traffic never leaves the premises. With the rule set to log, the "leakage" would be obvious to someone monitoring the logs as it would be shown blocked outbound on WAN.

Another less obvious consequence of having dual stack connectivity between networks is that differences in DNS can cause unintended routing to take place. Suppose IPv4 VPN connectivity exists between two sites, but there is no IPv6 VPN, only standard IPv6 connectivity at both locations. If a local host is set to prefer IPv6, and it receives a AAAA DNS response with the IPv6 IP address for a remote resource, it would attempt to connect over IPv6 first rather than using the VPN. In cases such as this, care would be needed to make sure that DNS does not contain conflicting records or that floating rules are added to prevent this IPv6 traffic from leaking out WAN. A more in-depth article on these kinds of traffic leakage can be found in the IETF draft named draft-gont-opsec-vpn-leakages-00.

20.6 VPN Scaling

The advice on this page is intended to help firewall administrators handle increased VPN volume when using pfSense® software, both in terms of throughput and number of connected users.

> **Warning:** The advice on this page is relayed from experience and from community members. The advice on this page may not apply to all environments or use cases, and has not been definitively proven to help, but is offered in case others find it useful.

See also:

This document won't cover all factors of choosing between VPN types or how to set up VPNs, that information can be found *elsewhere in this documentation* and in Hangouts Archive.

20.6.1 General Advice

No Artificial Limits

The firewall does not place any artificial limits on VPN connections. Any limitations encountered are due to settings, the hardware/environment, or the underlying technology.

Use External Authentication

For user-based authentication, the most efficient method of user management for large numbers of accounts is an external authentication source, such as a RADIUS server, LDAP server, Active Directory (Via LDAP or RADIUS/NPS), etc.

Check Logs

If additional users are unable to connect, look in the logs on both the client and server side for specific error messages before seeking assistance.

Use Hardware Acceleration

Using a *cryptographic accelerator* such as a QAT, IPsec-MB, AES-NI, or SafeXcel will help greatly with throughput and crypto-related tasks.

pfSense Plus software includes drivers for QAT, CESA, and SafeXcel hardware found on several Netgate Appliances. QAT is also available as an add-on card for certain models. IPsec-MB is also available on pfSense Plus for compatible CPUs. pfSense Plus software and pfSense CE software both include drivers for AES-NI.

For more information on these devices and their capabilities, including how to configure and test them, see *Cryptographic Accelerators*.

Use AES-GCM or ChaCha20-Poly1305

Using efficient encryption will increase security and performance. Authenticated Encryption with Additional Data (AEAD) ciphers combine encryption and authentication in one step, eliminating the need for additional hashing. AES-GCM and ChaCha20-Poly1305 are both AEAD ciphers. IPsec and OpenVPN can use both AES-GCM and ChaCha20-Poly1305, while the only cipher implemented by WireGuard is ChaCha20-Poly1305.

Client implementations vary by platform.

Use Accelerated Ciphers

Certain hardware may accelerate ciphers so that choices are faster or more efficient. For hardware sold by Netgate, see the Netgate Appliances page for performance data and recommendations.

Disable Performance-Limiting Mitigation Settings

While Netgate does not have data on if or by how much they may impact VPNs, CPU vulnerability mitigation methods such as Kernel PTI and MDS mode can potentially degrade total performance. The potential for exploitation is minimal since arbitrary code cannot be run on the firewall except by users which already have the equivalent of administrator-level access. To ensure this risk stays low, only allow trusted administrators to access the firewall GUI and shell (SSH or console). The settings to enable/disable these features are under **System > Advanced** on the **Miscellaneous** tab.

Check Tunnel Network/Virtual Address Pool Sizes

Both IPsec and OpenVPN can assign addresses to clients out of a pool for remote access/mobile VPNs. The sizing of this pool limits how many clients can connect. For example, the maximum number of users in a /24 pool is 252, but other settings may reduce that value. See the sections below for more specific advice.

Use "Secure Enough" Settings

While we do not recommend deliberately using weak configurations, in some cases trade-offs are made for security between two secure ciphers or settings where one may offer *even better* security, but the lower of the two is still secure. In these cases, using the "Secure Enough" option can provide efficiency vs increased security. So long as the decision is informed, there may be some performance gained without compromising security in an unacceptable way. For example, with AES-GCM a key length of 128 bits is still considered secure. A 256-bit key is more secure, but the 256-bit key could put more of a burden on the hardware.

Consider Split Tunneling

Configurations which send all client data over the VPN, including Internet-bound traffic, will consume more resources than those which only send traffic for specific subnets. There are plenty of valid reasons to use either kind of configuration, however, when resources are stretched thin, easing the traffic burden on the VPN may justify switching to split tunneling rather than tunneling everything. Depending on the type of VPN and client, this may require adjustments on the server, the client, or both. See the sections below for specific recommendations.

Use Multiple Firewalls

In some instances, the burden may be too great for any single firewall to handle. In cases like this, multiple firewalls can be used to handle the required number of clients or throughput, at a cost of greatly increased complexity. There is no way to automatically balance between nodes in this manner, but such a configuration could be manually managed. This would also likely require the capability to have multiple external addresses on the WAN, so each firewall can work in parallel, and also increases the complexity of routing on the internal side.

Use TNSR

TNSR® software is capable of vastly increased total IPsec and WireGuard throughput compared to pfSense software. If pfSense software is unable to reach the throughput needs for a given use case, see the TNSR product page for more information.

20.6.2 Scaling IPsec

IPsec is well-suited to high throughput by default, especially given the advice above, but there are additional IPsec-specific tweaks which may help.

Note: See the TNSR product page for information about using TNSR for even larger total site-to-site throughput needs.

Optimal Encryption Settings

- Use QAT, IPsec-MB, or AES-NI capable hardware.
- In Phase 1 (IKE) settings, use:
  - *AES128-GCM* with *128-bit* key length for the Algorithm
  - *AES-XCBC* for the hash, which in this case is effectively a Pseudo-Random Function (PRF). This will yield the highest performance in combination with *AES128-GCM* on hardware which can accelerate both (e.g. AES-NI)

> **Warning:** While this may lead to higher performance when hardware is capable of accelerating this algorithm, it is less secure than other choices such as SHA256.
>
> Using a more secure algorithm is better in most use cases. IKE negotiation does not happen frequently unless there are large quantities of tunnels or mobile clients, so it may not be critical to accelerate this function in the majority of deployments.

- In Phase 2 (Child SA) settings, use:
  - *AES128-GCM* with *128-bit* key length for the Algorithm
  - **Do not select any Hash Algorithms**. A hash algorithm is unnecessary for AES-GCM as it already includes authentication.

Enable Multiple Phase 1 and Phase 2 Proposals

Multiple Phase 1 and Phase 2 encryption proposals may be configured in the GUI. Enabling multiple combinations of settings will allow peers to choose the most optimal settings compatible with both endpoints.

Enable Asynchronous Cryptography

IPsec cryptography jobs can be dispatched multithreaded to run in parallel and increase performance. However, not all platforms and configurations fully implement this function. To enable this capability, check **Asynchronous Cryptography** under **VPN > IPsec** on the **Advanced** tab.

> **Warning:** Be on the lookout for IPsec traffic drops/failures to pass with this setting enabled. See https://redmine.pfsense.org/issues/8964 for more information.

Split Tunneling

As mentioned above, split tunneling would only send traffic for specific subnets across the VPN rather than sending all traffic. On IPsec, this can be done in some cases by listing the specific networks in Phase 2 entries for the Mobile IPsec P1 rather than `0.0.0.0/0`. On the mobile clients tab, set **Provide a list of accessible networks to clients**. Even with that set, certain cases such as Windows 10 may require additional changes to direct clients to send only specific traffic over the tunnel.

20.6.3 Scaling OpenVPN

Use Data Channel Offload (Plus Only)

OpenVPN Data Channel Offload (DCO), a pfSense Plus exclusive feature, can potentially increase performance of OpenVPN well beyond the capabilities of traditional OpenVPN connections. Under ideal conditions OpenVPN with DCO can match or exceed the performance of WireGuard and IPsec.

DCO operates optimally when it is enabled on all peers, but it can still offer performance benefits when enabled on only one side.

Use IPsec or WireGuard Instead

If DCO cannot be enabled on all OpenVPN peers, then check if the peers are capable of using IPsec or WireGuard instead. IPsec and WireGuard are much more efficiently integrated into the operating system than traditional OpenVPN (without DCO), and both are capable of much greater throughput in that situation.

Use UDP

UDP has less overhead for tunneled data, and if a client has to retransmit, it won't compound the problem by retransmitting both inside and outside the tunnel. Unless there are extenuating circumstances which require TCP, use UDP.

Use TLS for Authentication Only

OpenVPN can use TLS for both authentication and for encryption of the control channel. Performing control channel encryption adds more overhead, which can add up with many clients. If control channel encryption is not required, consider using TLS for only authentication instead. No matter which option is chosen, traffic carried by OpenVPN is encrypted.

Encryption Algorithm

Use a CPU with QAT, IPsec-MB, or AES-NI when possible, and use AES-GCM for the Encryption Algorithm when possible. ChaCha20-Poly1305 may also give a performance boost if hardware acceleration is unavailable. Note that for AEAD ciphers such as AES-GCM and ChaCha20-Poly1305, OpenVPN ignores the setting for **Auth Digest Algorithm**.

Note: For OpenVPN, AES-GCM and ChaCha20-Poly1305 can only be used in SSL/TLS mode with a tunnel network that enables client/server mode (larger than /30).

Use Data Cipher Negotiation

Data cipher negotiation can be used to set preferences so that more efficient ciphers can be preferred by clients where possible, but others can be used when necessary. Set high-priority selections such as *AES-128-GCM* first, followed by others like *AES-128-CBC*.

Split Tunneling

As mentioned in the general section above, split tunneling only sends traffic for specific subnets across the VPN rather than sending all traffic. With OpenVPN, this can be done by **Unchecking** the **Redirect IPv4/IPv6 Gateway** option(s) and configuring **IPv4/IPv6 Local Network(s)** entries instead. Clients may still override this behavior remotely, however, so check the client configurations as well.

Concurrent Connections

The firewall does not impose any connection limits by default, but an administrator may have chosen to configure a limit on the number of connections via the **Concurrent Connections** setting on servers. Ensure this is either unset or set high enough to accommodate the required number of users.

Disable Compression

Though using compression is tempting to squeeze extra throughput out of slower links, it is both inefficient and insecure. Most data sent across VPNs in modern environments is already encrypted or otherwise uncompressible, which wastes CPU when attempting to compress. Additionally, vulnerabilities such as VORACLE can allow attackers to glean information about encrypted data when it has been compressed. Disabling compression will mitigate that attack and also reduce CPU overhead. On the server, set **Compression** to *Disable Compression*.

Duplicate Connections

Normally, if an OpenVPN client connects using the same username or certificate CN, the older connection is broken in favor of the new connection. This is more secure, but does not allow any given user to connect multiple times. Circumstances may necessitate enabling this, and in some environments it's not possible to give every device a unique username and/or certificate. Check **Duplicate Connection** in the OpenVPN server settings to allow multiple connections from the same user.

Topology

OpenVPN defaults to *subnet* topology which uses addresses more efficiently, but if the VPN was configured initially on older versions, or if an older guide was followed, it may still be using *net30* topology. Using a common example tunnel network of `10.0.8.0/24`, with *subnet* topology, the VPN can have a maximum of 252 users but with *net30*, it can only have 63. This is because in *net30* mode, each user receives a /30 size subnet which utilizes four IP addresses for each user. In *subnet* mode, the server uses a single address and the client uses a single address, which is much more efficient.

Use UDP Fast I/O

This option is considered experimental by OpenVPN, but for those who have used it, it can result in much higher throughput. Not all platforms implement this feature, however.

Increase Send/Receive Buffer

The default buffer size is safe, but not optimal. Increasing the buffer size to 512KiB on both sides can result in greater throughput. Results will vary by platform, internet link speed, and other factors. May require experimenting with multiple values to find the most efficient setting for a given environment.

Use Multiple Servers

OpenVPN is not multithreaded so any single instance of OpenVPN is limited to using a single CPU. If a router has fast cores and not too many users, that may be OK, but it does not scale well. A workaround for this is to split users onto multiple servers. There are various means to reach this goal, including (but not limited to):

- Multiple servers on different WANs or ports, each with unique tunnel networks but otherwise identical settings (Same CA structure, encryption, etc.).
  - Administrators could choose to manually configure pools of clients to connect to specific servers, but that does not scale well.
  - Clients may connect to any server configured in this manner so long as their settings line up properly.
  - Multiple servers can be listed in a single client configuration with additional `remote` statements.
  - Add `remote-random` to the client configuration so that clients will pick a random server when starting, which avoids overloading whichever server is listed first.
  - Servers could be run on multiple WANs to overcome single-circuit throughput limits.
- Multiple servers with completely unique settings (Different CA structure, different clients, etc.)
  - More secure but more difficult to manage.
  - Clients must use different configurations to reach each server, there is no automated or built-in way to pick between them unless a specific client offers that functionality.
  - Good for isolating separate security levels (e.g. remote workers, remote administrators, vendors).

Process Efficiency

As a counterpoint to the above, each server will incur additional memory and other overhead to manage the process. When dealing with site-to-site VPNs, it is more efficient from a *memory* standpoint to use a single server with multiple clients (Peer to Peer SSL/TLS in client/server mode) vs servers for every node (Peer to Peer SSL/TLS with a /30 tunnel network). If memory is a limiting factor, use fewer servers. If CPU overhead is the limiting factor, use separate servers.

20.7 OpenVPN

20.7.1 OpenVPN Data Channel Offload (DCO)

OpenVPN Data Channel Offload (DCO) allows for huge performance gains when processing encrypted OpenVPN data by reducing the amount of context switching that happens for each packet. DCO accomplishes this by keeping most of the data handling tasks in the kernel rather than repeatedly switching between kernel and user space for encryption and packet handling. This makes the overall processing of each packet more efficient while also potentially taking advantage of hardware encryption offloading in the kernel. DCO also adds multithreaded encryption, allowing for even more performance gains.

Netgate worked with OpenVPN to develop and integrate OpenVPN Data Channel Offload (DCO) into FreeBSD and pfSense Plus software version 22.05 and later.

Warning: pfSense Plus software version 22.05 or later is required to use OpenVPN DCO. OpenVPN DCO is not available on pfSense CE Software.

DCO is not a change to the protocol, it is a change in how an endpoint processes encrypted data. Thus, DCO is beneficial even when only one endpoint is capable of DCO. That said, tunnels employing DCO on all peers will see the most benefit. With DCO on only one peer the performance improvement can still be notable but not as significant as the gains with DCO enabled on both endpoints.

Note: Some OpenVPN features and use cases are not compatible with DCO. See *Limitations* for a list of known DCO limitations.

See also:

OpenVPN Site-to-Site Configuration Example with SSL/TLS and DCO

Using OpenVPN DCO

DCO is a per-tunnel option, and it is **not** automatically enabled by default for new or upgraded tunnels. Existing tunnels will continue to function as they have in the past.

DCO can be enabled for both new and existing tunnels by using a simple checkbox option on OpenVPN server and client instances. The current best practice is to create a new tunnel with DCO to minimize the chance of problems with existing clients.

Limitations

There are a few limitations in OpenVPN DCO generally and in the current DCO implementation on FreeBSD/pfSense software, including:

- Encryption is limited to AES-256-GCM, AES-128-GCM, and ChaCha20-Poly1305.
- DCO requires a TLS-based tunnel, such as SSL/TLS, SSL/TLS+User Auth, or User Auth.
- DCO is only present in OpenVPN 2.6.0 and later.
- DCO is only compatible with UDP, it cannot be used with TCP.
- DCO is not yet able to utilize internal routing in OpenVPN (`iroute`). This means that although remote access use cases work, and site-to-site setups with one client per server work, it does not yet function with multiple site-to-site clients on a single server which require LAN-to-LAN routing.
- Using a `/30` or smaller tunnel network for peer-to-peer tunnels (one server with one client) is not compatible with DCO. There are problems with the code for this mode in OpenVPN which can lead to failed connections and instability.
- Compression is not compatible with DCO. The GUI disables compression options when DCO is enabled for an instance, but for a client instance the server could still push a compression option which would make the client fail to pass traffic.
- Some features are not compatible with DCO or are not relevant with DCO. These options include:
  - Explicit exit notify
  - Inactivity timeouts
  - UDP fast I/O
  - Send/receive buffer sizes
- Per-peer data usage is not tracked properly.

  Until this is resolved peer data usage on the OpenVPN status page will not reflect the actual amount of data transferred between peers.

DCO and Routing

DCO does not currently honor internal routes from client-specific overrides (i.e. `iroute`) for multiple site-to-site clients on a single server, but it does honor kernel route destinations that would normally be ignored by non-DCO OpenVPN.

Assign clients static addresses in overrides (after patching #13274) and then setup custom routes in OpenVPN custom options with complete destinations defined or even setup FRR and exchange routes via BGP.

DCO and Hardware Cryptographic Acceleration

For optimal performance with DCO, ensure a hardware cryptographic accelerator is present and enabled.

See also:

- *Cryptographic Accelerators*
- *Cryptographic Hardware*

QAT currently offers the highest performance for AES-256-GCM. If the hardware includes a QAT device, enable QAT.

If there is no QAT device available, but the CPU has SIMD instruction sets, then enable IPsec-MB and use AES-GCM, ChaCha20-Poly1305, or even AES-CBC. This can also benefit uses of these ciphers which are not yet accelerated by QAT.

If the hardware does not have QAT or IPsec-MB, but it does have AES-NI, ensure the AES-NI kernel module is **loaded** for optimal performance with AES-256-GCM. Though OpenSSL can utilize AES-NI without the module loaded, performance is poor in that state and can even be slower than with DCO disabled.

Note: pfSense Plus software can utilize ChaCha20-Poly1305 with OpenVPN DCO, but currently only IPsec-MB can accelerate that algorithm. At this time, neither AES-NI nor QAT can accelerate ChaCha20-Poly1305. Some newer QAT hardware may be capable of accelerating ChaCha20-Poly1305, but the current QAT drivers do not yet implement that encryption algorithm.

20.7.2 OpenVPN Configuration Options

This section describes available options for use with OpenVPN for servers, clients, and custom options.

Server Configuration Options

These options are available in one or more modes for OpenVPN server instances, managed from **VPN > OpenVPN**, on the **Servers** tab.

Note: The options available to clients overlap significantly, and differences are called out where appropriate. See *Client Configuration Options* later in this document for details.

General Information

Description

Enter a description for this VPN instance, for reference.

Disabled

Check this box and click **Save** to retain the configuration, but not enable this instance. The firewall will stop the process for this instance, which will disconnect all peers from the VPN. Any other active instances are unaffected.

Unique VPN ID

A read-only field which shows the internal VPN ID and interface name for this OpenVPN instance.

Mode Configuration

Server Mode

The role for the server, which specifies how peers connect to a server instance. Changing this also affects which options the GUI displays on the rest of the page.

Peer to Peer (SSL/TLS)
: A connection between local and remote networks that is secured by SSL/TLS.

  This choice offers increased security as well as the ability for the server to push configuration commands to the remote peer router when using a 1:many style setup. Remote peer routers can also have certificates revoked to remove access if they become compromised.

Peer to Peer (Shared Key)
: A connection between local and remote networks that is secured by a single shared key known to both peers.

  This choice is easier to set up, but is less secure. If a shared key is compromised, a new key must be generated and then copied to any router or client using the old shared key. This mode requires a separate server instance for each client.

> **Danger:** Shared key mode has been deprecated by OpenVPN as it is no longer considered sufficiently secure for modern requirements.
>
> **Shared key mode will be removed from future versions of OpenVPN**. Users **should not** create any new shared key tunnels and should **immediately** convert any existing shared key tunnels to SSL/TLS mode.
>
> When an SSL/TLS instance is configured with a `/30` tunnel network it behaves similarly to shared key mode. The primary difference is the need to create and distribute the certificate structure to peers. See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for information on configuring OpenVPN in SSL/TLS mode.

Remote Access (SSL/TLS)
: A mobile client setup with per-user X.509 certificates.

  As with the peer-to-peer SSL/TLS connection type, using this method offers increased security as well as the ability for the server to push configuration commands to clients. Mobile clients can also

have keys revoked to remove access if a key is compromised, such as a stolen or misplaced phone or laptop.

Remote Access (User Auth)
A remote access server configuration that does not use certificates, but requires the end user to supply a username and password to authenticate.

This is less secure than using certificates, but simpler to deploy as every client will use the same configuration file and only their credentials are different.

Remote Access (SSL/TLS + User Auth)
Requires SSL/TLS and user authentication to connect.

This is the most secure choice available. Not only does it get the benefits of other SSL/TLS choices, but it also requires a username and password from the client when it connects. Client access can be removed not only by revoking the certificate, but also by changing the password. Also, if a compromised key is not immediately discovered, the danger is lessened because it is unlikely that the attacker has the keys and the password.

The OpenVPN wizard uses this mode when it configures a remote access VPN.

DCO (Plus Only)

Check **Enable Data Channel Offload (DCO) for this instance** to enable *OpenVPN Data Channel Offload (DCO)*.

DCO is a pfSense Plus exclusive feature that can potentially increase performance of OpenVPN well beyond the capabilities of traditional OpenVPN connections.

To be compatible with DCO the instance must use SSL/TLS in client/server mode. This can either be remote access or peer-to-peer so long as the tunnel network is large enough for multiple clients (e.g. `/24`).

DCO can be enabled on one or more peers, but the greatest speed increase comes when all peers have DCO enabled.

Warning: Several OpenVPN options are not compatible with DCO, which are noted on *OpenVPN Data Channel Offload (DCO)* and elsewhere in the documentation where relevant.

Note: Some OpenVPN features and use cases are not compatible with DCO. See *Limitations* for a list of known DCO limitations.

Backend for authentication

One or more authentication servers to use when checking user credentials. This can be the *Local Database* (user manager) or it can be a RADIUS or LDAP server configured in the user manager.

Selecting multiple entries will use each one in turn. If authentication against an entry fails, the VPN will try the next server. It will continue this process until it has tried all selected entries.

This field is only present when using a **Server Mode** which requires user authentication.

Device Mode

OpenVPN can run in one of two device modes: *tun* or *tap*:

tun
Works on OSI layer 3 and performs routing on point-to-point interfaces.

tap
Can work at OSI layer 2 and can perform both routing and bridging if necessary.

Note: Not all clients implement *tap* mode. Clients such as those found on Android and iOS can only utilize *tun* mode in the apps most people can use. Some Android and iOS OpenVPN apps that require rooting or jailbreaking a device do allow *tap*, but the consequences of doing so can be a bit too high for most users.

Endpoint Configuration

Protocol

The IP protocol clients will use to connect to this VPN.

UDP on IPv4/IPv6 only
UDP is the most reliable and fastest choice for running OpenVPN, and the best practice is to always use UDP when possible.

Connectionless protocols such as UDP are always the best practice to use when tunneling traffic.

This mode binds to a single interface and limits OpenVPN to only accepting IPv4 or IPv6 exclusively, not both at the same time.

TCP on IPv4/IPv6 only
Using TCP for a VPN is slower and can be more problematic. In some rare cases TCP can be work around limitations of a client environment, such as bypassing firewalls by running an OpenVPN server on TCP port `443`.

TCP is connection oriented with guaranteed delivery, which means any lost packets are retransmitted. This sounds like a good idea on the surface but TCP retransmissions will cause performance to degrade significantly on heavily loaded Internet connections or those with consistent packet loss.

TCP traffic frequently exists within tunnels, and it is undesirable to retransmit lost packets of encapsulated VPN traffic. In cases where TCP is wrapped around TCP, such as a VPN tunnel using TCP as a transport protocol, when a packet is lost both the outer and inner lost TCP packets will be re-transmitted. Infrequent occurrences of this will be unnoticeable but recurring loss will cause significantly lower performance than UDP. If the traffic inside the tunnel requires reliable delivery, it will be using a protocol such as TCP which ensures that and will handle its own retransmissions.

This mode binds to a single interface and limits OpenVPN to only accepting IPv4 or IPv6 exclusively, not both at the same time.

> **Warning:** TCP mode is not compatible with DCO.

UDP on IPv4 and IPv6 on all interfaces (multihome)
This option binds to all interfaces using UDP for both IPv4 and IPv6. This mode allows OpenVPN to track incoming connections and respond back to clients using whichever IP address they used when connecting to the VPN. This works ideally with multiple WAN connections.

Since this mode can also work with both IPv4 and IPv6, clients of both types can be served by a single instance.

The primary downside of this mode is that since it binds to all interfaces, the port this instance uses cannot be used by any other service on the firewall. Also, since it does not bind specifically to a specific interface, return routing can be problematic in certain scenarios. Furthermore, outgoing connections will follow the routing table and cannot use specific interface addresses.

TCP on IPv4 and IPv6 on all interfaces (multihome)
Similar to the above, but using TCP instead of UDP.

Interface

Selects the interface, VIP, or failover group that the OpenVPN instance will use when communicating with peers. This also controls which interface the traffic from the instance will exit.

Note: The GUI hides this option if the **Protocol** is set to a multihome selection.

Several types of options are listed in the drop-down for **Interface**, and some have special behavior or use cases:

Interfaces
OpenVPN will bind to the interface address. If the interface is dynamic, such as DHCP, OpenVPN will automatically bind to the new address when it changes.

VIPs
OpenVPN will bind only to the specified VIP, which must be an IP alias or CARP type VIP.

Gateway Groups
For use with failover gateway groups. OpenVPN will bind to the address of the interface which contains the currently active gateway in the group. If that gateway becomes unreachable, it binds to the next one instead, and so on.

Localhost
Useful for multi-WAN deployments. Binding to localhost and utilizing port forwards to accept connections from several interfaces and/or ports is a versatile way to provide redundant OpenVPN connectivity for connecting clients.

Any
Binds to every address on every interface. Though tempting, this option is not recommended. When used with UDP, replies to Internet clients will always exit back out the default gateway WAN, which may be undesirable. If a use case calls for this type of behavior, consider selecting one of the **Protocol** options which uses multihome instead.

Local port

The port number upon which OpenVPN will listen for incoming connections from peers. Firewall rules must allow traffic to this port and this port must be specified in the client configuration.

The port for each server must be unique for each interface when using a standard UDP or TCP **Protocol** choice and must be globally unique if using a multihome **Protocol**.

Cryptographic Settings

This section controls how the VPN encrypts and validates traffic to and from peers.

Shared Key

> **Danger:** Shared key mode has been deprecated by OpenVPN as it is no longer considered sufficiently secure for modern requirements.
>
> **Shared key mode will be removed from future versions of OpenVPN**. Users **should not** create any new shared key tunnels and should **immediately** convert any existing shared key tunnels to SSL/TLS mode.
>
> When an SSL/TLS instance is configured with a `/30` tunnel network it behaves similarly to shared key mode. The primary difference is the need to create and distribute the certificate structure to peers. See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for information on configuring OpenVPN in SSL/TLS mode.

When using a shared key instance, either check the **Automatically generate a shared key** box to make a new key, or uncheck the box to paste in a shared key from an existing OpenVPN tunnel. When generating the key automatically, return to the edit screen for this tunnel later to obtain the key which may be copied to the remote router.

TLS Authentication

TLS, or Transport Layer Security, can provide session authentication and encryption to ensure the validity of peers and to protect the control channel.

> **Warning:** When using an SSL/TLS mode the best practice is to use TLS Authentication. In addition to the added security benefit from the key requirement, a TLS key also protects against some SSL-based attacks such as Heartbleed which could otherwise compromise the VPN using the control channel.

Use a TLS Key
: Check the box to make the VPN utilize a TLS key. When unchecked the GUI hides the remaining related options.

Automatically generate a shared TLS authentication key
: The GUI offers this option when there is no existing TLS key. Leave this checked, so the firewall will generate a new TLS key automatically when the instance is saved. Otherwise, uncheck the box to display the **TLS Key** field and paste in a TLS key.

  When generating the key automatically, return to the edit screen for this tunnel later to obtain the key, then copy the key to the remote peer.

TLS Key
: A text area containing the TLS key for this instance. This type of key is specific to OpenVPN and is not the same type of key used for an OpenSSL certificate.

  Tip: To generate a new key manually, use the command `openvpn --genkey secret /dev/stdout` from a shell command prompt (SSH, console, or GUI). On systems with older versions of OpenVPN, use `openvpn --genkey --secret /dev/stdout` instead.

TLS Key Usage Mode

TLS Authentication
In *TLS Authentication* mode OpenVPN uses the TLS key for HMAC authentication of control channel packets, protecting the peers from unauthorized connections.

TLS Encryption and Authentication
In *TLS Encryption and Authentication* mode OpenVPN uses the key for authentication, as above, but it also uses the key to encrypt control channel communication. This provides increased privacy and traffic control channel obfuscation.

Tip: In some cases this can also help OpenVPN avoid detection by network systems which identify OpenVPN traffic by its control packets.

TLS keydir Direction
The direction in which this VPN endpoint uses the TLS key. The TLS Key Direction must be set to complementary values on the server and clients. The default behavior uses `0` on server instances and `1` on client instances.

For example, if the server is set to `0`, clients must be set to `1`. Both peers may omit the direction in which case the TLS Key will be used bidirectionally.

Peer Certificate Authority

Select the certificate authority which signed the client or peer certificate(s) for this OpenVPN server instance.

Tip: If there are no entries in this list, first import or generate a certificate authority under **System > Certificates**, on the **CAs** tab.

Peer Certificate Revocation List

This optional field is for the Certificate Revocation List (CRL) this tunnel will use to check the validity of peer certificates.

A CRL is a list of certificates signed by a CA which are no longer valid. This could be due to a certificate being compromised or lost, such as from a stolen laptop, spyware infection, etc.

Tip: If there are no entries in this list, first import or generate a CRL at **System > Certificates**, on the **Certificate Revocation** tab.

OCSP Check

When set, OpenVPN will attempt to confirm client certificate validity using Online Certificate Status Protocol (OCSP) against the site listed in the **OCSP URL** field.

Server Certificate

The certificate used by the VPN instance to identify itself to peers. This certificate must contain appropriate properties marking it as a server certificate and not a user or client certificate.

Tip: If there are no entries in this list, first import or generate a certificate under **System > Certificates**, on the **Certificates** tab.

DH Parameters Length

The Diffie-Hellman (DH) key exchange parameters are used for establishing a secure communications channel. They may be regenerated at any time, and they are not specific to an OpenVPN instance. When importing an existing OpenVPN configuration these parameters do not need to be copied from the previous server nor does it require generating new parameters. The length of the desired DH parameters may be chosen from the drop-down box, either *1024*, *2048*, or *4096*. There is an additional choice for *ECDH Only* which disables traditional DH parameters and uses only ECDH.

Note: Due to the heavy computation involved in generating DH keys, the firewall contains a pre-generated set of keys for several common lengths as specified in RFC 7919. The DH group values specified in the RFC have been audited for security and are the safest values for use by end-users.

Should a use case require custom DH keys, generate new DH parameters manually by using the following shell commands:

```
# /usr/bin/openssl dhparam 2048 > /etc/dh-parameters.2048
# /usr/bin/openssl dhparam 3072 > /etc/dh-parameters.3072
# /usr/bin/openssl dhparam 4096 > /etc/dh-parameters.4096
```

ECDH Curve

Configures a specific elliptic curve to use for Elliptic Curve Diffie-Hellman key exchanges. This is only for use with ECDH TLS encryption.

OpenVPN uses the curve from the server certificate by default when configured with an ECDSA certificate. Otherwise, OpenVPN uses *secp384r1* as a fallback.

Encryption algorithm

These options define the cryptographic ciphers OpenVPN will use for this VPN.

Warning: At this time the only algorithm compatible with *OpenVPN Data Channel Offload (DCO)* is `AES-256-GCM`. When DCO is active these options are hidden to prevent invalid selections.

Data Encryption Negotiation
When set, OpenVPN will attempt to negotiate a compatible set of acceptable cryptographic data encryption algorithms from those selected in the **Data Encryption Algorithms** list.

This allows the client and server to agree on the most preferable cipher available without limiting the VPN to a single cipher. Legacy clients can still connect so long as the cipher they use is present in the list.

Note: Disabling this feature is deprecated. The option remains as a mechanism for legacy compatibility, but it should not be disabled unless absolutely necessary.

Data Encryption Algorithms

The list of Data Encryption Algorithms OpenVPN may use for this VPN, in order of preference. The default selection uses AES-GCM in 256 and 128 bit varieties as well as ChaCha20-Poly135.

The best practice is to use AEAD ciphers such as AES-GCM and ChaCha20-Poly135. These ciphers combine encryption and authentication and thus do not require a separate hash algorithm. Aside from offering strong security, they are typically much faster than other ciphers as well.

Note: This feature is only compatible with client/server mode, which means it only works with SSL/TLS modes where the tunnel network is large enough for multiple clients (e.g. larger than a `/30`).

In shared key mode or when using a `/30` tunnel network, OpenVPN only uses the value of **Fallback Data Encryption Algorithm**.

Fallback Data Encryption Algorithm

The Data Encryption Algorithm OpenVPN will use when it cannot negotiate an algorithm automatically. OpenVPN uses this value for shared key tunnels and for SSL/TLS configurations only capable of using a single client (`/30` tunnel network).

OpenVPN also uses this algorithm for older legacy clients which not only cannot negotiate a data encryption algorithm but have been compiled for a "small footprint", such as embedded devices.

The default is *AES-128-CBC*, which is AES 128 bit Cipher Block Chaining. This is a fine choice for most scenarios.

Auth Digest Algorithm

Selects the message digest algorithm OpenVPN uses for HMAC authentication of incoming packets. This is used for the data channel and if also for the control channel when the tunnel uses a TLS key. The GUI default of SHA256 is a good balance of security and speed.

When using AEAD ciphers such as AES-GCM, OpenVPN ignores this value for the data channel since AEAD ciphers already perform authentication. Even with an AEAD cipher, OpenVPN still uses this algorithm to authenticate the control channel when the tunnel uses a TLS key.

Note: OpenVPN defaults to SHA1 when this option omitted from its configuration. Unless both sides are set to a known value, use SHA1 here.

Certificate Depth

This option limits the valid length of a certificate chain. The default value limits a chain to *One (Client+Server)*. With that value, if an unauthorized intermediate CA signs a certificate, certificates signed by the rogue intermediate would fail validation. In cases when the certificate stricture requires chaining with intermediates, raise this limit to accommodate the longest allowed chain.

Strict User-CN Matching

Controls whether the firewall will enforce a strict match between the username supplied by the user and the Common Name of their user certificate when the firewall authenticates a user. When enabled, authentication fails if the two values do not match.

This prevents users from using their own credentials with the certificate from a different user and vice versa.

Client Certificate Key Usage Validation

When set, the authentication process verifies that a certificate supplied by a client contains the appropriate certificate properties to act as a client. This means that the certificate must include the extended key usage attribute for "TLS Web Client Authentication".

This prevents using certificates made for different purposes, such as e-mail signing or acting only as a server, from being used as a VPN client certificate.

Tunnel Settings

The tunnel settings section governs how traffic flows between the server and clients, including routing and compression.

IPv4/IPv6 Tunnel Network

These are the subnets used by the VPN. The exact usage depends on the mode and size of the subnet. The firewall uses these addresses for direct communication between tunnel endpoints even when connecting two existing remote networks. The value must be an unused, non-overlapping subnet. It must not be in use locally or at any remote site. One or both of **IPv4 Tunnel Network** and **IPv6 Tunnel Network** may be entered, or in the case of a tap bridge, neither.

For SSL/TLS modes with subnets large enough for multiple clients (e.g. **IPv4 Tunnel Network** is larger than `/30`), OpenVPN uses a client/server mode. In this situation, these values are the pools of addresses the OpenVPN server assigns to clients. The server end of the OpenVPN configuration will use the first address in this pool for itself, and assign additional addresses to connected clients as needed. The specific method of assignment varies based on the chosen value for **Topology**. In OpenVPN client/server mode the server can push settings to clients, and client-specific overrides can influence how clients behave. For site-to-site VPNs in this mode, an override must contain **IPv4/6 Remote Network/s** values which route subnets to specific clients (`iroute`).

For a site-to-site SSL/TLS server using IPv4 and an **IPv4 Tunnel Network** value of `x.x.x.x/30`, or a shared key server, OpenVPN uses peer-to-peer mode. In peer-to-peer mode, each server can only have one client. In this mode, clients do not require client-specific overrides or `iroutes` configuration, however, the server also cannot push routes or settings to the client.

Warning: These specific types of peer-to-peer modes are not compatible with *OpenVPN Data Channel Offload (DCO)*.

This field also accepts the name of a network type alias, but it must only contain a single entry.

See also:

See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for more information on a site-to-multi-site example using a large tunnel network and `iroutes`.

Bridging Options

When using *tap* mode for a remote access SSL/TLS VPN, the GUI offers additional options to control bridging behavior in OpenVPN and client address assignment.

See also:

Bridging OpenVPN Connections to Local Networks

Bridge DHCP
When selected, OpenVPN passes through DHCP to the bridged. In the most common scenario, this is *LAN*.

Using this method connecting clients receive IP addresses from the same DHCP pool used by clients on the LAN.

Bridge Interface
This setting indicates to OpenVPN which interface will be used for the bridge. In most cases this is *LAN*.

> **Warning:** This does not create a bridge in the operating system, that must be done separately.

This option controls which existing IP address and subnet mask are used by OpenVPN for the bridge. Setting this to *none* will cause the **Server Bridge DHCP** settings below to be ignored.

Bridge Route Gateway
Makes OpenVPN push the **Bridge Interface** IPv4 address to connecting clients as a route gateway.

When the **IPv4 Tunnel Network** in OpenVPN is empty for a bridged VPN, connecting clients cannot automatically determine a server-side gateway for use with routes pushed to clients by the server.

Server Bridge DHCP Start/End
Defines a DHCP range from which OpenVPN can assign addresses to clients. If these settings are left blank, OpenVPN passes DHCP through to the bridge interface, and it ignores the interface setting above.

When set, this range should be within the same subnet as the **Bridge Interface**, but it should not overlap an existing in-use portion of the subnet, such as the current DHCP pool.

This allows a range of IP addresses to be set aside for use only by OpenVPN clients so they may be contained within a portion of the internal network rather than consuming IP addresses from the existing DHCP pool.

Redirect IPv4/IPv6 Gateway

When a **Redirect IPv4/IPv6 Gateway** option is selected the server pushes a message to clients instructing them to forward *all* traffic for that address family, including Internet traffic, over the VPN tunnel.

This option only works in SSL/TLS client/server modes (tunnel network larger than `/30`).

IPv4/IPv6 Local network(s)

These fields specify local networks reachable by VPN clients, if any. Entries must be in CIDR or prefix format (e.g. `192.168.56.0/24`) and they can be a host or network type alias name. The server pushes a route for these networks to clients when they connect.

To specify multiple subnets of a particular address family, enter the subnets separated by a comma, e.g. `192.168.2.0/24, 192.168.56.0/24` or use an alias.

This function relies upon the ability to push routes to the client. For IPv4 it is only valid in SSL/TLS client/server mode with a tunnel network larger than `/30`. It will always work for IPv6 provided the VPN does not use a similar too-small prefix.

Warning: If the server does not push any routes, clients will not receive a remote gateway value, which they may require. For example, pfSense® software will use the remote gateway value when creating a gateway for gateway monitoring and policy routing. If the server does not need to push any routes to the client, use a custom option to push the gateway value to clients, for example: `remote-gateway x.x.x.1`; where the IP address is the IP address of the tunnel on the server.

IPv4/IPv6 Remote Network

This option only appears for Peer-to-Peer type connection, and it is not available for remote access servers.

OpenVPN adds operating system route table entries for the specified subnets which hand the traffic over to this OpenVPN instance for processing. Entries must be in CIDR or prefix format (e.g. `192.168.56.0/24`) and they can be a host or network type alias name.

To specify multiple subnets of a particular address family, enter the subnets separated by a comma, e.g. `192.168.2.0/24, 192.168.56.0/24` or use an alias.

For Peer-to-peer SSL/TLS servers in client/server mode (tunnel network larger than `/30`), each network listed in this field **must** also be present in a *client specific override* as a remote network entry for a client. Without that entry, OpenVPN has no way to determine which client should receive traffic for a network. The remote networks listed here in the server configuration inform the operating system routing table to deliver the traffic to OpenVPN, while the entries in an override associate networks with specific remote clients.

Shared key and SSL/TLS servers in peer-to-peer mode (`/30` tunnel network) do not need overrides as there can only be a single client per server.

Concurrent Connections

Specifies the number of clients that may be simultaneously connected to this OpenVPN server instance at any given time.

> **Warning:** This is a collective limit for **all** connected clients, not a per-user setting.

Compression

> **Warning:** Compression for encrypted traffic is dangerous and should be disabled when possible!
>
> The OpenVPN project has deprecated compression, and they will remove it from future versions.
>
> Read the text in this section carefully! Compression can potentially increase throughput but may allow an attacker to extract secrets if they can control compressed plain text traversing the VPN (e.g. HTTP). Before enabling compression, consult information about the VORACLE, CRIME, TIME, and BREACH attacks against TLS to decide if the use case for this specific VPN is vulnerable to attack.

When compression is enabled OpenVPN attempts to compress traffic crossing VPN before it performs encryption. This has a potential to save bandwidth usage for some traffic at the expense of decreased security and increased CPU utilization on both the server and client.

For high speed connections such as the usage of OpenVPN across a LAN, high speed low/latency WAN, or local wireless network, compression adds unnecessary overhead as the delay added by the compression may be more than the delay saved in transmitting the traffic.

If nearly all the traffic crossing the OpenVPN connection is already encrypted (such as SSH, SCP, HTTPS, among many other protocols), do not enable LZO compression because encrypted data is not compressible and the LZO compression will cause slightly more data to be transferred than would be without compression. The same is true if the VPN traffic is almost entirely data that is already compressed.

Allow Compression
: Controls whether compression is allowed on the VPN, and how it is handled.

  Decompress incoming, do not compress outgoing
  : Asymmetric compression. Accepts compressed packets from the remote peer but will not perform any compression on outgoing packets. This behavior allows for a smoother transition when peers use older versions of OpenVPN.

  Refuse any non-stub compression
  : When set, OpenVPN will refuse any attempt to compress packets. Compressed packets from peers are dropped.

  Compress packets
  : When set, OpenVPN allows compression and both accepts compressed packets from peers and also compresses packets before sending them to peers.

Compression
: This selector controls the handling of LZO compression for this OpenVPN instance when the settings allow compression. There are several possible settings each with slightly different behavior.

  Disable Compression
  : Omits compression directives from the OpenVPN configuration entirely. OpenVPN

will not perform compression, but other methods such as Client-Specific overrides or advanced options may override this behavior.

Disable Compression, retain compression packet framing
OpenVPN will not perform compression, but keeps the packet framing for compression. This allows compression to be changed later, for example by pushed settings. When set this way, OpenVPN will advertise LZ4 and LZO compression functionality to peers using `IV_` variables. Always adds one extra framing byte to packets for compression.

Enable Compression (stub)
Similar to the previous setting, but it does not advertise compression to peers.

Enable Compression (stub v2)
Similar to the previous setting, but it utilizes a better framing that does not increase packet overhead when it cannot compress a packet.

LZ4 Compression
A newer compression algorithm offering the best compression performance and lowest CPU consumption. May not be compatible with older clients.

LZ4 Compression v2
The same as LZ4 compression but it utilizes a better framing that does not increase packet overhead when it cannot compress a packet.

LZO Compression
An older compression algorithm with wider compatibility, especially with older clients.

Omit Preference + Disable Adaptive LZO Compression
Does not specify a compression directive in the OpenVPN configuration and also explicitly disables adaptive compression.

Adaptive LZO Compression
OpenVPN will periodically check the efficiency of data compression for VPN traffic and disable compression if it is performing poorly. This prevents OpenVPN from compressing already compressed or encrypted data.

LZO Compression (Legacy Style)
Enables LZO compression using the deprecated `comp-lzo yes` directive. This can help interoperation with older legacy clients which cannot be updated. When present in the configuration, the server can push a different value to clients as needed.

No LZO Compression (Legacy Style)
Disables LZO compression using the deprecated `comp-lzo no` directive. This can help interoperation with older legacy clients which cannot be updated. When present in the configuration, the server can push a different value to clients as needed.

Push Compression
When set, OpenVPN will push the selected compression settings to connecting clients.

Type-of-Service

When this option is enabled OpenVPN sets the Type-of-Service (TOS) IP header value of tunnel packets to match the encapsulated packet value. This may cause important traffic to be handled faster over the tunnel by intermediate hops at the cost of minor information disclosure.

The most common example is VoIP or video traffic. If the TOS bit is set to reflect the priority of the traffic it can help QoS along the path, but someone intercepting the traffic could see the TOS bit and gain some knowledge about the contents of the traffic inside the tunnel. For those who rely on TOS bits for QoS, the benefit may outweigh the information leak.

Inter-Client Communication

This option controls whether connected clients are able to communicate with one another. To allow this behavior, check the option. When unchecked, clients can only send traffic to the server or destinations beyond the server such as routed networks or the Internet.

Typically, in remote access style deployments it is unnecessary for clients to reach each other, but there are use cases when it can be helpful. One example is remote web developers working together and running test servers on their local workstations. With this option activated, the developers can reach the other self-hosted test servers for collaborative development.

Duplicate Connections

Controls whether OpenVPN will allow multiple connections from the same user to work simultaneously.

This is primarily for security reasons, so the same certificate cannot be used by multiple people or devices simultaneously. The best practice is to use a unique certificate for each connecting device, or at least for each user. Otherwise, if a client is compromised there is no way to revoke that one client alone; certificates must be reissued to all clients that share the same certificate.

By default, OpenVPN will associate an IP address from its tunnel network with a specific certificate or username for a given session. If the same certificate connects again, the server assigns it the same IP address and will either disconnect the first client or cause an IP conflict where neither client will receive proper data.

If a setup that uses the same certificate in multiple locations is an absolute requirement and cannot be avoided, check **Duplicate Connections** to allow the non-standard behavior of multiple clients using the same certificate or username.

Duplicate Connection Limit

When active, the server will limit the number of simultaneous connections from the same client. New connections from the same client that exceed this limit will be denied by the server. To free up unused sessions and allow new connections from the same client more quickly, set lower values for **Ping settings**.

Note: Due to unstable connectivity in practice, clients may sometimes initially create multiple sessions on the server when establishing the tunnel. These extra temporary sessions will count towards the limit until they time out.

Client Settings

These settings control behavior of clients connecting to this sever.

Dynamic IP

Checking this box adds the `float` configuration option to the OpenVPN configuration. This allows clients to retain their connection if their IP address changes, similar to MOBIKE for IKEv2 in IPsec.

For clients on Internet connections where the IP address changes frequently, or mobile users who commonly move between different Internet connections, check this option for more stable connectivity. Where the client IP address is static or rarely changes, not using this option offers a small security improvement.

Topology

Sets the method OpenVPN uses to allocate addresses for clients in a client/server setup on *tun* device mode VPNs. The **Topology** option is relevant only when supplying a virtual adapter IP address to clients using *tun* mode on IPv4. Some clients may require this even for IPv6, such as OpenVPN Connect, though in reality IPv6 always runs with a subnet topology even when IPv4 uses *net30*. OpenVPN instances using *tap* mode always use *subnet* topology as well.

subnet
: Uses the first IP address in the subnet for the server and allocates one IP address per client in a single shared subnet.

  By default, OpenVPN on pfSense® software prefers a topology style of *subnet* when using a **Device Mode** of *tun*. This is the only available style when using the *tap* **Device Mode**.

Note: Some ancient clients may not implement this mode on certain platforms, such as before 2.3.x which was over 12 years ago. These clients are rare in modern environments.

Always make sure the client and associated drivers are fully up-to-date when using a *subnet* topology.

---

net30
: OpenVPN allocates a `/30` CIDR network (four IP addresses, two usable) to each connecting client, including one for the server itself. This style has a longer history, but can be confusing for administrators and users alike.

> **Warning:** The OpenVPN project has declared the *net30* style as deprecated, indicating it will be removed in future versions. Avoid using it when possible.

Ping Settings

Inactive

The amount of time, in seconds, which a client can be inactive before OpenVPN disconnects it for inactivity. OpenVPN bases activity on the last incoming or outgoing data channel packet, not control channel packets.

The default value is `300`. A value of `0` disables this feature.

Tip: For Peer-to-Peer SSL/TLS servers in client/server mode, the best practice is to set this to `0` so that site-to-site VPN tunnels stay up indefinitely.

Warning: This option is ignored in peer-to-peer modes, such as shared key mode and SSL/TLS with a blank or `/30` tunnel network. In those cases the option can cause the process to exit and not restart, resulting in a loss of service.

This option is not compatible with *OpenVPN Data Channel Offload (DCO)*.

Ping Method

The **Ping Method** controls OpenVPN monitoring of peers through the control channel and how it deals with unresponsive peers. There are two methods available: **Keepalive** and **Ping**:

Use Keepalive Helper
This method uses the **Interval** and **Timeout** values to automatically set common useful values for OpenVPN `ping` and `ping-restart` rather than defining behavior manually. The values are used locally and pushed to peers when possible.

Note: This is the best practice in nearly all use cases, as most environments do not necessitate the extra complexity of configuring the behavior manually.

Interval
Sets the interval, in seconds, between control channel pings as well as the idle period for the data channel before OpenVPN will send a control channel ping. The default value is `10` seconds.

Note: All peers must send pings at the expected intervals as OpenVPN does not echo responses.

Timeout
The amount of time, in seconds, OpenVPN will wait for a ping from a peer before it considers the peer to be down. The default value is `60` seconds.

In client/server mode, on the server this value is multiplied by `2`, and it disconnects an individual session for a client; on the client the value is used as-is, and it restarts the VPN process. In peer-to-peer mode this restarts the VPN process.

Define Ping Manually
This method offers more flexibility in how OpenVPN will send pings and expect responses from peers, but it is also more complicated.

Ping
Sets the interval, in seconds, between control channel pings as well as the idle period for the data channel before OpenVPN will send a control channel ping. The default value is `10` seconds.

Push ping to client
Controls whether the value of **Ping** is pushed to clients when OpenVPN is in client/server mode.

Ping restart or exit
Chooses between whether the OpenVPN process will restart on failure or exit.

ping-restart
In client/server mode, on a server this disconnects a client session when the client does not respond. In client/server mode on the client, as well as in peer-to-peer mode, it restarts the OpenVPN process when a peer fails to respond.

ping-exit
Causes the VPN process to exit entirely when a peer fails to respond.

> **Warning:** The VPN will not recover automatically when using this option. It requires manual intervention to start the VPN process again.

Ping restart or exit seconds
The amount of time, in seconds, OpenVPN will wait for a ping from a peer before it considers the peer to be down. The default value is `60` seconds.

Push to client
Controls whether **Ping restart or exit** and its associated value are pushed to clients when OpenVPN is in client/server mode.

Advanced Client Settings

DNS Default Domain

Configures a default domain name which clients will append to DNS requests. This can be helpful to ensure name resolution works properly for hosts on the local network where DNS name resolution is used.

For Microsoft Active Directory environments, this would usually be the Active Directory domain name.

DNS servers

When checked, the GUI allows configuring up to four DNS servers for use by clients while connected to the VPN.

For Microsoft Active Directory environments, this is typically the Active Directory Domain Controllers or DNS servers for proper name resolution and authentication when connected via OpenVPN.

Block Outside DNS

Makes Windows 10 clients block access to DNS server except across OpenVPN while connected, forcing clients to use only VPN DNS servers.

This is only relevant on Windows 10 clients using OpenVPN version 2.3.9 and later as they are the only clients prone to leak DNS requests in this way. The option has no effect on other platforms and they will ignore the directive.

Force DNS Cache Update

When checked, the OpenVPN server pushes a set of commands to Windows clients which flush and restart DNS caching to improve client handling of updated DNS servers from the VPN.

NTP servers

When checked, the GUI allows configuring one or two NTP servers which OpenVPN will push to clients for time synchronization. These values can be an IP address or FQDN.

NetBIOS Options

The **Enable NetBIOS over TCP/IP** option controls whether the GUI displays several other NetBIOS and WINS related options.

Node Type

The NetBIOS node type controls how Windows systems function when resolving NetBIOS names. The best practice is to leave this set to *none* which accepts the default value from Windows.

The available options include:

b-node
: Use broadcasts for NetBIOS name resolution. This would only be used in the case of a *tap* bridge as otherwise OpenVPN does not transmit or receive broadcast messages.

p-node
: Point-to-point name queries to a WINS server. WINS has been deprecated on modern networks, so this option is not useful in most Windows networks.

m-node
: Broadcast then query name server. Similar to *b-node* but will fall back to DNS.

h-node
: Query name server first, then use broadcast. This option is the most likely to succeed in a current network with proper, functional, DNS.

Scope ID

A NetBIOS Scope ID provides an extended naming service for NetBIOS over TCP/IP. The NetBIOS scope ID isolates NetBIOS traffic on a single network to only those nodes with the same NetBIOS scope ID.

WINS Servers

Checking this box allows defining two WINS servers which provide name resolution for clients accessing and browsing NetBIOS resources across the VPN. WINS has been largely deprecated and removed from use, so it is unlikely that most modern environments would benefit from this behavior.

Advanced Options

Custom options

While the GUI offers many commonly used options, OpenVPN contains many more options that are unavailable in the GUI which certain use cases may require.

Custom options may be added in using the **Custom option** box with **each directive separated by a semicolon** (``;``). The firewall will pass the custom directives to OpenVPN.

See also:

These options are described further in *Custom Configuration Options*.

> **Warning:** Use with extreme caution. Due to the nature of how this field operates, the firewall cannot validate its contents. Invalid combinations of directives will cause the OpenVPN instance to fail.

Username as Common Name

Controls whether OpenVPN will use the username given by the client in place of the certificate common name for purposes such as determining Client Specific Overrides. This is only relevant when user authentication is enabled. This is typically the best practice, but not a requirement.

UDP Fast I/O

Controls whether OpenVPN will use fast I/O operations with UDP writes to its *tun* or *tap* device. This behavior optimizes the packet write event loop, improving CPU efficiency by 5% to 10%.

> **Warning:** OpenVPN considers this option experimental, and it may not be implemented on all platforms. This option is not compatible with OpenVPN bandwidth limiting.
>
> This option is also not compatible with *OpenVPN Data Channel Offload (DCO)*.

Exit Notify

Controls whether OpenVPN will send an explicit exit notification to connected UDP clients or peers when restarting or shutting down. This notification allows peers to immediately disconnect rather than wait for a timeout. This is only relevant to UDP modes as TCP handles this naturally on its own.

In SSL/TLS Server modes, clients may be directed to reconnect or use the next server.

Disabled
: Does not send an exit notification.

Reconnect to this server / Retry Once
: In server mode this directs clients to reconnect to the same server. This is useful if there is only one server, and it will be available again shortly.

  For clients, it directs them to retry sending the notification once before giving up.

Reconnect to next server / Retry Twice
: In server mode this directs clients to reconnect to the **next** server if the client configuration contains multiple servers. This is useful to nudge clients to an alternate server if this server could be down for an extended period.

  For clients, it directs them to retry sending the notification twice before giving up.

Warning: The firewall ignores this option in Peer-to-Peer Shared Key mode and in SSL/TLS mode with a blank or `/30` tunnel network as it will cause the server to exit and not restart.

This option is not compatible with *OpenVPN Data Channel Offload (DCO)*.

Send/Receive Buffer

Configures a Send and Receive Buffer size for OpenVPN. The default buffer size can be too small in many cases, depending on hardware and network uplink speeds. Finding the best buffer size can take experimentation. To test the best value for a site, start at *512KiB* and test higher and lower values until testing results in peak performance.

Note: For remote access VPNs this may take experimentation with multiple types of clients on different devices, networks, and so on.

Warning: This option is not compatible with *OpenVPN Data Channel Offload (DCO)*.

Gateway Creation

Controls which types of gateways the firewall will automatically create for this VPN instance when assigned as an interface. The default behavior will create *both* IPv4 and IPv6 gateways but if the VPN will only ever carry one type of traffic, this option can limit that behavior, so the GUI will not display an unnecessary gateway entry.

Verbosity level

Configures the amount of detail OpenVPN will log for this instance, which is useful for troubleshooting problems. Higher numbers will result in higher amounts of detail in the log. During normal operation the *default* selection is ideal.

Note: When set to higher levels, the OpenVPN status page and dashboard widget will cause additional logging as they interact with the Management process to poll information from the OpenVPN daemons.

Client Configuration Options

These options are available in one or more modes for OpenVPN client instances, managed from **VPN > OpenVPN**, on the **Clients** tab.

Many of these options are identical to the server options mentioned in *Server Configuration Options*. This section only notes the differences.

Server mode

For client instances, the only server mode choices are *Peer to Peer (SSL/TLS)* and *Peer to Peer (Shared Key)*. These choices pair with the server options of the same name and type.

> **Danger:** Shared key mode has been deprecated by OpenVPN as it is no longer considered sufficiently secure for modern requirements.
>
> **Shared key mode will be removed from future versions of OpenVPN**. Users **should not** create any new shared key tunnels and should **immediately** convert any existing shared key tunnels to SSL/TLS mode.
>
> When an SSL/TLS instance is configured with a `/30` tunnel network it behaves similarly to shared key mode. The primary difference is the need to create and distribute the certificate structure to peers. See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for information on configuring OpenVPN in SSL/TLS mode.

Interface

This option selects the interface, VIP, or failover group that the OpenVPN client instance will use for outgoing connections.

When a CARP type VIP is selected for the Interface on OpenVPN Client instances, the firewall will stop the OpenVPN instance when the CARP VIP is in a backup state. This prevents the secondary HA node from maintaining invalid routes or attempting to make outbound connections which can interfere with the active connection on the primary HA node.

Note: To ensure that a connection uses out the correct WAN when selecting a WAN which is not the default gateway, add a static route for the server IP address. Depending on the protocol in use and contents of the routing table, the firewall may not be able to send traffic out the correct WAN without a static route.

Local Port

For clients, the local port should be blank in nearly every case so that OpenVPN will use a randomized local port. This behavior is more secure, but some server configurations may require a specific source port.

Server host or address

The IP address or fully qualified domain name for the server.

Note: When using a hostname for the remote server address, OpenVPN will resolve the server host name on each connection attempt.

Server Port

The port on which the server is listening, typically `1194`.

Proxy Settings

Proxy Host or Address
: The IP address or fully qualified domain name for a proxy server through which this client must connect.

Proxy Port
: The port on which the proxy is listening for connections.

Proxy Auth Extra Options
: Extra authentication options. When set to *basic* or *ntlm* the GUI presents **Username** and **Password** fields to configure proxy authentication.

User Authentication Settings

Configures authentication option for SSL/TLS mode. This may be optional, depending on the server configuration.

Username and Password
: The user credentials to send, if required by the server.

Authentication Retry
: When set, OpenVPN will not retry a connection when authentication fails; the OpenVPN process will exit if it receives an authentication failure message instead. The default behavior is to retry authentication.

Cryptographic Settings

The settings in this section are identical to those on their corresponding options on the server side except for the new **Client Certificate** option. This option sets the certificate for use by this client.

Note: The client certificate, its key, and the associated CA certificate must all be imported to the firewall using the certificate manager before OpenVPN can use them.

Shared Key / TLS Authentication

These options work similar to the server side counterparts, but be aware that the key from the server **must be copied here exactly**. Do not generate a new key on the client if the server already has a key.

Limit Outgoing Bandwidth

This option limits the speed of outgoing VPN traffic to the given amount, specified in **bytes per second**. The value must be either empty or between `100` and `100000000`.

OpenVPN will not limit traffic when the field is empty.

Don't Pull Routes

When checked, the client ignores routes pushed from the server. This is useful in cases when the server pushes a default gateway redirect when this client does not need one, or if the server pushes routes for networks that this client prefers to handle in other ways.

Don't Add/Remove Routes

When checked, OpenVPN will not manage route table entries for this VPN. In this case, the routes must be managed manually. Routes that OpenVPN would normally add are instead passed to `--route-up script` using environmental variables.

Pull DNS

If this option is set, the firewall will use DNS servers assigned by the remote OpenVPN server for its own purposes, similar to if it had received a DNS server from a dynamic WAN.

Client Specific Overrides

Client specific overrides define custom settings which apply to only certain clients connecting to an SSL/TLS server in client/server mode. These settings are determined by the way a connecting client authenticates, either by their username or the common name of their certificate.

These options are additive, meaning the client will receive the options pushed by the server configuration and then the options defined on this client-specific override entry. Some settings will override the values received from the server, such as the Tunnel Networks, while others will combine and use both server and local values, such as Local Network definitions.

Tip: The "Prevent this client from receiving any server-defined client settings." option will make the client ignore any pushed settings from the server which are not defined on the override entry.

Client specific overrides are managed at **VPN > OpenVPN** on the **Client Specific Overrides** tab.

Purpose

Depending on the use cases, overrides may be a required part of a deployment, such as a site-to-site VPN with multiple clients connecting to a single server.

Overrides also enable special behavior such as configuring different routes for different clients, static IP addresses when connecting to the VPN, exceptions to default VPN behaviors, and many more scenarios.

Configuration

The following settings are available when configuring client specific overrides.

Description
Text describing the override entry, such as a user or site name, or its purpose.

Disable
Whether this override is enabled. When checked, the override is not active.

Common Name
The name of the user which OpenVPN will match when a client connects.

When using SSL/TLS authentication alone this matches the common name field of the certificate.

When using user authentication alone, the common name is undefined by default and will not match anything unless the **Username as Common Name** option is enabled on the OpenVPN server. With that set, this field will match the username.

With SSL/TLS and user authentication both are active, the behavior is determined by the **Username as Common Name** option on the OpenVPN server. When checked, it matches the username. When unchecked, it matches the common name of the certificate. If multiple users share the same certificate, check that option.

The special name `DEFAULT` will trigger if the connecting user does not match any other existing overrides. This can be useful when adding options which are only possible on overrides, such as **Connection Blocking**.

Connection Blocking
When set, OpenVPN rejects clients matching this override. This can be used as an additional means of blocking a specific user, though this use case should only be temporary.

Note: Do not use this to permanently block a specific user if their credentials have been compromised or terminated. Instead, use a CRL to revoke the client certificate or make changes to the user account such as removing the account or changing its password.

Some administrators use this option to selectively allow clients on specific servers when they share a common CA structure. The best practice is to use a separate certificate structure for each VPN server, but that is not always possible.

Tip: To only allow users with an override to connect, create an override for the `DEFAULT` common name with this option set. With that in place, OpenVPN will reject all clients by default and will allow connections from clients matching other overrides which do not have this option set.

Server List
The OpenVPN servers for which this override will be active. Select one or more OpenVPN instances

which will utilize this override. By default, with no entries selected, overrides are active for all servers.

Many use cases may call for using an override with only a specific VPN. For example if a client has a static IP address inside the VPN, that address may only be valid on one VPN server. Attempting to use that static address on a different VPN will fail.

Note: Multiple overrides can use the same common name, but such entries are only useful if they are active on different VPN instances.

IPv4/IPv6 Tunnel Network
A specific static virtual IPv4 network or network type alias with a single entry used for private communications between this client and the server expressed using CIDR notation (e.g. `10.0.8.5/24`).

On a server with *subnet* topology, or for IPv6, the client IP address and the subnet/prefix mask must match the **Tunnel Network** on the server.

On a server with *net30* topology, OpenVPN assumes the first IPv4 network address of the `/30` is the server address, and it assigns the second network address to the client.

Note: These options express a preference, **not** a reservation. OpenVPN dynamically allocates addresses to connecting clients near the start of its tunnel network. An override will not prevent another client from using this address if it happens to be assigned the address dynamically.

IPv4/IPv6 Local Network/s
Networks located on the **server** side for which OpenVPN will push routes to this client. This can be a comma-separated list of networks in CIDR notation, and it can also be a host or network type alias.

Note: This is functionally identical to the same fields on the OpenVPN server configuration. Values in these fields are added to the lists pushed by the server, so there is no need to duplicate those values here. Only exceptions or differences from the default entries should be in these fields.

IPv4/IPv6 Remote Network/s
Networks which can be reached **through this client**. OpenVPN will internally route traffic destined for these networks to this client (`iroute`).

The remote networks listed in the *server* configuration inform the operating system routing table to deliver the traffic to OpenVPN, while the entries in an override associate networks with specific remote clients.

> **Warning:** Unlike local networks, entries in the remote network fields **must** be set in **both** the OpenVPN server **and** an override. For full routing functionality to a client network, both the operating system routing table and OpenVPN must know how to reach the network.

Redirect Gateway
When set, OpenVPN pushes a default gateway to the client so it will send all of its traffic, including Internet traffic, through this VPN.

Server Definitions
When set, OpenVPN will not push options from the server configuration to this client.

Remove Server Routes
When set, OpenVPN will not push routes to this client, but it will push other options.

DNS Default Domain
When set, the GUI presents a field in sets an alternate default DNS search domain which OpenVPN will push to this client.

DNS Servers
When set, the GUI presents four fields for alternative DNS servers OpenVPN will push to this client.

NTP Servers
When set, the GUI presents two fields for alternative NTP servers OpenVPN will push to this client.

NetBIOS Options
When set, the GUI presents fields for alternate NetBIOS options which OpenVPN will push to this client.

Advanced
Additional custom options for OpenVPN to apply to this client.

Each directive must be separated by a semicolon (`;`).

See also:

These options are described further in *Custom Configuration Options*.

Custom Configuration Options

OpenVPN offers a vast array of configuration options, many more than the most commonly used fields in the GUI. The **Custom Options** box enables using directives in OpenVPN which are not available directly in the GUI.

> **Warning:** Each directive **must** be separated by a semicolon (`;`).

This section covers a few custom options users have found useful, but which are not common enough to add to the GUI. There are many more, and the OpenVPN man page details them all.

> **Warning:** Exercise caution when adding custom options. The GUI cannot perform input validation on directives in this field. If an option is incorrect or invalid, the OpenVPN instance may not start.
>
> View the OpenVPN logs under **Status > System logs** on the **OpenVPN** tab to ensure the options used are valid. Any invalid options will result in a log message, followed by the option that caused the error:
>
> ```
> Options error: Unrecognized option or missing parameter(s)
> ```

Additional Servers

The `remote <address> <port> <protocol>` directive specifies servers to which a client can connect. This is primarily used on client instances to define multiple servers for redundancy. Clients will try the server defined in the GUI settings first and then any additional servers in the order given.

The address can be an IP address or FQDN. The port number defaults to `1194` and may be omitted if it is the default. The protocol can be either `udp` or `tcp`, and optionally can end in `4` or `6` to limit an FQDN server to either IPv4 or IPv6 respectively, if DNS contains records for both.

There are two primary strategies for which administrators use this type of configuration. One is for redundancy between multiple servers, and the other is for redundancy between multiple ports. The latter can be important for working around limitations on client networks, such as networks which only allow specific outbound ports. The two strategies can also be mixed as needed.

To specify multiple additional servers, consider a set of entries such as this:

```
remote vpn2.example.com 1194 udp4;
remote vpn3.example.com 1194 udp4;
remote vpn4.example.com 1194 udp4;
```

Contrast this with a set of servers which are crafted to work around network limitations:

```
remote vpn1.example.com 53 udp;
remote vpn1.example.com 123 udp;
remote vpn1.example.com 443 tcp;
remote vpn1.example.com 80 tcp;
remote vpn1.example.com 8080 tcp;
```

Another way the `remote` directive can be used is in the custom options of a peer-to-peer server instance (shared key or SSL/TLS with a tunnel network of `/30`). If both peers are defined as a server and each has a `remote` directive pointing to the other, then they will attempt connections in both directions and whichever connects first is used. This operates closer to IPsec where both peers can initiate. In practice this is not very useful as it's typically better to have one designated initiator, but there may be a rare use case which calls for this behavior.

The `remote-random` directive tells clients to connect to a random server from the list instead of the next available choice. This is particularly useful when trying to load balance clients between multiple servers, such as with public VPN providers.

Renegotiation Time

The `reneg-sec <seconds>` directive controls how often OpenVPN renegotiates authentication with clients. The default time is `3600` seconds (one hour). In most cases the clients renegotiate and continue on without interruption, however with multifactor authentication (MFA) this can disrupt clients. With MFA, clients would need to utilize a fresh token each time OpenVPN renegotiates the VPN, which can range from inconvenient to impossible. With MFA mechanisms such as OTP or Google Authenticator there is no mechanism to supply a new code, so the VPN disconnects after an hour and the client must manually reconnect.

In these cases, it is common for administrators to disable the renegotiation:

```
reneg-sec 0;
```

This is less secure, but more convenient than forcing users to reauthenticate once per hour.

Alternately, the time limit can be raised to a higher value which is less inconvenient, such as `28800` (8 hours) for a typical workday, or `86400` (24 hours) to make it once per day.

Recursive Routing

The `allow-recursive-routing` directive allows OpenVPN to send non-OpenVPN traffic to the VPN server itself over the VPN. Certain rare use cases call for this behavior where the VPN server and a public service are both hosted on the same server, but portions of the public service are only available to clients connecting over the VPN. OpenVPN used to allow this by default, but now it must be explicitly enabled by this directive.

Control Pushed Options

The `push-remove <name>` directive selectively filters options pushed by OpenVPN servers. This allows clients to ignore certain directives that would normally be sent by servers, such as routes (`route` or `route-ipv6`), `keepalive/ping` values, compression options, default gateway (`redirect-gateway def1`), DNS options, and more.

In a client-specific override context this can be used to skip pushing certain items to a specific client and then supply a new value in its place.

Routing options

Note: Using the custom option box is not necessary to add most routes. To add additional routes for a particular OpenVPN client or server, use the **Local Network** and **Remote Network** boxes, which allow multiple networks using comma-separated lists.

The `route` custom configuration option adds routes locally for networks that are reachable *through* the VPN, but is not necessary in most cases as the GUI **Remote Network** fields for IPv4 and IPv6 accomplish the same goal. Some users prefer to enter the routes in this box instead, however. It can also be useful for cases where the routing is ambiguous, such as in bridged VPNs, to manually define specific routes with gateways that cannot be automatically determined by OpenVPN.

The following example adds a route for `10.50.0.0/24`:

```
route 10.50.0.0 255.255.255.0;
```

To add a route with a specific gateway, add it after the netmask:

```
route 10.50.0.0 255.255.255.0 10.0.1.1;
```

To add multiple routes, separate them with a semicolon:

```
route 10.50.0.0 255.255.255.0;
route 10.254.0.0 255.255.255.0;
```

An OpenVPN server configuration using SSL/TLS in client/server mode can push additional routes to clients. The GUI can configure these using the **Local Network** field. To push the routes manually for `10.50.0.0/24` and `10.254.0.0/24` to all clients, use the following custom configuration option:

```
push "route 10.50.0.0 255.255.255.0";
push "route 10.254.0.0 255.255.255.0";
```

Note: Note the placement of the double quotes in these directives.

Redirecting the default gateway

OpenVPN can also redirect the default gateway to the VPN, so all non-local traffic from a client is sent through the VPN. This is great for untrusted local networks such as wireless hotspots, as it provides protection against numerous attacks that are a risk on untrusted networks. This is configurable in the GUI using the **Redirect Gateway** checkbox in the OpenVPN instance configuration.

To do this manually for IPv4, add the following custom option:

```
push "redirect-gateway def1";
```

The same value may be used as a custom option on the client side by entering `redirect-gateway def1` without specifying `push`. (Note the option is the letters `def` followed by the digit *one*, **not** the letter L.)

To do the same for IPv6, use:

```
push "redirect-gateway ipv6";
```

20.7.3 OpenVPN Firewall Rules

Permitting traffic to the OpenVPN server

A firewall rule must permit traffic to the OpenVPN server or clients will not be able to connect. Add a rule as follows:

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  UDP

  Source
  any

  Destination
  WAN Address

  Destination port
  `1194`, or whichever port the server is using to listen

  Description
  `Allow traffic to OpenVPN Server`

- Click **Save**
- Click **Apply changes**

This rule is depicted in Figure *OpenVPN Server WAN Rule*.

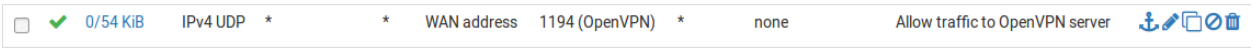


Fig. 1: OpenVPN Server WAN Rule

If the client source addresses are known and do not change, then the source of the rule could be altered to limit traffic from only those clients. This is more secure than leaving the server exposed to the entire Internet, but that is often

necessary to accommodate clients with dynamic IP addresses, roaming clients, and so on. The risk of leaving the service exposed with most OpenVPN configurations is minimal, especially in cases where OpenVPN servers utilize TLS Authentication. With certificate-based authentication there is less risk of compromise than password-based solutions which can be susceptible to brute forcing. This presumes a lack of security holes in OpenVPN itself, which to date has a solid security track record.

Allowing traffic over OpenVPN Tunnels

By default, all traffic is blocked from entering OpenVPN tunnels. To allow traffic from remote OpenVPN hosts to make connections to resources on the local side through the VPN, add firewall rules under **Firewall > Rules**, on the **OpenVPN** tab.

As with other aspects of the firewall these rules only match traffic coming **into** the firewall from remote sources, they do not control traffic leaving from this firewall. Craft the rules accordingly.

Add an OpenVPN rule which passes all traffic as follows:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  any

  Source
  any

  Destination
  any

  Description
  `Allow all on OpenVPN`

- Click **Save**
- Click **Apply changes**

To limit the traffic to only specific sources and destinations, adjust the rule(s) as needed. A strict ruleset is more secure, but more difficult to create.

Tip: Rules on the OpenVPN tab apply to **all** OpenVPN server and client instances. The OpenVPN interface may also be assigned (*Assigning OpenVPN Interfaces*) in which case there will be a separate firewall rule tab for that VPN, upon which rules can pass traffic for that specific VPN.

Rules on assigned OpenVPN interface tabs are processed *after* rules on the **OpenVPN** tab. Rules on the **OpenVPN** tab must not pass or block traffic too broadly or the firewall will stop processing before it gets to the assigned interface tab rules.

20.7.4 OpenVPN clients and Internet Access

For OpenVPN Remote Access clients to reach the Internet through the OpenVPN connection using IPv4, Outbound NAT must translate their traffic to a WAN IP address on the firewall.

The default automatic outbound NAT rules cover this scenario. If the firewall is using manual outbound NAT then manual rules must exist to perform outbound NAT on traffic from sources which include the OpenVPN tunnel network and remote network(s).

See also:

Outbound NAT

20.7.5 Assigning OpenVPN Interfaces

Assigning an OpenVPN interface as an OPT interface allows the firewall to perform complex NAT, policy routing, and tunnel-specific filtering.

Benefits of assigning an OpenVPN instance as an interface include:

- Adds a firewall tab under **Firewall > Rules**
- Adds `reply-to` to rules on the VPN interface tab to help with return routing
- Adds a gateway entry for the far side of the VPN for policy routing
- Allows the interface to be selected elsewhere in the GUI and packages for binding services and other tasks
- Allows fine-grained control of port forwards and outbound NAT for the VPN

Note: This does not change the functionality of OpenVPN, only how the firewall treats the interface.

Interface assignment and configuration

- Navigate to **System > Routing**
- Set the **Default gateway** options to a specific gateway or group, as long as they are not left at *Automatic* (*Managing the Default Gateway*)

  > **Warning:** If the default gateway remains set to *Automatic* the firewall may end up using the OpenVPN interface as the default gateway, which is unlikely to be the desired outcome.

- Navigate to **Interfaces > Assignments**
- Set the **Available network ports** field to the appropriate `ovpns` or `ovpnc` interface

  The GUI prints description of the VPN next to the interface name for reference.

- Click  **Add** to create the interface assignment

  The firewall assigns the interface an automatic `OPTx` interface name (e.g. `OPT1`)

  Figure *Assign OpenVPN Interface* shows `ovpns1` assigned as OPT1.

- Navigate to the interface configuration page, **Interfaces > OPTx**
- Check **Enable**

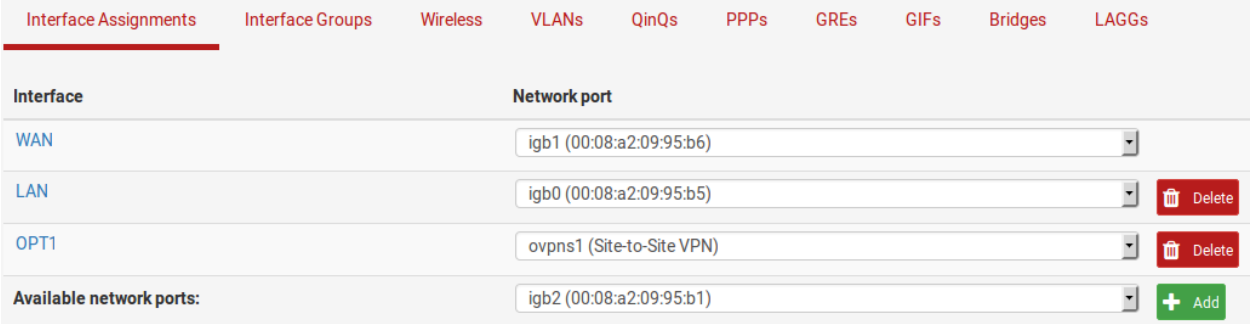


Fig. 2: Assign OpenVPN Interface

- Enter an appropriate **Description** which will become the interface name (e.g. `VPNServer`)
- Click **Save**
- Click **Apply Changes**
- Navigate to **VPN > OpenVPN** and edit the newly assigned instance using the appropriate tab (**Servers** or **Clients**)
- Do not make any changes
- Click **Save** to refresh the VPN configuration and restart its process

> **Warning:** This reinitialization is necessary for the VPN to recover from the assignment process.

Filtering with OpenVPN

When an OpenVPN interface is assigned the GUI contains a tab for the interface under **Firewall > Rules** dedicated to the specific VPN instance.

Rules on this tab govern traffic coming in from the remote side of the VPN and these rules also get the `reply-to` keyword which ensures traffic entering this VPN interface will exit back out the same interface. This can help with advanced NAT and routing configuration scenarios.

Note: Rules on assigned interface tabs are processed *after* rules on the OpenVPN tab. To match the rules on an assigned VPN tab, the traffic **must not match** any rules on the OpenVPN tab. Remove any "Allow All" or "Block all" style rules from the OpenVPN tab and craft more specific rules instead.

See also:

Firewall

Policy Routing with OpenVPN

The firewall automatically creates dynamic gateways for assigned and enabled OpenVPN interfaces. These gateways can be found under **System > Routing**, on the **Gateways** tab.

The firewall will create both IPv4 and IPv6 gateways by default but the **Gateway creation** option on OpenVPN instances can limit this behavior to either IPv4 or IPv6.

Firewall rules can use these gateways to direct traffic into the VPN using the **Gateway** field on LAN or other internal interface rules. These gateways can also be included in gateway groups for failover or load balancing.

Reaching Internet sites through the VPN may require more configuration. Either outbound NAT must be performed on the VPN interface before it leaves the firewall (e.g. For VPN services such as PIA, StrongVPN and similar) or the NAT must be performed on the remote side before it reaches the Internet. To perform outbound NAT on this firewall, switch to **Hybrid** or **Manual** outbound NAT mode and add outbound NAT rules on the assigned OpenVPN interface matching the appropriate traffic sources.

See also:

- *Policy routing*
- *Outbound NAT*

Warning: Do not use this automatic gateway for static routes, use the **Remote Network** field in the VPN configuration instead. Automatic gateways for assigned OpenVPN interfaces **do not** work properly for static routes.

NAT with OpenVPN

Assigned OpenVPN interfaces can utilize any type of NAT rules. This is useful when connecting two conflicting subnets or for making NAT rules specific to this a single VPN connection (outbound NAT, port forwards, or 1:1 NAT)

20.7.6 OpenVPN and Multi-WAN

OpenVPN is multi-WAN capable, with some caveats in certain circumstances. This section covers multi-WAN considerations with OpenVPN server and client configurations.

See also:

- *Multiple WAN Connections*
- *Routing*
- *OpenVPN Site-to-Site with Multi-WAN and OSPF*

OpenVPN assigned to a Gateway Group

A Gateway Group (*Gateway Groups*) may be selected as the **Interface** for an OpenVPN instance. Gateway groups for this purpose must be failover only (one gateway per tier), not load balancing.

Gateway groups can also be set to use a VIP for a specific gateway. When OpenVPN uses a gateway group set this way on a server instance, it will use the interface or VIP of the Tier 1 gateway in the group first. If that gateway goes down, it moves to tier 2, and so on. If the tier 1 gateway comes back up, the VPN will resume operating on that WAN immediately.

When used for a VPN server the server can only be active on one WAN at a time. Some of the other methods described below may be a better fit for most common circumstances, such as needing both WANs usable concurrently with the VPN. When used with OpenVPN clients, the outbound interface will be switched according to the gateway group tiers.

OpenVPN servers and multi-WAN

OpenVPN servers can use any WAN connection, though the methods vary depending on the specifics of a given configuration.

OpenVPN server using TCP

TCP is not the best practice protocol for OpenVPN. However, TCP can make multi-WAN OpenVPN easier to configure when the VPN is using an interface setting of *any*. OpenVPN servers using TCP will work properly on all WANs where the firewall rules allow the traffic to the OpenVPN server. A firewall rule is required for each WAN interface.

This works because of the connection-oriented nature of TCP. OpenVPN can reply to the other end with the proper source preserved since it is part of an open connection.

Warning: This method should be considered a last resort. Using a *protocol* choice which includes *multihome* works properly for UDP on multiple WANs and is a better alternative. Only use TCP if the other methods are not viable.

OpenVPN server using UDP

OpenVPN servers with UDP are also multi-WAN capable, but with some caveats that aren't applicable with TCP.

The *protocol* choice for *UDP on IPv4 and IPv6 on all interfaces (multihome)* will work properly on all WANs and respond back using the address clients expect.

These other UDP modes in OpenVPN are limited by the connectionless nature of UDP. In these cases, the OpenVPN instance replies to the client, but the Operating System selects the route and source address based on what the routing table believes is the best path to reach the peer. For non-default WANs, that will not be the correct path or the address the peer used when contacting this VPN.

Multiple Server Method

In certain cases each WAN may require its own OpenVPN server. The server instances may all use the same certificates. Only two parts of the OpenVPN configuration must change:

Tunnel Network
: Each server must have a unique **Tunnel Network** that does not overlap with any other tunnel network or internal subnet.

Interface
: Each OpenVPN server must specify a different WAN **Interface**.

Port forward method

An easier and more flexible option is to bind the OpenVPN server to *Localhost* or the *LAN* interface and use port forwards on each WAN to direct the OpenVPN port to the service. This method takes advantage of the `reply-to` functionality in PF which returns traffic flows back to the proper source via the interface from which the packet originated.

Note: This method requires minor manual intervention when used with the client export package. The **Host Name Resolution** option must be set to one of the automatic port forwarding methods otherwise the default export settings would have clients connecting to the wrong address. See *OpenVPN Client Export Package* for details.

Automatic Failover for Clients

OpenVPN clients can use multiple remote servers. If a client cannot reach the first server, it will attempt a connection to the second server, and so on until it runs out of servers. Then it starts over again.

See also:

Additional Servers describes this concept and its configuration in more detail.

This behavior can be used in combination with a multi-WAN OpenVPN server deployment to provide automatic failover for clients. If the OpenVPN servers use IP addresses of `198.51.100.3` and `203.0.113.5` with port `1194`, the `remote` lines in the client configuration file would be:

```
remote 198.51.100.3 1194 udp
remote 203.0.113.5 1194 udp
```

For clients configured on pfSense® software, the first `remote` is set by the **Server Host or Address** field in the GUI. Additional `remote` statements must be in the *Custom options* field.

This method has three notable behaviors that some may find undesirable:

- Clients will take at least 60 seconds to detect a failure and switch to the next server.

  This can be fine-tuned by adjusting the keep alive parameters that the server pushes to clients, but making it too sensitive will also be problematic.

- Any connection failure will cause clients to try the next server, even if it is not a WAN failure.
- Clients will **not** automatically reconnect to the first server when it recovers.

  Once a client connects to another server it will stay there until it gets disconnected again.

OpenVPN Clients and Multi-WAN

To use an OPT WAN interface, select it as the **Interface**. OpenVPN clients configured on the firewall will bind to the chosen **Interface**, but may require a manual static route for the server address to ensure traffic takes the correct path.

If the interface is set to *any*, the client will automatically follow the system routing table when selecting the interface and IP address it uses when connecting to the server.

Tip: For some cases, using the same gateway group for both the OpenVPN client and the system default route will result in the best failover behavior.

20.7.7 OpenVPN and High Availability

OpenVPN works well with high availability (HA) on pfSense® software. To provide an HA OpenVPN solution, configure the OpenVPN server or client to use a CARP VIP as its **Interface**. For HA server instances, configure clients to connect to the CARP VIP.

See also:

- *High Availability*
- *Troubleshooting VPN Connectivity to a High Availability Secondary Node*

When XMLRPC Configuration Synchronization settings are enabled, OpenVPN instances will automatically synchronize from the primary node to the secondary. The connection state is not retained between hosts, so clients must reconnect when failover occurs, but OpenVPN will detect the connection failure and reconnect automatically.

When an OpenVPN client instance has a CARP VIP for its **Interface** the firewall will automatically shut down the client as needed while a CARP node is in a BACKUP state. This prevents OpenVPN from making unnecessary outbound connections and also prevents it from placing potentially conflicting information into the routing table. When the CARP VIP status transitions to MASTER, the firewall starts OpenVPN clients automatically.

20.7.8 Sharing a Port with OpenVPN and a Web Server

To be extra sneaky (or careful) with an OpenVPN server, take advantage of the `port-share` capability in OpenVPN which allows it to pass any non-OpenVPN traffic to another IP address behind the firewall.

The usual use case for this is to run the OpenVPN server on TCP port `443` while letting OpenVPN hand off HTTPS traffic from browsers to a web server in place of a port forward.

Warning: This requires using TCP for OpenVPN, and thus is likely to result in reduced VPN performance.

Locked-down networks frequently only allow traffic outbound to common ports such as `80` and `443` for security reasons. Running OpenVPN instances on these allowed ports can help users reach the VPN from restricted networks.

Note: Port sharing is only necessary if these two services must share the same port and IP address. If the firewall has multiple public IP addresses to use, or if there is no public HTTPS web server, then running OpenVPN on TCP port `443` directly without port sharing is a better practice to get the same net effect.

To configure port sharing:

- Move the firewall GUI from port `443` to an alternate port such as `4433`
- Configure an OpenVPN server to listen on TCP port `443`
- Add a firewall rule to pass traffic to the WAN IP address or VIP used for OpenVPN on port `443`
- Add the following to the *Custom options* of the OpenVPN instance:

```
port-share x.x.x.x 443;
```

  Replace `x.x.x.x` with the internal IP address of the web server to which OpenVPN will forward non-VPN traffic.

If an OpenVPN client is pointed to the public address it will connect to the VPN, while a web browser connecting to the same IP address will be connected to the web server.

20.7.9 Controlling Client Parameters via RADIUS

When using RADIUS as an authentication source for a VPN, pfSense® software can receive certain client configuration parameters from the RADIUS server as reply attributes.

Inbound firewall rules

Inbound firewall rules to govern traffic from the client to the server.

```
Cisco-AVPair = <IP_PROTO>:inacl#<NUM>=<rule>
```

- <IP_PROTO> is the address family / IP protocol (`ip` or `ipv6`)
- <NUM> is a rule number
- <rule> is a rule string in Cisco-style ACL format.

  Note: Subnet masks must be **wildcard** style, not CIDR or traditional netmasks.

  The firewall replaces the template strings `{clientip}` and `{clientipv6}` in rules with the Tunnel IP addresses of the connecting client.

FreeRADIUS example:

```
Cisco-AVPair = "ip:inacl#1=permit tcp host 192.168.5.10 host 192.168.6.3 eq 80",
Cisco-AVPair += "ip:inacl#2=permit udp host {clientip} host 192.168.33.4 eq 53",
Cisco-AVPair += "ip:inacl#3=permit ip 192.168.5.0 0.0.0.255 host 192.168.6.4",
Cisco-AVPair += "ipv6:inacl#1=permit icmp host {clientipv6} host 2001:DB8::10",
Cisco-AVPair += "ipv6:inacl#2=permit udp host 2001:DB8::4444 host 2001:DB8::7 range 1024␣
↪65535"
```

Outbound Firewall Rules

Outbound firewall rules to govern traffic from the server to the client.

```
Cisco-AVPair = <IP_PROTO>:outacl#<NUM>=<rule>
```

Aside from the `outacl` keyword, the format is the same as inbound rules.

DNS Servers

DNS servers that OpenVPN will push to this client.

```
Cisco-AVPair = dns-servers=x.x.x.x y.y.y.y
```

Separate multiple servers with spaces.

Routes

Additional route statements OpenVPN will push to the client.

```
Cisco-AVPair = route=x.x.x.x y.y.y.y
```

Specified as `x.x.x.x y.y.y.y` where the first parameter is a network address and the second is a subnet mask.

Static IP Address

A specific IP address OpenVPN will assign to the client.

```
Framed-IP-Address=x.x.x.x
Framed-IP-Netmask=255.255.255.0
```

If the OpenVPN server uses a *subnet* style **Topology** the RADIUS server must also send back an appropriate `Framed-IP-Netmask` value matching the VPN **Tunnel Network**.

When using a *net30* style **Topology**, the client receives this IP address and the server side is set as one IP address **lower** than the address given to the client.

Note: This currently only works for IPv4. The software does not implement the `Framed-IPv6-Address` reply attribute at this time.

20.7.10 OpenVPN Adapter Address ICMP Behavior

When using the older *net30* topology, OpenVPN will not respond to ping on certain virtual addresses used solely for routing endpoints. Thus, in *net30* topology mode pinging the OpenVPN endpoint addresses is unreliable as a means of determining if the tunnel is passing traffic properly. Instead, ping something inside the remote subnet, such as the LAN IP address of the peer.

In contrast, subnet topology in OpenVPN does not use virtual addresses and per-client subnets so it does not have these limitations. OpenVPN servers and clients using *subnet* topology are capable of responding to ping requests on their interface addresses.

The OpenVPN FAQ addresses this behavior in the section titled *Why does OpenVPN's "ifconfig-pool" option use a /30 subnet (4 private IP addresses per client) when used in TUN mode?*:

> As 192.168.1.5 is only a virtual IP address inside the OpenVPN server, used as an endpoint for routes, OpenVPN doesn't bother to answer pings on this address, while the 192.168.1.1 is a real IP address in the servers O/S, so it will reply to pings.

This may seem a little counter-intuitive as an OpenVPN server interface may look like this in `ifconfig`:

```
tun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> metric 0 mtu 1500
        inet6 fe80::202:b3ff:fe03:8028%tun0 prefixlen 64 scopeid 0xc
        inet 192.168.100.1 --> 192.168.100.2 netmask 0xffffffff
        Opened by PID 27841
```

While the client shows:

```
tun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> metric 0 mtu 1500
        inet6 fe80::202:b3ff:fe24:978c%tun0 prefixlen 64 scopeid 0xa
```

(continues on next page)

(continued from previous page)

```
        inet 192.168.100.6 --> 192.168.100.5 netmask 0xffffffff
        Opened by PID 1949
```

In this case, `.5` or `.1` will **not** likely respond to ping. `.5` because it is a virtual address and `.1` because the client has no route to reach it directly. The `.5` and `.6` addresses are part of a `/30` subnet that encompasses `.4` to `.7` and trying to ping `.1` would go out the default route instead.

There are many cases where the far side of an OpenVPN tunnel **can** respond to ping, but not the local. This is also counter-intuitive, but works especially in cases where there is a site-to-site link. If the server shows its tun addresses as `x.x.x.1 -> x.x.x.2` and the client shows the reverse – `x.x.x.2 -> x.x.x.1`, then the far side will likely respond to ping from both ends.

See also:

- *OpenVPN Client Export Package*
- *OpenVPN Server and Client Status*
- *OpenVPN Logs*
- *Connecting OpenVPN Sites with Conflicting IP Subnets*
- *OpenVPN Remote Access Configuration Example*
- *Authenticating OpenVPN Users with FreeRADIUS*
- *Authenticating OpenVPN Users with RADIUS via Active Directory*
- *Installing OpenVPN Remote Access Clients*
- *Installing the OpenVPN Client on iOS*
- *Adding OpenVPN Remote Access Users*
- *OpenVPN Site-to-Site Configuration Example with SSL/TLS*
- *Routing Internet Traffic Through A Site-To-Site OpenVPN Tunnel*
- *Bridging OpenVPN Connections to Local Networks*
- *OpenVPN Site-to-Site with Multi-WAN and OSPF*
- *Troubleshooting OpenVPN*
- *Troubleshooting OpenVPN Internal Routing (iroute)*
- *Troubleshooting Windows OpenVPN Client Connectivity*

OpenVPN is an open source VPN solution which can provide access to remote access clients and enable site-to-site connectivity. OpenVPN clients exist for a wide range of operating systems including all the BSDs, Linux, Android, macOS, iOS, Solaris, Windows, and even some VoIP handsets.

Every OpenVPN connection consists of a server and a client, for both remote access and site-to-site deployments. In the case of site-to-site VPNs, one firewall acts as the server and the other as the client. In most cases it does not matter which firewall acts in a particular role. Typically, the location of the primary firewall will provide server connectivity for all remote locations, whose firewalls are configured as clients. This is functionally equivalent to the opposite configuration the primary location configured as a client connecting to servers running on the firewalls at the remote locations. In practice, the servers are nearly always run on a central location.

OpenVPN implements several types of authentication methods:

X.509 (also known as TLS, SSL, or PKI)
: Utilizes a certificate structure (CA, certificates, and keys). This offers strong security as it cannot be guessed or brute forced.

User authentication
Clients authenticate using credentials such as a username and password which are checked against a local user database, LDAP, or RADIUS server.

Some authentication sources also implement multifactor authentication via mechanisms such as mOTP.

X.509 and User authentication together
The most secure combination, combining multiple factors of authentication that the user has (e.g. certificates, keys) with something they know (credentials).

Shared key
Client and server share a single shared key known to both parties.

Danger: Shared key mode has been deprecated by OpenVPN as it is no longer considered sufficiently secure for modern requirements.

Shared key mode will be removed from future versions of OpenVPN. Users **should not** create any new shared key tunnels and should **immediately** convert any existing shared key tunnels to SSL/TLS mode.

When an SSL/TLS instance is configured with a `/30` tunnel network it behaves similarly to shared key mode. The primary difference is the need to create and distribute the certificate structure to peers. See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for information on configuring OpenVPN in SSL/TLS mode.

Note: While OpenVPN utilizes TLS it is not a "clientless" SSL VPN in the sense that commercial firewall vendors commonly state. The OpenVPN client must be installed on all client devices and it is not browser-based. In reality no VPN solution is truly "clientless", and this terminology is nothing more than a marketing ploy. For more in depth discussion on SSL VPNs, this post from Matthew Grooms, an IPsec tools and former pfSense® software developer, in the mailing list archives provides some excellent information.

For general discussion of the various types of VPNs available in pfSense software and their pros and cons, see *Virtual Private Networks*.

See also:

Hangouts Archive to view the September 2014 Hangout on Advanced OpenVPN Concepts and the September 2015 Hangout on Remote Access VPNs

20.7.11 OpenVPN and Certificates

The best practice is to use certificates for remote access and site-to-site VPNs because it allows access to be revoked for individual clients or sites. Ideally, certificates should be unique per device or at least per user.

If a client machine is compromised, stolen, or lost, or otherwise must be revoked, a shared certificate must be re-issued to all clients. If a client with an individual certificate is compromised, or access needs to be revoked for any other reason, simply revoke that certificate. No other clients are affected.

The pfSense software GUI includes a certificate management interface that is fully integrated with OpenVPN. Certificate authorities (CAs) and server certificates are managed in the Certificate Manager in the web interface, located at **System > Certificates**. User certificates are also managed in the web interface, as a part of the built-in user manager found at **System > User Manager**. Certificates may be generated for any user account created locally on the firewall except for the default admin account.

See also:

For further information on creating a certificate authority, certificates, and certificate revocation lists, see *Certificate Management*.

20.8 IPsec

IPsec provides a standards-based VPN implementation that is compatible with a wide range of clients for mobile connectivity and other devices for site-to-site connectivity. It can interoperate with numerous third party devices, and is being used in production with devices ranging from consumer grade Linksys routers all the way up to IBM z/OS mainframes, and everything imaginable in between.

See also:

For general discussion of the various types of VPNs available in pfSense® software and their pros and cons see *Virtual Private Networks*.

pfSense software offers IPsec with IKEv1 and IKEv2, policy-based and route-based tunnels, multiple phase 2 definitions for each tunnel, NAT traversal, NAT on Phase 2 definitions, a wide variety of encryption and hash options, and many more options for mobile clients including EAP and xauth.

20.8.1 IPsec Terminology

Before delving too deeply into configuration there are a few terms used throughout the chapter which require explanation. Other terms are explained in more detail upon their use in configuration options.

IKE

IKE stands for *Internet Key Exchange* and comes in two different varieties: IKEv1 and IKEv2. Nearly all IPsec-enabled devices can use IKEv1. Most modern implementations also include IKEv2. The IKEv2 protocol is an updated version of IKE that solves some of the difficulties present in the earlier version. For example, IKEv2 has MOBIKE which is a standard for mobile clients that allows them to switch addresses dynamically. It also has built-in NAT traversal and standard mechanisms for reliability similar to DPD. In general IKEv2 provides a more stable and reliable experience provided both ends implement it sufficiently.

ISAKMP Security Association

ISAKMP stands for Internet Security Association and Key Management Protocol. It gives both parties a mechanism by which they can set up a secure communications channel including exchanging keys and providing authentication.

An ISAKMP Security Association (ISAKMP SA or IKE SA) is a one-way policy which defines how traffic will be encrypted and handled. Each active IPsec tunnel will have two security associations, one for each direction. The ISAKMP Security Associations are set up between the public IP addresses for each endpoint. Knowledge of these active security associations is kept in the Security Association Database (SAD).

Security Policy

A security policy manages the complete specifications of the IPsec tunnel. As with security associations these are one-way, thus for each tunnel there will be one in each direction. These entries are kept in the Security Policy Database (SPD). The SPD is populated with two entries for each tunnel connection as soon as a tunnel is added. By contrast SAD entries only exist upon successful negotiation of the connection.

In pfSense software security policies for policy-based IPsec tunnels control which traffic will be intercepted by the kernel for delivery via IPsec.

Phase 1

There are two phases of negotiation for an IPsec tunnel. During phase 1 the two endpoints of a tunnel set up a secure channel between using ISAKMP to negotiate the IKE SA entries and exchange keys. This also includes authentication, checking identifiers, and checking the pre-shared keys (PSK) or certificates. When phase 1 is complete the two ends can exchange information securely, but they have not yet decided which traffic will traverse the tunnel or its encryption.

Phase 2

In phase 2 the two endpoints negotiate how to encrypt and send the data for the private hosts based on security policies. This part builds an entry referred to as a "Child SA". This forms the connection used to transfer data between the endpoints and clients whose traffic is handled by those endpoints. If the policies on both side agree and a phase 2 child SA is successfully established the tunnel will be up and ready for use.

Mobile IPsec

Mobile IPsec refers to IPsec connections from individual client devices rather than site-to-site connections. This is also commonly called a "Road Warrior" or "Remote Access" style VPN.

The main purpose of a mobile IPsec VPN is for users who are not in the office who need to connect back to the main network. Common use cases are for employees working from home, sales personnel using Wi-Fi on a business trip, or even the boss from his cabin via LTE modem.

Most of these use cases are forced to deal with dynamic IP addresses, unknown IP addresses, NAT (regular and Carrier Grade NAT), and other complications. Without a router or firewall that can use IPsec, a traditional IPsec tunnel will not work.

In telecommuting scenarios, it's usually undesirable and unnecessary to connect an entire home networks to the office network, and doing so can reduce security and introduce routing complications. This is where IPsec Mobile Clients are most useful.

Instead of relying on a fixed address for the remote end of the tunnel, Mobile IPsec uses authentication to distinguish between authorized users. For example, this could be a username and password with IKEv2 and EAP, a per-user Identifier and Pre-Shared Key pair, or a certificate.

20.8.2 IPsec Configuration

IPsec on pfSense® software offers numerous configuration options which influence the performance and security of IPsec connections. For most users performance is the most important factor. When crafting a configuration, carefully select options to ensure optimal efficiency while maintaining strong security and compatibility with equipment on both ends of a tunnel.

Tip: For low-to-moderate bandwidth usage deployments the options may not have significant impact on performance. Even so, take care not to use insecure options such as SHA1 or weak pre-shared keys.

Warning: pfSense software offers several options which are weak from a security standpoint. These are included for compatibility with third party vendors and equipment which do not implement stronger options. The GUI includes warnings against using these options.

The next sections outline how to design an IPsec tunnel and the options available.

IPsec Tunnel Design

Before configuring an IPsec tunnel, a few general decisions must be made about how the tunnel will operate.

IPsec Modes

pfSense software offers several primary modes of IPsec operation:

Policy-based IPsec
: This mode uses policies to match specific combinations of traffic which are grabbed by the kernel and pushed through an IPsec tunnel. It also uses special "trap" policies to detect when traffic intends to use IPsec so that it can bring the tunnel up automatically. Only traffic specifically matching phase 2 child SA entries can use IPsec, and all traffic matching those entries will be taken over by IPsec.

  This mode is the most common and is compatible with most third party IPsec implementations.

Route-based IPsec (VTI)
: Routed IPsec uses a special Virtual Tunnel Interface (VTI) for each IPsec tunnel. The VTI interface is assigned and used like other interfaces. Phase 2 entries define addresses for the tunnel interface itself rather than policies which direct traffic to IPsec. Arbitrary traffic may cross VTI IPsec tunnels as traffic follows the system routing table. Static routes or dynamic routing daemons can control which traffic crosses a tunnel.

  Routed IPsec compatibility varies by vendor.

  By default, traffic for VTI tunnels is filtered on the IPsec tab and cannot use per-interface rules, NAT, or `reply-to`. This can be changed in *Advanced IPsec Settings* using the **IPsec Filter Mode** option. Read the consequences of that option carefully before changing the behavior.

Mobile IPsec
: Similar to policy-based mode, but for remote access/mobile clients.

Transport Mode
: This mode encrypts all traffic from the external IP address on this firewall to the external IP address on the far side as defined in the phase 1 settings. Since all traffic sent between the two nodes will be encrypted, other tunneling methods that do not employ encryption, such as a GIF or GRE tunnel, can be safely used by the firewall between the endpoints.

By default, traffic for transport mode tunnels may experience problems with keeping state appropriately among other filtering quirks. This can be improved by using the **IPsec Filter Mode** option in *Advanced IPsec Settings*. Read the consequences of that option carefully before changing the behavior.

Interface Selection

In many cases, the **Interface** option for an IPsec tunnel will be WAN, since the tunnels are connecting to remote sites. However, there are plenty of exceptions, the most common of which are outlined in the remainder of this section.

High Availability Environments

IPsec works well with high availability, with some caveats. See *IPsec in High Availability Environments* for details.

IP Alias VIP

If multiple IP addresses are available on an interface using IP Alias type VIPs, they will also be available in this list. To use one of those IP addresses for the VPN instead, select it here.

Multi-WAN Environments

IPsec can utilize multiple WANs in multiple configurations. See *IPsec in Multi-WAN Environments* for details.

Wireless Internal Protection

When configuring IPsec to add encryption to a wireless network as described in *Additional protection for a wireless network*, choose the OPT interface which corresponds to the wireless card. When using an external wireless access point, pick the interface which is connected to the wireless access point.

IPsec Tunnels Tab

IPsec VPN tunnels are managed by the **Tunnels** tab at **VPN > IPsec**. The page contains a list of tunnels with a brief summary of their settings along with various management controls.

See also:

See *Managing Lists in the GUI* for details on how to create and edit items in lists using the GUI.

Each IPsec tunnel contains a single phase 1 definition and one or more phase 2 definitions.

The  **Add P1** button creates a new IPsec phase 1 entry.

See also:

- *Phase 1* (Definition)
- *Phase 1 Settings*

After adding a phase 1 definition, click  **Show Phase 2 Entries** underneath a phase 1 entry to display and manage its phase 2 entries.

See also:

- *Phase 2* (Definition)
- *Phase 2 Settings*

Continue on for detailed information about IPsec settings.

Phase 1 Settings

The settings here control the phase 1 negotiation portion of the tunnel, as described previously.

General Information

Description
A name or brief description of the tunnel. The best practice is to enter a few words to describe the purpose of this VPN tunnel or about the remote end of the tunnel. For example `Remote Office`, `HQ Tunnel`, or `ATX to NYC`.

This functions as a reminder for anyone managing the firewall as to who or what will be using the tunnel. This description is also reflected in the IPsec status which makes it easier to match up status entries with a specific tunnel.

Disabled
Controls whether this tunnel (and its associated phase 2 entries) are active and used.

IKE ID
This is a read-only field containing the IKE identifier for this tunnel.

IKE Endpoint Configuration

Key Exchange Version
This can be *IKEv1*, *IKEv2*, or *Auto*.

See also:

The differences between IKEv1 and IKEv2 are discussed in *IKE*.

IKEv1
IKEv1 is more common and widely compatible but has problems handling common modern issues such as NAT and mobile clients.

IKEv2 (Default)
An updated version of the protocol which has increased capabilities and security, as well as built-in compatibility with mobile clients and NAT.

Tip: IKEv2 is the best choice when available on both endpoints.

Auto
This option uses IKEv2 when initiating, but will accept either IKEv2 or IKEv1 when responding.

Internet Protocol
The protocol for the *outside* of the tunnel. That is, the protocol that will be used between the outside peer addresses. In most cases this will be IPv4, but if both ends are capable of IPv6 that may be used instead. Whichever protocol is chosen here will be used to validate the **Remote Gateway** and the associated identifiers.

Note: On current versions of pfSense software IPv4 and/or IPv6 traffic may be carried inside a tunnel no matter which type of Key Exchange Version or Internet Protocol is used outside the tunnel.

Some third party vendors may only be compatible with this when using IKEv2, and on IKEv1 the inner and outer traffic might need to match when connecting to those third party implementations.

Interface
This determines which part of the network will be the termination point (end point) for the IPsec tunnel. If the tunnel will be connecting to a remote server, then WAN is likely the desired setting. This can also be a virtual IP address. A gateway group can also be used for automatic failover.

See also:

See *Interface Selection* earlier in this document for details on selecting the appropriate interface.

Remote Gateway
The address for the peer to which the tunnel will be established. This is most likely the WAN IP address of the remote device.

This may be set to an IP address or a fully qualified domain name (FQDN). When set to an FQDN the firewall periodically resolves the name using DNS and updates the tunnel when it detects a change.

To allow connections from any remote endpoint with tunnel mode phase 2 entries, use `0.0.0.0` for IPv4 or `::` for IPv6. When allowing connections from any remote endpoint the **Child SA Start Action** must be set to *None* and the **Peer Identifier** cannot be set to *Peer IP Address*.

Note: VTI tunnels require a specific remote endpoint address for their interface configuration. They cannot allow connections from "any" or unknown endpoints. Using an FQDN for the remote gateway is a viable workaround provided the remote peer is capable of using dynamic DNS.

Phase 1 Proposal (Authentication)

Authentication Method
An IPsec phase 1 can be authenticated using a pre-shared key (PSK) or certificates. The **Authentication Method** selector chooses which of these methods will be used for authenticating the remote peer. Fields appropriate to the chosen method will be displayed on the phase 1 configuration screen.

Mutual PSK
The peer is validated using a pre-defined string known to both endpoints. Since it is simple a string there is a possibility that it can be guessed. For this reason a long and complex key is more secure when using PSK mode.

Mutual Certificate
A TLS CA and certificate are used to verify the peer. During the phase 1 exchange each peer sends its certificate to the other peer and then validates it against a CA. The CA and certificate must be created for the tunnel before attempting to set up the phase 1.

See also:

IPsec Site-to-Site VPN Example with Certificate Authentication

Mutual Certificate (PKCS#11)
Similar to *Mutual Certificate* but the certificate is read from a locally attached PKCS#11 device.

Note: This option is only available when **Enable PKCS#11** is checked on the **Advanced** IPsec settings.

Mutual PSK+Xauth
Used with mobile IPsec and IKEv1. This selection enables xauth username and password verification along with a shared (or "group") pre-shared key.

See also:

IPsec Remote Access VPN Example Using IKEv1 with Xauth

Mutual Certificate+Xauth
Used with mobile IPsec and IKEv1. This selection enables xauth username and password verification along with certificate authentication using certificates on both the client and server.

Hybrid Certificate+Xauth
Used with mobile IPsec and IKEv1. This selection enables xauth username and password verification along with a certificate only on the server side. It is not quite as secure as *Mutual Certificate+Xauth*, but it is easier on the clients.

EAP-TLS
Used with mobile IPsec and IKEv2. EAP-TLS verifies that certificates on the client and server are from the same shared CA, similar to *Mutual Certificate*. The client and server certificates require special handling:

- The server certificate must have the firewall hostname as it exists in DNS listed in its Common Name and again as a Subject Alternative Name (SAN). The firewall IP address must also be listed in a SAN.
- The identifier in phase 1 must also be set to match the firewall hostname as listed in the Common Name of the certificate.
- The client certificate must have the username listed as the common name and then again as a SAN.

The CA and server certificates must be generated before attempting to configure EAP-TLS. The CA and user certificate must be imported into the client.

See also:

IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS

EAP-RADIUS
Used with mobile IPsec and IKEv2. This selection performs CA verification along with username and password authentication via RADIUS. A RADIUS server must be selected on the **Mobile Clients** tab. Though user certificates are not necessary EAP-RADIUS still requires that a CA and server certificate be present using the same attributes mentioned under *EAP-TLS*. The CA must be imported to, or globally trusted by, the client, but no user certificate.

See also:

IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS

EAP-MSCHAPv2
Used with mobile IPsec and IKEv2. EAP-MSCHAPv2 works identically to EAP-RADIUS except the usernames and passwords are defined on the **Pre-Shared Key** tab under **VPN > IPsec** with the **Secret type** set to *EAP*. It also requires a CA and server certificate with the same properties listed previously. The CA must be imported to, or globally trusted by, the client, but no user certificate.

See also:

IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2

Negotiation Mode
(IKEv1 only) The type of authentication security used by this tunnel. This can be either *Main* or *Aggressive*.

Main
Main is the most secure mode as it is strict in its validation. It requires several packets between the peers to accomplish a successful negotiation due to the strict validation procedures, thus is slower.

Using *Main* mode is the best practice where possible due to its higher security.

Aggressive
Aggressive is generally the most compatible mode to use with other vendors and is the fastest mode. It is more forgiving with identifier types and tends to be more successful when negotiating with third-party devices.

Aggressive mode is faster because it sends all the identifying information in a single packet, which also makes it less secure because the verification of that data is not as strict as that found in main mode.

Avoid aggressive mode due to its weaker security unless it is required for interoperability with a third party IPsec implementation.

Identifiers
Identifiers are used by IPsec to identify remote peers and associate specific peers with a tunnel and its related settings, such as authentication components (keys, certificates, etc.).

My Identifier
Identifier type and value sent by this firewall to the far side. It is best left at *My IP Address* and the firewall will fill it in as needed. In some cases an FQDN or similar may be entered so that the value is constant.

Most settings are viable so long as both sides agree on the identifier type and value.

Peer Identifier
Identifies the remote peer on the other side of the tunnel. It is best left at *Peer IP Address* and the firewall will fill it in as needed. In some cases an FQDN or similar may be entered so that the value is constant.

Most settings are viable so long as both sides agree on the identifier type and value.

Identifier Types

My IP Address / Peer IP address
This choice is a macro that will automatically use the IP address on the interface, or the selected VIP, as the identifier.

For peers this is the IP address from which incoming IPsec packets are received by this firewall, which should match the value used in **Remote Gateway**.

IP Address
A specific static and manually-entered IP address to be used as the identifier. One potential use for this is when the firewall is behind NAT where the real external IP address could be used in this field.

Fully qualified domain name
A fully qualified domain name (FQDN) such as `host.example.com`. Enter a value in that format into the box. This value is sent as-is and is not resolved to an IP address.

When using *Mutual Certificate* authentication this can be used to match the CN or SAN of a certificate to ensure the correct certificate is associated with this tunnel endpoint.

User fully qualified domain name / E-mail
An e-mail address such as `vpn@example.com`.

ASN.1 Distinguished Name
When using *Mutual Certificate* authentication this can be set to the subject of the certificate which also ensures the correct certificate is associated with this tunnel endpoint.

For example: `/CN=host.example.com/C=US/ST=Texas/L=Austin/O=Example Co`

KeyID Tag
An arbitrary string to use as the identifier.

Dynamic DNS
A hostname to resolve and use as the identifier. This is mostly useful if the firewall is behind NAT and has no direct knowledge of its external IP address aside from a dynamic DNS hostname. This is only available for **My identifier**.

Any
In cases when the remote identifier is unknown or cannot be matched the **Peer Identifier** may be set to **Any**. This is more common on certain types of mobile configurations, but it is a much less secure choice than matching the identifier properly.

Pre-Shared Key
(*Mutual PSK* authentication only) A string known by both peers used as a key to authenticate the tunnel, similar to a password. This key is case-sensitive and must be exactly the same on both endpoints.

The best practice is to make this as long and complex as possible. There is little incentive to keep the value simple as this only gets entered once on each peer and there is no need for a human to remember its contents.

Click  **Generate new Pre-Shared Key** to populate the field with a random string suitable for use as a key.

> **Warning:** This key must be as random as possible to protect the contents of the tunnel.

My Certificate
(*Mutual Certificate* authentication only) Defines the certificate which identifies this firewall. The CA which signed this certificate must be copied to the peer. If one is not shown, create or import it under **System > Cert Manager** on the **Certificates** tab.

Peer Certificate Authority
(*Mutual Certificate* authentication only) Defines the CA which has signed the certificate sent by the peer. This is used to validate the peer certificate. If it does not show in the list, import it under **System > Certificates** on the **Certificate Authorities** tab.

Phase 1 Proposal (Encryption Algorithm)

There are many options for encryption algorithms on both phase 1 and phase 2.

Encryption choices depend on the device to which the tunnel will connect and the hardware. AES-GCM is the most desirable cipher for both speed and security, but may not be widely implemented by other vendors. AES with a 128-bit key length is a common choice on endpoints which cannot use AES-GCM.

Multiple combinations of these options can be defined using the  **Add Algorithm** button to add another line. The order of the entries is the order of preference so configure the strongest and/or most desirable settings first.

Encryption Algorithm
Use the strongest available option available on both endpoints. If both sides can use AES-GCM, use *AES128-GCM* with a *128*-bit **Key Length**. This will combine strong encryption and hashing together and can be accelerated by AES-NI. Failing that, use *AES* With a **Key Length** of *128* or whichever option is strongest in common between both sides.

Hash Algorithm
Hash algorithms are used with IPsec to verify the authenticity of packet data and as a Pseudo-Random Function (PRF). These hash algorithms may also be referred to with HMAC (Hash Message Authentication Code) in the name in some contexts, but that usage varies depending on the hardware or software in use.

A tunnel configured for AES-GCM uses this solely as a PRF because AES-GCM performs hashing internally. The fastest choice to use in combination with AES-GCM on hardware capable of accelerating AES, such as AES-NI, is *AES-XCBC*. However, *AES-XCBC* is less secure than other algorithms such as *SHA256* and may not be available on other platforms. For different types of **Encryption Algorithms**, use *SHA256* if possible. If the peer does not offer any of these, use the strongest option available.

PRF
Specifically set a manual Pseudo-Random Function different from the one the IPsec daemon would choose automatically based on the **Hash Algorithm**. This control is hidden by the GUI unless **PRF Selection** is enabled in the **Advanced Options** section at the bottom of the page.

DH Key Group
The best practice is to use DH Group *14 (2048 bit)* or higher when possible. As with the other options, if the suggested value is not available on the peer, use the strongest available option.

> **Warning:** Avoid using groups 1, 2, 22, 23, and 24 as they do not provide sufficient security.

Note: IPsec uses this DH group for the first child SA when initially building a tunnel. The PFS option of a phase 2 entry is used for subsequent child SA entries and when rekeying.

Expiration and Replacement

The total lifetime for phase 1 defines how often the connection will be rekeyed or reauthenticated by the IPsec daemon.

Warning: Take care when crafting these values. Incorrect or suboptimal values can lead to problems such as tunnels failing to renegotiate in a timely manner or multiple duplicate security associations.

See *Troubleshooting Duplicate IPsec SA Entries* for more details.

The specific values for these fields depend on the IKE mode and which mechanisms are available on both endpoints. In most cases setting only the **Life Time** value will allow the firewall to choose the best options.

Life Time
The hard IKE SA lifetime, in seconds, after which the IKE SA will be expired. This value **must** be larger than **Rekey Time** and **Reauth Time** and cannot be set to the same value.

`28800` total seconds is a good balance of frequent rekeying without being too aggressive.

Tip: Set one endpoint to this recommended value, but use a higher **Life Time** on the other endpoint by at least 10% (e.g. `31680`) to help avoid overlap.

If left empty the value defaults to 110% of **Rekey Time** or **Reauth Time**, whichever is higher.

Rekey Time
Time, in seconds, before the IPsec daemon attempts to establish a new set of keys for the IKE SA. Only compatible with IKEv2, and the best choice for use with IKEv2.

Rekey works without interruption and allows both endpoints to seamlessly change to new keys on the fly. This is optimal, but implementation quality varies by vendor.

Leave blank to automatically calculate the value based on 90% of **Life Time**. Enter a value of `0` to disable rekeying.

Normally both sides will rekey as needed. If the tunnel often fails when a rekey event occurs, try disabling this feature on one side.

Note: Some clients, especially Windows clients behind NAT, misbehave when they receive a rekey request. In those cases it is safer to allow the client to initiate the rekey by disabling the option on the server.

Reauth Time
Time, in seconds, before an IKE SA is torn down and recreated from scratch by the IPsec daemon. This also includes a new round of authentication. Can be used with IKEv1 and IKEv2, but should be avoided with IKEv2 where possible.

This process can be disruptive to traffic flow unless all peers are compatible with IKEv2 make-before-break (*Advanced IPsec Settings*) and overlapping IKE SA entries.

Leave blank to automatically calculate the value based on 90% of **Life Time**. Enter a value of `0` to disable reauthentication.

Rand Time
Introduces randomness into the rekey or reauthentication process to avoid both endpoints attempting to renegotiate simultaneously.

A random value up to this amount will be subtracted from **Rekey Time** or **Reauth Time** for each scheduled renegotiation to reduce the chance of collisions.

If left empty, the value defaults to 10% of **Life Time**. Enter `0` to disable randomness.

Warning: Disabling **Rand Time** increases the likelihood of simultaneous renegotiation which can lead to duplicate security associations.

See *Troubleshooting Duplicate IPsec SA Entries* for more details.

Advanced Options

Child SA Start Action

This option forces specific behavior performed by the IPsec daemon when loading a phase 2 configuration ("Child SA") during initial service startup. This happens at boot and when the service is restarted for any reason.

Default

Automatically chooses behavior based on other settings.

None (Responder Only)

The IPsec daemon will not attempt to initiate the tunnel. The tunnel will only be established by an initiation attempt from the far side. If DPD detects that the tunnel has failed, the tunnel will be left down rather than restarted. The far side must reconnect.

This is the default behavior for mobile IPsec and tunnels with unknown remote endpoints.

Initiate at Start (VTI or Tunnel Mode)

The firewall will attempt to establish the IPsec tunnel immediately when the IPsec daemon starts.

This is the default behavior for VTI mode.

Initiate on Demand (Tunnel mode only)

Traffic which matches the networks in phase 2 definitions for this tunnel ("interesting traffic") will trigger initiation of the tunnel.

This is the default behavior of tunnel mode.

Note: Finding "interesting" traffic relies on trap policies to function and trap policies are not compatible with VTI mode.

Child SA Close Action

Controls how the IPsec daemon behaves when a child SA (P2) is unexpectedly closed by the peer.

Default

Retains the default behavior based on other settings for the tunnel.

Close connection and clear SA

Removes the child SA and does not attempt to establish a new SA. This is the desired behavior when acting in a **Responder Only** or mobile IPsec role.

Restart/Reconnect

Immediately attempts to reconnect the child SA. This ensures that the tunnel reestablishes properly in cases that are not compatible with trap policies such as routed IPsec

(VTI). Set this on **one side only** if the tunnel does not reconnect after it disconnects, rekeys, or reauthenticates.

> **Warning:** This option **must not** be set on both peers! Both peers will attempt to initiate and hold open multiple copies of each child SA.

Close connection and reconnect on demand
Clears the child SA and reinstalls trap policies to watch for interesting traffic. Will reestablish the tunnel on demand when traffic attempts to cross the tunnel.

This option requires the use of trap policies, so it cannot be used with modes incompatible with trap policies, such as routed IPsec (VTI).

NAT Traversal
(IKEv1 Only) Also known as NAT-T. NAT Traversal encapsulates ESP traffic for IPsec inside of UDP packets to more easily function in the presence of NAT. If this firewall or the firewall on the other end of the tunnel is behind a NAT device then NAT Traversal will likely be necessary for the tunnel to function properly.

Auto
(Default) Allows the IPsec daemon to detect and use NAT Traversal automatically when it determines one or both peers is behind NAT.

Force
Instructs the IPsec daemon to always use NAT Traversal for the tunnel. This can help if there is a known issue detecting NAT. It can also help with issues carrying ESP traffic between the two endpoints even when neither side is behind NAT.

IKEv2 integrates NAT Traversal natively, so the option is unnecessary in that case.

MOBIKE
An extension to IKEv2 which handles multi-homed clients and clients which roam between different IP addresses. This is primarily used with mobile clients to allow them to switch remote addresses while keeping a connection active. Leave disabled unless the remote peer must change addresses dynamically.

Gateway Duplicates
Allows multiple phase 1 configurations to use the same remote endpoint address.

> **Warning:** This option also disables automatic static routes to the peer via specific WAN gateways. Traffic will follow the default route, not the selected tunnel interface, unless manual static routes redirect the traffic.

Split Connections
(IKEv2 Only) By default when an IKEv2 tunnel has multiple phase 2 definitions the settings are collapsed in the IPsec configuration such that all phase 2 combinations are held in a single child SA.

Split Connections changes this behavior to be more like IKEv1 where each phase 2 entry is configured by the daemon as its own separate child SA.

Certain scenarios require this behavior, such as:

- The remote peer does not properly handle multiple addresses in single traffic selectors. This is especially common in Cisco, Checkpoint, Fortinet, and Juniper equipment.

- Each child SA must have unique traffic selector or proposal settings. This could be due to the peer only allowing specific combinations of local/remote subnet pairs or different encryption options for each child SA.

PRF Selection
Enables a GUI control to specifically set a Pseudo-Random Function (PRF) rather than allow the IPsec daemon to choose one automatically based on the selected Hash Algorithm. Can be useful in combination with AEAD encryption algorithms such as AES-GCM.

Custom IKE/NAT-T Ports
In rare situations the remote endpoint may be running IPsec on alternate port numbers for IKE and NAT-T. These settings can accommodate such endpoints.

Remote IKE Port
The UDP port for IKE on the remote gateway. Leave empty for the default automatic behavior (Port `500` for IKE and `4500` for NAT-T)

Remote NAT-T Port
The UDP port for NAT-T on the remote gateway.

Note: If **Remote IKE Port** is empty and NAT-T contains a value the tunnel will use only NAT-T.

Dead Peer Detection
Dead Peer Detection (DPD) is a periodic check that the host on the other end of the IPsec tunnel is still alive. This detects when an IPsec peer has lost connectivity or is otherwise unreachable. If a DPD check fails the tunnel is torn down by removing its associated SAD entries and a fresh negotiation is attempted.

The default settings are sufficient for most connections. Increase the values for bad quality links to avoid tearing down a usable, but lossy, tunnel.

Delay
Time between DPD probe attempts. The default of `10` is best.

Max Failures
Number of failures before the peer is considered down. The default of `5` is best. This per-tunnel value is only honored for IKEv1.

Note: The default values of `10` seconds and `5` failures will result in the tunnel being considered down after approximately one minute.

Phase 2 Settings

The phase 2 settings for an IPsec tunnel govern how the tunnel handles traffic (e.g. policy-based or route-based, see *IPsec Modes*) as well as the encryption of that traffic.

Phase 2 entries are used in a few different ways, depending on the IPsec configuration:

- For policy-based IPsec tunnels this controls which subnets will enter IPsec. Multiple phase 2 definitions can be added for each phase 1 to allow using multiple subnets inside a single tunnel.
- For route-based IPsec this controls the VTI interface addresses.
- For mobile IPsec this primarily controls the encryption for phase 2. It can also optionally be used by the IPsec daemon or export utilities to generate a list of networks to the clients for use in split tunneling.

Each phase 2 entry has the following options:

General Information

Description
A description for this phase 2 entry. Shows up in the IPsec status for reference.

A name or brief description for this entry. The best practice is to enter a few words to describe the purpose of this phase 2 entry or about the remote end of the tunnel. For example `TNSR VTI`, `DC Management`, or `ATX DMZ to NYC DMZ`.

This functions as a reminder for anyone managing the firewall. This description is also reflected in the IPsec status which makes it easier to match up status entries with a specific tunnel.

Disabled
An on/off switch for this phase 2 entry only.

Mode
The IPsec Mode for this phase 2 entry, which controls how the tunnel handles traffic. See *IPsec Modes* for more detailed explanations of each type of mode.

Tunnel IPv4
A policy-based tunnel that will carry traffic between IPv4 networks matching the specified **Local Network** and **Remote Network**.

Tunnel IPv6
A policy-based tunnel that will carry traffic between IPv6 networks matching the specified **Local Network** and **Remote Network**.

Transport
Encrypts all traffic between the endpoints. The **Local Network** and **Remote Network** are not set for transport mode, the addresses are based on the phase 1 settings.

Routed (VTI)
Routed IPsec using Virtual Tunnel Interfaces (e.g. `ipsecX`). The **Local Network** and **Remote Network** define the addresses used by the firewall for the VTI interface. Typically, only one phase 2 entry is present for each address family (e.g. one for IPv4, one for IPv6)

See *Routed IPsec (VTI)* for more information.

Phase 1
A link to the IPsec phase 1 entry under which this phase 2 entry exists, along with its IKE ID.

P2 Reqid
The unique phase 2 request ID for this entry. It is used by IPsec in various contexts, such as for the IPsec VTI interface name.

Networks

Local Network

Tunnel Mode
Defines which subnet or host can be accessed from the other side of the VPN tunnel. This is typically the LAN or other internal subnet for the VPN. It can also be a single IP address if only one client needs to use the tunnel. The **Type** selector is preloaded with choices for each interface (e.g. *LAN subnet*), as well as *Address* and *Network* choices that allow entering an IP address or subnet manually.

Most often this is set to *LAN subnet* which means the entire LAN will be accessible from the remote network.

NAT/BINAT
Sets a *different* subnet or address which is used by IPsec to perform NAT on the local network addresses to make them appear to the remote peer as a different subnet.

To disable NAT for the tunnel, set this to *None*.

See also:

For more details see *NAT with IPsec Phase 2 Networks*.

Routed (VTI) Mode
Sets the local IP address and subnet mask of the `ipsecX` interface.

Remote Network

Tunnel Mode
(Non-mobile only) Specifies the IP Address or Network which exists on the remote side of the VPN. This field operates similarly to the **Local Network** option.

Routed (VTI) Mode
Sets the remote IP address for the `ipsecX` interface tunnel network (the remote address of the VTI).

Phase 2 Proposal (SA/Key Exchange)

Protocol
Controls how IPsec protects its traffic.

ESP (Encapsulating Security Payload)
Encrypts traffic before sending it to the peer.

ESP is the correct choice in nearly all circumstances.

AH (Authenticated Header)
Provides assurance the traffic came from a trusted source but does not provide encryption. Rarely used in practice.

Note: With automatic VPN rules (*Disable Auto-added VPN rules*) the firewall automatically passes the appropriate ESP or AH protocol traffic from the remote endpoint. If automatic VPN rules are disabled, add manual rules to pass the traffic instead.

Encryption algorithms
Sets the encryption algorithms used when negotiating phase 2 child SA entries with peers. Must match values available to and configured on the peer.

In systems with AES-NI the fastest and most secure choice is AES-GCM if it is available on the peer. When using AES-GCM do not select any options for **Hash Algorithms** in phase 2.

If both peers cannot use AES-NI then use AES with a 128-bit or higher key length.

This set of controls allows for multiple selections so that multiple choices will be accepted when acting as a responder or proposed when working as an initiator. The best practice is to only select a single desired cipher on both peers. In some cases, such as mobile clients, selecting multiple will allow a tunnel to work with a variety of clients. It can also be better when acting in both a responder and initiator role.

Hash algorithms
Controls which hash algorithms are used when negotiating phase 2 child SA entries with peers. Must match values available to and configured on the peer.

As with the Encryption Algorithms multiple hashes may be selected. The best practice is to select a single desired choice if possible.

See also:

For more discussions on the quality of the various hash types see *Phase 1 Settings*.

With AES-GCM as the **Encryption Algorithm** no **Hash Algorithm** should be selected as AES-GCM performs hashing on its own. When not using AES-GCM, the optimal choice for speed and security is **SHA256**.

PFS key group
Perfect Forward Secrecy (PFS) provides keying material with greater entropy which improves the cryptographic security of the connection. This comes at a cost of higher CPU usage during rekeying.

The options have the same properties as the DH key group option in phase 1 (See *DH key group*) and some products also refer to them as "DH" values even in phase 2.

The optimal choice for speed and security is *14 (2048 bit)*.

Note: When an IPsec tunnel is initially connected the first child SA will use the DH value from phase 1 in this context. The PFS value from phase 2 is used for subsequent child SA entries and when rekeying. This can lead to a misconfiguration not being noticed initially since it may work at first but fail after the first child SA expires.

Expiration and Replacement

Life Time
The hard child SA lifetime, in seconds, after which the child SA will be expired. This value **must** be larger than **Rekey Time** and cannot be set to the same value.

`3600` total seconds is a good balance of frequent rekeying without being too aggressive.

Tip: Set one endpoint to this recommended value but use a higher **Life Time** on the other endpoint by at least 10% (e.g. `5400`) to help avoid overlap.

If left empty the value defaults to 110% of **Rekey Time**. If both **Life Time** and **Rekey Time** are empty it defaults to `3960`.

Rekey Time
Time, in seconds, before the child SA establishes a new set of keys. This works without interruption and allows both endpoints to seamlessly change to new keys on the fly.

Leave blank to automatically calculate the value based on 90% of **Life Time**. If both **Life Time** and **Rekey Time** are empty it defaults to `3600`. Enter a value of `0` to disable rekeying.

Note: If rekeying is disabled connections can be interrupted while a new child SA is negotiated after an old entry expires.

Rand Time
Introduces randomness into the rekey process to avoid both endpoints attempting to renegotiate simultaneously.

A random value up to this amount will be subtracted from **Rekey Time** for each scheduled renegotiation to reduce the chance of collisions.

If left empty the value defaults to 10% of **Life Time**. Enter `0` to disable randomness.

Warning: Disabling **Rand Time** increases the likelihood of simultaneous renegotiation which can lead to duplicate security associations.

See *Troubleshooting Duplicate IPsec SA Entries* for more details.

Keep Alive

Methods which will attempt to keep this phase 2 entry up and active over time.

See also:

See *Configuring IPsec Keep Alive* for additional discussion on these options.

Automatically ping host
For use on non-mobile tunnel mode entries. This option tells the firewall to initiate a ping periodically to the specified IP address. This option only works if the firewall has an IP address inside the **Local Network** for this phase 2 entry and the value of the ping host here is inside the **Remote Network**.

This option will not initiate a tunnel if its phase 1 **Child SA Start Action** is set to *Responder Only*. This option will also not initiate VTI mode tunnels as they are not compatible with trap policies.

Keep Alive
Enables a periodic check to see if the child SA is connected and initiates when it is down. This function does not send traffic inside the tunnel.

This option works for VTI and tunnel mode phase 2 entries.

For IKEv2 without split connections this only needs to be enabled on the first phase 2 entry.

See also:

See *Configuring IPsec Keep Alive* for additional information.

IPsec Mobile Clients Tab

The **Mobile Clients** tab under **VPN > IPsec** contains settings which influence the authentication and configuration of mobile clients. These are specific to mobile tunnels and separate from the typical phase 1 and phase 2 negotiation.

See also:

- *Mobile IPsec*
- *Choosing a Mobile IPsec Style*
- *Remote Access Mobile VPN Client Compatibility*

Warning: The behavior of these fields varies by the type of mobile connection in use. Check type-specific documentation such as those linked on *Choosing a Mobile IPsec Style* for recommendations relevant to a given use case.

Enable

IKE Extensions
When checked, enables additional options for mobile IPsec clients in the GUI. The GUI will prompt to create an IPsec phase 1 entry for mobile connections if one does not already exist.

Extended Authentication

User Authentication
Specifies which authentication sources will be used when authenticating mobile users.

This list contains *Local Database* which is users from the *User Manager* or the *IPsec Pre-Shared Keys Tab*, as well as RADIUS and LDAP servers defined on the firewall.

See also:

See *Authentication Servers* for information on managing entries for this list.

Group Authentication
When set, the authentication process checks group membership of the user and their privileges. The groups are checked for either "User - VPN: IPsec with Dialin" or "WebCfg - All pages" privileges.

Authentication Groups
A list of available *groups from the user manager*. Only members of the selected groups will pass authentication.

When using EAP-RADIUS, group membership is determined by responses from a RADIUS server. See *RADIUS Groups* for details.

Note: The IPsec daemon only allows the specification of a single group for a user in the `Class` attribute, while pfSense® software allows specifying multiple semicolon-delimited groups. See the strongSwan eap-radius Plugin documentation for more details.

RADIUS Accounting
When enabled, the IPsec daemon will attempt to send RADIUS accounting data for **all** tunnels, not only connections associated with mobile IPsec. The RADIUS server must be selected in the **User Authentication** list above.

> **Warning:** Do not enable this option unless the selected RADIUS servers are always online and capable of receiving RADIUS accounting data. Tunnels will be disconnected if RADIUS accounting data is enabled and fails to send, even if they are not mobile clients.

Client Configuration

Virtual Address Pool
Defines the pool of IP addresses from which dynamic client addresses are assigned. For example, `10.3.200.0/24`.

Note: This subnet must not already be in use on an interface or elsewhere on the firewall or local network.

Virtual IPv6 Address Pool
Same as above, but for IPv6 addresses.

RADIUS IP Address Priority
When set, the IPv4/IPv6 address pool is used if address is not supplied by RADIUS server. When unset, a client only receives an IP address if RADIUS provides one.

RADIUS Advanced Parameters
Advanced options for tuning RADIUS behavior. Typically, these are only required when under high load or when using two-factor or similar out-of-band authentication methods with RADIUS.

Retransmit Base
Base to use for calculating exponential back off. Default value is `1.4`.

Retransmit Timeout
Timeout in seconds before sending first retransmit. Default value is `2`.

Retransmit Tries
Number of times to retransmit a packet before giving up. Default value is `4`.

Sockets
Maximum number of sockets (ports) to use when communicating with a RADIUS server. Increase for high load environments with many frequent authentication requests. Default value is `1`.

Network List
Controls whether the client will attempt to send all of its traffic across the tunnel or only traffic for specific networks.

If this option is checked the networks defined in the **Local Network** options for the mobile phase 2 definitions will be sent to the client. If this option is unchecked clients will attempt to send **all** of their traffic, including Internet traffic, across the tunnel.

Note: Not all clients and mobile IPsec modes respect this option. Some, such as Windows, require routes to be added on the client side in certain configurations like IKEv2.

Save Xauth Password
When checked, clients compatible with this control message will allow the user to save their credentials when manually entered during Xauth authentication.

This is mainly respected by Cisco-based Xauth clients like those found on iOS and macOS.

DNS Default Domain
When checked, the GUI offers a text box for a value the firewall will push to clients as their default domain suffix for DNS requests.

For example if this is set to `example.com` and a client requests `host`, then the DNS request will be for `host.example.com`.

Split DNS
Controls how the client will send DNS requests to the DNS Server supplied (if any).

The following behaviors are available:

- If this option is unchecked, the client will send all of its DNS requests to the provided DNS Server(s).
- If the option is checked and the text box is empty, and a **DNS Default Domain** is set, then only requests for that domain name will go to the provided DNS Server(s).

- If the option is checked and a value is entered in the text box, then only requests for the domain(s) entered in the box will be forwarded to the provided DNS Server(s).

DNS Servers
When **Provide a DNS server list to clients** is checked and DNS server IP addresses are entered, the firewall sends these values to clients for making DNS requests while the VPN is connected.

Note: If mobile clients will route to the Internet over the VPN, ensure the clients get a DNS Server from the firewall using this option and that they do not have **Split DNS** enabled.

Without this configuration clients will send DNS requests to servers they were assigned by their ISP, however, clients will route the request across the tunnel and the queries will most likely fail.

WINS Servers
Works similar to DNS servers, but for WINS.

WINS is rarely used on modern networks and is best left disabled.

Phase 2 PFS Group
Overrides the PFS setting for all mobile phase 2 entries.

The best practice is to set this value on the phase 2 entries individually.

Login Banner
A brief bit of text sent to the client for display after the login process succeeds. Optional and only works on some Xauth clients.

IPsec Pre-Shared Keys Tab

The **Pre-Shared Keys** tab under **VPN > IPsec** defines key and identifier pairs which are used for authenticating IPsec tunnels. Primarily this is intended for use with mobile IPsec, but there are occasional use cases for site-to-site tunnels as well.

Identifier
A string used to identify a peer. This is typically a username, a hostname, an E-mail address, or an IP address.

Secret Type
The type of secret to associate with this identity. It can be one of two types:

PSK
A traditional pre-shared key for use with most IKEv1 mobile IPsec configurations, site-to-site tunnels, and similar use cases.

EAP
An EAP key for use with IKEv2 mobile IPsec EAP-MSCHAPv2 authentication.

Pre-Shared Key
The contents of the key. As with a pre-shared key on an IPsec tunnel, this should be as long and complex as feasible. However, since this may be manually entered by a human in a manner similar to a password it might need to be more user-friendly than the key for a site-to-site tunnel.

Warning: The contents of these passwords must be known to the IPsec daemon, and thus they must be stored in plain text (*Password Storage Security Policies*). If this is not acceptable, consider using RADIUS-based authentication instead.

Additional options are available for EAP type keys:

Identifier Type
Manually sets the type of the **Identifier** field to override automatic behavior.

See also:

See *Phase 1 Proposal (Authentication)* for explanations of the different identifier types.

Virtual Address Pool
A static IP address to assign to this particular peer. Leave blank to assign a random address from the pool defined on the **Mobile Clients** tab.

> **Warning:** This configuration creates a new pool which **must not** overlap existing pools. As such, this address must be **outside** the pool defined on the **Mobile Clients** tab and **different from** any other pool defined on other PSK tab entries or Mobile Group Pool entries (Plus only).

DNS Server
A DNS server that the firewall will push to only this peer. Leave blank to use the DNS server value(s) from the **Mobile Clients** tab.

IPsec Mobile Group Pools Keys Tab (Plus Only)

The **Mobile Group Pools** tab under **VPN > IPsec** defines additional group mappings to address pools. This is most useful when utilizing EAP-RADIUS and supplying a Class attribute to identify pool membership for an authenticated user. Additional Mobile Group Pools are only used when **Group Authentication** is enabled in the **Mobile Clients** tab under **VPN > IPsec**

Groups
A string of comma separated group names. Must not overlap with the Authentication Groups selected on the **Mobile Clients** tab or the Groups of any previously configured Mobile Group Pool.

Virtual Address Pool
A subnet of addresses to assign to peers which map to this Mobile Group Pool.

> **Warning:** This configuration creates a new pool which **must not** overlap existing pools. As such, this address must be **outside** the pool defined on the **Mobile Clients** tab and **different from** any other pool defined on other PSK tab entries or Mobile Group Pool entries (Plus only).

DNS Server
A DNS server that the firewall will push to peers in this group. Leave blank to use the DNS server value(s) from the **Mobile Clients** tab.

Advanced IPsec Settings

The **Advanced Settings** tab under **VPN > IPsec** contains options which control IPsec daemon behavior and how traffic is handled with IPsec.

IPsec Logging Controls
These options control which areas of the IPsec daemon generate log messages and their level of detail. For information on viewing the log, see *IPsec Logs*.

In most cases the optimal settings are the default: **IKE SA**, **IKE Child SA**, and **Configuration Backend** set to *Diag*, and all others set to *Control*.

Configure Unique IDs as
Controls how the IPsec daemon treats new connections with an identifier which matches an existing connection. In most cases a new connection is intended to replace an older connection, but certain use cases such as mobile clients may require multiple connections from the same remote identifier.

Yes (Replace)
The new connection is accepted by the IPsec daemon, and it replaces the old connection, which is disconnected.

No
The new connection is accepted, and the old connection is replaced only if the peer sends an `INITIAL_CONTACT` notification.

Never
The new connection is always allowed, and `INITIAL_CONTACT` notifications are ignored.

Keep
The new connection is rejected, and the old connection remains active.

IPsec Filter Mode
Experimental. Controls how the firewall filters IPsec traffic.

Filter IPsec Tunnel, Transport, and VTI on IPsec tab (enc0)
The default behavior. Rules on the IPsec tab filter all IPsec traffic, including tunnel mode, transport mode, and VTI mode.

This is limited in that it **does not** allow for filtering on assigned VTI or transport mode interfaces, and does not allow for NAT or `reply-to` to function for VTI rules. It can also cause problems with connection state tracking for transport mode traffic.

> **Warning:** This mode requires special changes to the rules to work around incompatibilities between the default firewall state policy and the way VTI traffic is handled by the OS. See *IPsec VTI Filtering* for details.

Filter IPsec VTI and Transport on assigned interfaces, block all tunnel mode traffic
Enables firewall rules for assigned VTI and transport mode interfaces, NAT on VTI interfaces, and `reply-to` for rules on assigned VTI interface tabs. This also allows transport mode to properly filter traffic in both directions, such as with GRE tunnels protected by transport mode IPsec.

However, when set to filter on assigned VTI interfaces, **all tunnel mode traffic is blocked**.

> **Warning:** Do not set this option unless all IPsec tunnels are using VTI or Transport Mode.
>
> This option is incompatible with mobile IPsec as mobile IPsec is only capable of using tunnel mode.

IP Compression
Propose IPComp content compression for data carried across the VPN.

> **Warning:** Though the option is present in the GUI, the underlying operating system does not yet fully implement IP compression.

Enable PKCS#11 Support
When set, enables PKCS#11 token functionality in IPsec. This includes activating the `pcscd` daemon and enabling GUI controls in IPsec phase 1 for activating PKCS#11 authentication.

Strict Interface Binding
When set, the IPsec daemon configuration binds only to the interfaces required by the configuration, rather than binding to all interfaces.

This option is more secure but is known to break with interfaces which have dynamic IP addresses. Only enable this option in environments where it has been lab tested and proven to work as intended.

Unencrypted Payloads in IKEv1 Main Mode
Some IPsec implementations send the third Main Mode message unencrypted, probably to find the PSKs for the specified ID for authentication. This is similar to Aggressive Mode, and has the same security implications: A passive attacker can sniff the negotiated Identity, and start brute forcing the PSK using the HASH payload. The best practice is to keep this option disabled unless the implications are fully understood and compatibility to such devices is required (for example, some SonicWall devices).

Maximum IKEv1 Phase 2 Exchanges
IKEv1 phase 2 can initiate multiple rekey events for one VPN gateway in parallel. By default, it only allows `3` parallel rekey events. Undersized values can break VPN connections with many phase 2 definitions. If unsure, set this value to match the largest number of phase 2 entries on any phase 1.

Enable Cisco Extensions
Enables the Unity plugin which provides compatibility with Cisco Extensions such as `Split-Include`, `Split-Exclude`, and `Split-DNS` for IKEv1 XAuth mobile clients. This allows compatible clients to obtain values automatically when connecting to a mobile IPsec VPN.

Strict CRL Checking
When set, the IPsec daemon requires availability of a fresh CRL for peer authentication based on certificate signatures to succeed. Primarily useful when the CRL is obtained dynamically (e.g. OCSP).

> **Warning:** If there is no CRL available for a CA, validation will fail.

Make Before Break
Controls whether IKEv2 Reauthentication uses Make-before-Break or Break-before-Make ordering when an IKE Security Association (SA) expires.

Only relevant for IKEv2 tunnels using reauthentication, it does not affect IKEv1 tunnels or IKEv2 tunnels set to rekey.

Break-before-Make (Unchecked, Default)
Deletes IKE and Child SAs before reauthenticating and making a new set of SAs. This behavior is standard and widely compatible, but disruptive as there is a small gap between the old and new SA set in which IPsec connectivity is unavailable.

Make-before-Break (Checked)
Reauthenticates and makes a new SA set before deleting the old SA set. This eliminates the connectivity disruption, but requires that both endpoints allow overlapping IKE and Child SA entries.

Asynchronous Cryptography
Allows cryptographic framework jobs to be dispatched in a multithreaded manner to increase performance. Jobs are handled in the order they are received so that packets will be re-injected in the correct order.

> **Warning:** This option can increase performance, but may be unstable on certain hardware. When enabling this option, test connectivity during a maintenance window to ensure proper behavior. See Bug #8964 for details.

Custom Ports
Rare situations may require the firewall to listen for inbound IPsec packets on alternate port numbers for IKE and NAT-T. These settings can accommodate such cases, but affect every tunnel on the firewall.

Leave empty for the default behavior, which is to use UDP port `500` for IKE and `4500` for NAT-T.

Auto-exclude LAN Address
Set up an automatic IPsec bypass for traffic to and from the LAN subnet, so it does not get captured by policy-based IPsec.

Additional IPsec Bypass
Configures additional manual IPsec bypass behavior. When set, the GUI exposes the **IPsec Bypass Rules** control.

IPsec Bypass Rules
Custom rules which allow traffic matching combinations of **Source Address** and **Destination Address** pairs to be excluded from IPsec policies.

Source Address
The source address or network to exclude, and its mask.

Destination Address
The corresponding destination address or network to exclude, and its mask.

Note: These values are considered *together*. A packet must match **both** the source and destination to bypass IPsec policies.

These rules are useful to exclude traffic between multiple local networks, especially when a policy-based IPsec tunnel is set to use `0.0.0.0/0` as the remote network.

20.8.3 Choosing a Mobile IPsec Style

Only one type of mobile IPsec may be configured at a time, though there are multiple different styles to choose from.

- *IKEv2 with EAP-MSCHAPv2* for local username and password authentication
- *IKEv2 with EAP-RADIUS* for remote username and password authentication
- *IKEv2 with EAP-TLS* for per-user certificate authentication
- *Xauth+PSK* for local or remote username and password authentication
- Xauth+RSA for certificates and local or remote username and password authentication
- *Pre-Shared Key* for basic IPsec connectivity from older clients
- *L2TP/IPsec* for local or remote username and password authentication with clients that are not compatible with one of the above methods.

See also:

- *Configuring IPsec IKEv2 Remote Access VPN Clients on Windows*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on Ubuntu*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on Android*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on macOS*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on iOS*

IKEv2 is the best available choice. As of this writing, most **current** operating systems natively offer IKEv2 clients or can use an app or add-on client. Windows 7 and later, Android 11 and later, macOS 10.11 (El Capitan) and later, iOS 9 and later, and most Linux distributions have built in clients compatible with IKEv2. There is also a simple-to-use strongSwan IKEv2 app for various operating systems including older versions of Windows as well as Android 4.x and later.

Note: All IKEv2 types require a certificate structure including at least a Certificate Authority and a Server Certificate, and in some cases user certificates. For more information on Certificates, see *Certificate Management*. Clients can be picky about certificate attributes, so pay close attention to this chapter when creating the certificate structure.

Warning: When generating a **Server Certificate** for use with IKEv2, the **Common Name** of the certificate must be the hostname of the firewall as it exists in DNS. The name must be repeated as an FQDN type Subject Alternative Name (SAN). The SAN step is handled automatically by the Certificate Manager on current versions of pfSense® software. The IP address of the firewall must also be present as an IP Address type SAN when possible. This information will be repeated later when relevant, but requires extra emphasis due to its importance. See *Create a Server Certificate*

IKEv2 with EAP-MSCHAPv2

IKEv2 with EAP-MSCHAPv2 is an ideal choice for current operating systems. Though there are several variations, EAP-MSCHAPv2 is the easiest to configure since it does not require generating or installing per-user certificates and does not require a working RADIUS server. The CA Certificate must still be installed onto the client as a trusted root certificate.

Tip: Many current clients will also work with a server certificate generated by the *ACME Package*. ACME certificates are already trusted by most clients natively, thus using ACME will bypass the need to install a CA entry on those clients.

EAP-MSCHAPv2 allows for username and password authentication using passwords stored on the **Pre-Shared Keys** tab under **VPN > IPsec** (*IPsec Pre-Shared Keys Tab*). These passwords are stored in plain text, so it is not as secure as using a RADIUS server, though it is more convenient.

See also:

- *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2*

IKEv2 with EAP-RADIUS

EAP-RADIUS works identically to EAP-MSCHAPv2 except that user authentication happens via RADIUS. When EAP-RADIUS is chosen, a RADIUS server must be selected on the **Mobile Clients** tab. The RADIUS server must accept and understand EAP requests, and it must also allow MSCHAPv2. Password security is left up to the RADIUS server. VPN access can be optionally limited by *RADIUS group* membership using the **Group Authentication** options on the **Mobile Clients** tab.

EAP-RADIUS is typically the best choice when a RADIUS server is available.

See also:

- *IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS*

IKEv2 with EAP-TLS

EAP-TLS uses per-user certificate authentication instead of username and password authentication. As such, EAP-TLS requires generating certificates for each user, which makes it a bit more cumbersome from an administration standpoint. Certificates are validated against the CA similar to OpenVPN. The CA certificate, user certificate and its associated key must all be imported to the client properly.

Warning: When creating user certificates, the username must be used as the certificate common name and as a DNS/FQDN type Subject Alternative Name. If the same name is not present in both places, clients may not be validated properly. This is handled automatically by the Certificate Manager on current versions of pfSense software.

See also:

- *IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS*

IKEv1 with Xauth and Pre-Shared Keys

Xauth+PSK works on a majority of platforms, the notable exception being current versions of Android. Windows XP through Windows 8 can use the Shrew Soft client, but Windows 10 does not currently work with any available client. macOS and iOS can use their built-in client to connect.

Note: When using Xauth, local users must exist in the **User Manager** and those users must have the `User - VPN - IPsec Xauth Dialin` privilege.

See also:

- *IPsec Remote Access VPN Example Using IKEv1 with Xauth*

IKEv1 with Xauth and RSA Certificates

Xauth+RSA works in most of the same conditions as Xauth+PSK, though it does in fact work on current Android devices. Certificates must be made for each user, and the certificates must be imported into the clients.

IKEv1 with Pre-Shared Keys Only

Pre-Shared Key only IPsec VPNs for mobile IPsec have become rare in modern environments. Implementations were rare, it was only found in the Shrew Soft client, some very specific Android versions such as those from Motorola, and in other third-party clients. They are not very secure, and are not a good practice. The only time they may be needed is in cases when client devices cannot use any other method.

See also:

- *IPsec Remote Access VPN Example Using IKEv1 with Pre-Shared Keys*

L2TP/IPsec (IKEv1)

L2TP/IPsec is a unique combination that, unfortunately, does not work very well in practice. In this style of setup mobile IPsec is set up to accept Transport Mode connections which secure all traffic between the public IP address endpoints of clients and the firewall. An L2TP connection is made across this transport mode channel to tunnel user traffic in a more flexible way. Though clients for this method are found in most versions of Operating Systems, they are all implemented in different and incompatible ways.

For example, the Windows client does not work properly when the client system is behind NAT, which is the most common place that a VPN client would find itself. The problem is in an interaction between the client and the IPsec daemon used on pfSense, strongSwan. The strongSwan project states that it is a bug in the Windows client, but it is unlikely to be fixed since both strongSwan and Windows have focused their mobile client efforts on more modern and secure implementations such as IKEv2 instead.

Warning: L2TP/IPsec should be avoided when possible.

See also:

- *L2TP/IPsec Remote Access VPN Configuration Example*

20.8.4 NAT with IPsec Phase 2 Networks

pfSense® software can utilize NAT on policy-based IPsec phase 2 entries to make the local network appear to the remote peer as a different subnet or address. This can work around subnet conflicts or connect to vendors without renumbering a local network.

Warning: NAT is not currently compatible with route-based VTI IPsec tunnels without configuring an **IPsec Filter Mode** which is incompatible with tunnel-based IPsec. See *Advanced IPsec Settings* for details.

Configuration

NAT is configured by the **NAT/BINAT Translation** options on an IPsec phase 2 entry in tunnel mode, in combination with the **Local Network** settings.

Local Network
Values of **Type** and **Address** specify the **actual** local network (e.g. LAN subnet).

NAT/BINAT Translation
Values of **Type** and **Address** specify the **translated** network visible to the far side.

NAT Types

There are two main modes for NAT with IPsec:

Binat - 1:1 NAT
When both the actual and translated local networks use the same subnet mask, the firewall will directly translate the networks to one another inbound and outbound. Can also be used for single addresses.

This allows remote host to directly contact local hosts using their equivalent NAT addresses, provided that IPsec rules allow the traffic to pass.

NAT - Overload/PAT Style
If the Local Network is a subnet, but the **NAT/BINAT Translation** address is set to a single IP address, then a 1:many NAT (PAT) translation is set up that works like an outbound NAT rule on WAN. All outbound traffic will be translated from the local network to the single IP address in the NAT field.

Note: Inbound traffic from the remote network to individual local hosts is not possible in this mode.

Warning: NAT+IPsec cannot be configured between two different sized subnets (e.g. It cannot NAT a /24 size subnet to a /27 size subnet).

Example

Consider an IPsec tunnel to a Vendor which requires `172.16.5.0/24` for the network on this firewall. However, the LAN is actually `192.168.1.0/24`, and renumbering is not feasible.

To accommodate this scenario, set the phase 2 values as follows:

Local Network

Type
Network

Address
`192.168.1.0/24`

NAT/BINAT Translation

Type
Network

Address
`172.16.5.0/24`

Firewall Rules

NAT is processed before firewall rules, so firewall rules on the IPsec tab refer to the network in **Local Network**.

Remote End Notes

The far side of the tunnel does not need any knowledge of the actual **Local Network**. Their tunnel is built between their local network and the **NAT/BINAT Translation** value.

Packet Capturing Quirk

In a packet capture, the **Local Network** addresses are shown on outbound traffic, not the translated address. This does not indicate any problem.

20.8.5 Routed IPsec (VTI)

Route-based IPsec is an alternative method of managing IPsec traffic. It uses `if_ipsec(4)` from FreeBSD for Virtual Tunnel Interfaces (VTI) and traffic is directed using the operating system routing table. It does not rely on strict kernel security association matching like policy-based (tunnel mode) IPsec.

A routed IPsec tunnel creates an `ipsecX` interface at the operating system level and this interface has its own IP address. The `ipsecX` interface must be assigned so it can be used for purposes such as static or dynamic routing, daemon binding, traffic monitoring, and so on.

Once assigned, the IPsec interface also gains an automatic gateway which provides policy routing and gateway group capabilities.

Note: Routed IPsec is not replacing traditional tunnel mode IPsec or transport mode. The mode choice is up to the user when creating an IPsec phase 2 entry. Any mode may be used at the same time, subject to the caveats listed later in this document.

See also:

The Hangouts Archive contains a video which covers Routed IPsec.

Prerequisites

First pick a transit network. This is similar to choosing a tunnel network for a WireGuard, GRE, or OpenVPN instance. Typically, this is a `/30` network in an unused subnet. This example uses `10.6.106.0/30`.

IPsec Configuration

- Create an IPsec Phase 1 entry as usual
- Create a Phase 2 entry under this Phase 1 with the following settings:

  Description
  Some useful relevant text (e.g. `HQ VTI Tunnel`)

  Mode
  Routed (VTI)

  Local Network Address
  `10.6.106.1`

  Remote Network Address
  `10.6.106.2`

  Proposal
  Set as needed to match the other end.

- Click **Save**
- Click **Apply Changes**

In most cases only a single phase 2 entry is necessary as all traffic for a specific address family can be carried over a single interface. A common use case for a second phase 2 is to configure an IPv6 address if the first is configured for IPv4 (or vice versa).

When using IKEv2 without split connections these addresses will all use the same underlying `ipsecX` interface. With IKEv1 or with IKEv2 and split connections active then each phase 2 will be a separate `ipsecX` interface which must be assigned and used independently.

Tip: The default case of IKEv2 without split connections is the best practice. Avoid using the other cases unless required by a third party.

IPsec Interface Assignment

- Navigate to **System > Routing**
- Set the **Default gateway** options to a specific gateway or group, as long as they are not left at *Automatic* (*Managing the Default Gateway*)

> **Warning:** If the default gateway remains set to *Automatic* the firewall may end up using the IPsec VTI interface as the default gateway, which is unlikely to be the desired outcome.

- Navigate to **Interfaces > Assignments**
- Pick the new `ipsecX` interface from the **Available Network Ports** list
- Click **+ Add**
- Note the new interface name, e.g. *OPT1*
- Navigate to **Interfaces > [New Interface Name]**
- Check **Enable**
- Give the interface a more suitable name using the **Description** field (e.g. `VTI_HQ`)
- Click **Save**
- Click **Apply Changes**

The firewall creates a gateway automatically which can be used for static routing, policy routing, and so on.

At this point the interface is available for use like any other interface. It can be used for packet captures, traffic graphs, binding daemons, routing protocols, and so on.

Routing

No traffic will attempt to cross the IPsec tunnel until routing is configured except for gateway monitoring probes (if enabled).

Static Routes

To set up static routes navigate to **System > Routing**, **Static Routes** tab. Add new routes there using the assigned IPsec interface gateway.

Typically, there will be one static route per remote destination network, similar to how there would be one phase 2 entry per remote destination network with tunnel mode IPsec.

Dynamic Routes

As an alternative to manually managing static routes, assigned IPsec VTI interfaces can be used with the *FRR Package* for dynamic routing such as BGP and OSPF.

Policy Routes

To policy route traffic across a routed IPsec tunnel, use the assigned IPsec interface gateway in firewall rules as usual for policy routing.

Note: This may not work as expected without NAT and/or `reply-to`, which require special settings. See *Routed IPsec Firewall Rules* for details.

See also:

Policy Routing Configuration

Routed IPsec Firewall Rules

By default, routed IPsec traffic appears to the OS on both the per-tunnel `ipsecX` interface and the `enc0` interface. When set this way **traffic must be passed on the IPsec tab**.

This can be changed, however. The behavior of firewall rules for traffic inside an IPsec tunnel depends on the **IPsec Filter Mode** option in the *Advanced IPsec Settings*.

To utilize features such as per-interface rules, NAT, and `reply-to` with routed IPsec the **IPsec Filter Mode** option must be set to filter on assigned interfaces. This option is not compatible with tunnel mode so it is only feasible if all tunnels on the firewall are using VTI or transport mode.

> **Warning:** Staying on the default filter mode requires special changes to the rules to work around incompatibilities between the default firewall state policy and the way VTI traffic is handled by the OS. See *IPsec VTI Filtering* for details.

See also:

- *IPsec and firewall rules*
- *Advanced IPsec Settings*

Caveats

Routed IPsec works best when both sides are compatible with routed IPsec. It can still work when only one side implements routed IPsec, but most of its benefits are lost.

Rather than managing IPsec Phase 2 entries, routes must be managed instead. Since this can be automated with dynamic routing protocols this is not a large concern.

Firewall rule processing can be confusing as mentioned in *Routed IPsec Firewall Rules*. Some features such as NAT require special settings to function.

20.8.6 IPsec and firewall rules

Outer IPsec Traffic

pfSense® software automatically adds hidden firewall rules which allow traffic required to establish enabled IPsec tunnels. The traffic required to establish a tunnel includes:

- UDP port `500` (or a custom configured **Remote IKE Port** on a tunnel)
- UDP port `4500` (or a custom configured **Remote NAT-T Port** on a tunnel)
- The *ESP* protocol

The automatic rules restrict the source to the **Remote Gateway** IP address (where possible) destined to the Interface IP address specified in the tunnel configuration. When IPsec mobile client options are enabled, the same firewall rules are added except with the source set to *Any*.

To override the automatic addition of these rules check **Disable all auto-added VPN** rules under **System > Advanced** on the **Firewall & NAT** tab. When that box is checked firewall rules must be manually added to allow appropriate traffic on the correct interface(s) from the expected source(s).

Tunneled IPsec Traffic from Remote to Local

The behavior of firewall rules for traffic **inside** an IPsec tunnel depends on the **IPsec Filter Mode** option in the *Advanced IPsec Settings*.

Filtered on IPsec Tab

By default, traffic passed inside a tunnel from the remote end is filtered by rules configured under **Firewall > Rules** on the **IPsec** tab (`enc0`). Those rules allow and restrict resources made accessible to remote IPsec users.

Note: By default, all traffic from remote VPN hosts is blocked as there are no rules on the IPsec tab until they are manually added by a firewall administrator.

In this default mode traffic for transport and VTI mode tunnels does not always behave in a desirable way. This mode prevents VTI from using per-interface rules, NAT, or `reply-to`; transport mode can have issues tracking state properly.

Filtered on Assigned IPsec Interfaces

If all tunnels on the firewall are VTI or transport mode, then set the **IPsec Filter Mode** to filter on assigned interfaces instead. When set this way, assigned VTI interfaces can use per-interface rules, NAT, and `reply-to` as one would typically expect. Additionally, transport mode filtering works as expected with rules on the interfaces involved in transport mode (e.g. WAN, tunneling protocols like GRE, etc.).

The downside of this mode is that all tunnel mode traffic is dropped and only VTI or transport mode traffic can be filtered as it is handled on separate interfaces (e.g. `ipsec1`, not the shared `enc0` interface).

Tunneled IPsec Traffic from Local to Remote

To control traffic in the other direction, from local networks to remote IPsec VPN connected devices or networks, use rules on the **local** interface where the **local** device resides. For example, connectivity from hosts on LAN to VPN destinations is controlled by rules on the **LAN** tab.

20.8.7 Using IPsec with Multiple Subnets

pfSense® software handles multiple IPsec networks using separate IPsec phase 2 entries which define source and destination pairs to pass through a tunnel.

For example, to accommodate the table below, define two Phase 2 entries on both sides:

Site A	Site B
`172.16.0.0/24`	`10.0.0.0/24`
`172.16.1.0/24`	

On the Site A Firewall:

- `172.16.0.0/24` to `10.0.0.0/24`
- `172.16.1.0/24` to `10.0.0.0/24`

On the Site B Firewall:

- `10.0.0.0/24` to `172.16.0.0/24`
- `10.0.0.0/24` to `172.16.1.0/24`

This works for any additional networks on either side, such as multiple local interfaces, mobile VPN clients, networks on the other end of VPNs connected to the remote router, etc.

Supernetting Example

If the equipment to which the tunnel connects does not allow multiple phase 2 entries, it may be necessary to employ supernetting/CIDR summarization to fit the networks into a single phase 2.

Tip: This technique can also be used to keep the phase 2 list shorter and more manageable.

It is also more reliable to use fewer phase 2 entries as there is less to negotiate when building or rekeying tunnels. Note that this only applies to IKEv1 or IKEv2 with Split Connections, as IKEv2 uses a single child SA by default.

For example, consider the networks in the following table:

Site A	Site B
`192.168.0.0/24`	`10.0.0.0/24`
`192.168.1.0/24`	
`192.168.2.0/24`	

Due to the fact that the subnets are close to each other they can be grouped into a larger network in the tunnel definition: `192.168.0.0/22`.

Note: This larger subnet also includes `192.168.3.0/24` due to the way subnet math and boundaries work. If that is unacceptable, then an alternate solution may be required, such as isolating that subnet with firewall rules or using different equipment at the other endpoint.

20.8.8 Configuring IPsec Keep Alive

There are two methods which can make the firewall attempt to keep a non-mobile IPsec tunnel up and active at all times: automatic ping and periodic check. These options are available in the settings for each IPsec phase 2 entry.

See also:

See *Keep Alive* for additional details on these settings.

Automatic Ping

This method utilizes ICMP echo requests sent to a specific remote host across the VPN to match policies which will start a tunnel and keep it active.

For tunnel mode (policy-based) IPsec tunnels traffic destined to the **Remote Network** will attempt to initiate the tunnel when it is down. This is because the generated ping will match trap policies in the kernel and be considered "interesting traffic" for IPsec.

Warning: Due to the reliance on policies this method is not capable of initiating a VTI mode tunnel. It can send periodic traffic across a VTI mode tunnel if a use case requires that behavior.

This option will also not initiate a tunnel if its phase 1 **Child SA Start Action** is set to *Responder Only*.

Unlike other mechanisms such as DPD, this periodic traffic sent across the tunnel is treated like other traffic crossing the tunnel. This traffic would count as tunnel activity and reset any idle counters on the far side.

Note: Any IP address within the **Remote Network** of the phase 2 definition may be used. It does not have to reply or even exist.

Warning: For this feature to work the firewall **must** have an IP address assigned inside the **Local Network**. Otherwise, it cannot generate the necessary traffic to match the phase 2 policies and traffic cannot enter the tunnel.

Periodic Check

This method utilizes a periodic status check which looks at the list of connected IPsec tunnels and will initiate entries which are not currently connected.

As this does not rely on tunnel traffic or trap policies it is compatible with any IPsec tunnel mode, including VTI mode.

IKEv1 vs IKEv2

Whether this option should be enabled on every phase 2 entry for a tunnel depends on the tunnel configuration.

IKEv1 or IKEv2 with Split Connections
: In these modes each phase 2 entry results in a separate child SA entry which can be connected separately. In this case, the keep alive options may be set on each phase 2 entry individually as needed. If all phase 2 entries must stay connected, then it must be enabled on every entry.

IKEv2 without Split Connections
: In this mode the phase 2 entries are combined into a single child SA entry and all combinations of phase 2 entries are connected as a single group. In this case the keep alive options need only be enabled on the **first** phase 2 entry for a tunnel.

See also:

See *IPsec phase 1 Advanced Options* for more information on how the **Split Connections** option works.

20.8.9 Testing IPsec Connectivity

The easiest test for an IPsec tunnel is a ping from one client station behind the firewall to another on the opposite side. If that works, the tunnel is up and working properly.

As mentioned in *Accessing Firewall Services over IPsec* traffic initiated from pfSense® software will not normally traverse a tunnel without extra routing. That said, there is a quick way to test the connection from the firewall itself by manually specifying a source address when issuing a ping.

There are two methods for performing this test: the GUI, and the shell.

Specifying a Ping Source in the GUI

In the GUI, a ping may be sent with a specific source as follows:

- Navigate to **Diagnostics > Ping**
- Fill in the settings as follows:

  Host
  Enter an IP address which is on the remote router within the remote subnet listed for the tunnel phase 2 (e.g. `10.5.0.1`)

  IP Protocol
  The address family of the host being used (e.g. *IPv4* for `10.5.0.1`)

  Source Address
  Select an interface or IP address on the local firewall which is inside the local Phase 2 network (e.g. Select *LAN* for the LAN IP address)

  Maximum number of pings
  Set an appropriate value which will be high enough to be meaningful yet low enough that it doesn't take too long to run. The default value of `3` ideal.

- Click **Ping**

If the tunnel is working properly ping replies will be received by the firewall from the LAN address at Site B. If replies are not received, move on to the *Troubleshooting IPsec VPNs* section.

Note: Typically, the first ping or two may be lost during tunnel negotiation, so the best practice is to use at least `3`.

If the first attempt did not produce any results, try again. If it still fails, try once more with a slightly higher **Maximum number of pings** value.

Specifying a Ping Source in the Shell

Using the shell on the console or via ssh, the ping command can be run manually, and a source address may be specified with the `-S` parameter. Packets generated by `ping` will not attempt to traverse the tunnel without using `-S` or a static route.

The syntax for a proper test is:

```
# ping -S <Local LAN IP Address> <Remote LAN IP Address>
```

Where the *Local LAN IP Address* is an IP address on an internal interface within in the local subnet definition for the tunnel, and the *Remote LAN IP Address* is an IP address on the remote router within the remote subnet listed for the tunnel.

In most cases this is the LAN IP address of the respective firewalls. For example, if the LAN IP address at site A is `10.3.0.1` and the LAN IP address at site B is `10.5.0.1`, then the following command would send a test ping from site A to site B:

```
# ping -S 10.3.0.1 10.5.0.1
```

If the tunnel is working properly, ping replies will be received by the firewall from the LAN address at Site B. If replies are not received, move on to the *Troubleshooting IPsec VPNs* section.

20.8.10 Client Routing and Gateway Considerations

When the VPN endpoint is the default gateway for a network there are normally no problems with routing. When a client PC sends traffic it will go to its default gateway, over the tunnel, and out the other end. However, if the firewall is not the default gateway for a given network, then other routing measures will need to be taken.

As an example, imagine that the firewall is the gateway at Site B but not Site A. This is illustrated in Figure *Site-to-Site IPsec Where the VPN Endpoint is not the Gateway*. The device PC1 at Site B sends a ping to PC2 at Site A. The packet leaves PC1, travels through the firewall at Site B, traverses the tunnel, exits the firewall at Site A, and on to PC2. What happens on the way back? The gateway on PC2 is a different router entirely. The reply to the ping will be sent to the gateway router and most likely be tossed out, or even worse, it may be sent out the Internet link and be lost that way.

There are several ways around this problem. The best practice depends on the circumstances of a given use case.

- A static route could be entered into the gateway router that will redirect traffic destined for the far side of the tunnel to the VPN endpoint.

  Even with this route, additional complexities are introduced because this scenario results in asymmetric routing as covered in *Bypass Firewall Rules for Traffic on Same Interface*.

- A static route could be added to the client systems individually so that they know to send that traffic directly to the VPN endpoint and not via their default gateway.

  Unless there are only a very small number of hosts that need to access the VPN, this is a management headache and should be avoided.

- Make the VPN endpoint the default gateway and let it handle the Internet connection instead of the existing gateway.

  In some situations this is easier than attempting to set up complicated routing, but it may not be acceptable to network administrators or management.

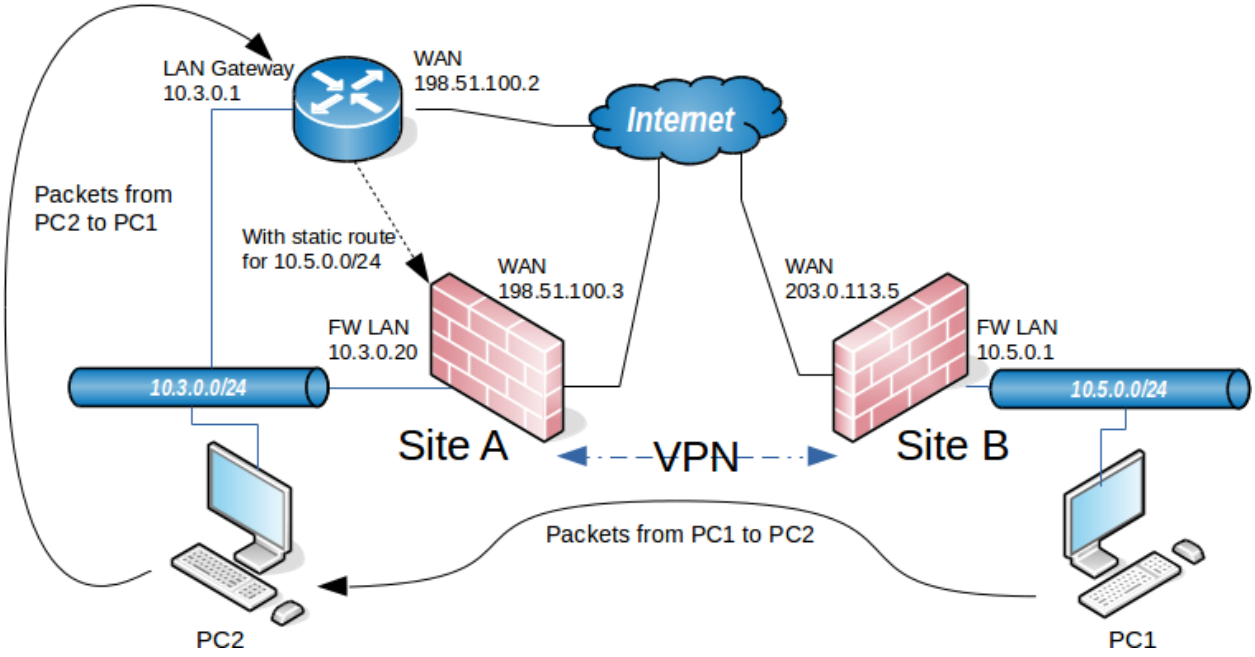


Fig. 3: Site-to-Site IPsec Where the VPN Endpoint is not the Gateway

20.8.11 Configuring Third Party IPsec Devices

Nearly any device compatible with IPsec VPNs can connect to a pfSense® software installation. pfSense software is used in production in combination with numerous vendors' equipment, and will most likely interoperate with any IPsec-capable devices.

Note: Connecting devices from two different vendors can be troublesome regardless of the vendors involved because of configuration differences between vendors, in some cases bugs in the implementations, and the fact that some of them use proprietary extensions. Some examples are provided at the end of this chapter for several common Cisco devices.

To configure an IPsec tunnel between pfSense software and a device from another vendor, the primary concern is to ensure that the phase 1 and 2 parameters match on both sides. For the configuration options on pfSense, where it allows multiple options to be selected, only select one of those options and ensure the other side is set the same. The endpoints will attempt to negotiate a compatible option when multiple options are selected, however that is frequently a source of problems when connecting to third party devices. Configure both ends to what are believed to be matching settings, then save and apply the changes on both sides.

Once the settings match on both ends of the tunnel, attempt to pass traffic over the VPN to trigger its initiation then check the IPsec logs on both ends to review the negotiation. Depending on the situation, the logs from one end may be more useful than those from the opposite end, so it is good to check both and compare. The pfSense software side typically provides better information in some scenarios, while on other occasions the other device provides more useful logging. If the negotiation fails, determine whether it was phase 1 or 2 that failed and thoroughly review the settings accordingly, as described in *Troubleshooting IPsec VPNs*. The side that is initiating often cannot see why, so check the logs on the responding side first.

Terminology Differences

Another frequent source of failures is differences in terminology between vendors. Here are a few common things to look out for:

Policy-Based VPN/IPsec
: The type of IPsec used by pfSense software in tunnel mode. Policies are defined which control traffic entering the tunnel (e.g. Phase 2 entries).

Encryption Domain, Policy, Proxy ID
: A network definition used in Phase 2 to control which traffic will be handled by IPsec.

  On some platforms, including Palo Alto, these "Proxy ID" definitions are used in Phase 2 (IPsec) negotiation to inform the remote peer about traffic selectors for networks involved in the VPN. In these cases, they help with negotiation between a device which only implements route-based IPsec and a device that only implements policy-based IPsec.

Route-Based VPN/IPsec
: The type of IPsec used by pfSense software in VTI mode. There is an IPsec interface which routes similar to other interfaces and obeys the routing table, rather than relying on policies.

S2S or L2L
: Short for Site-to-Site or LAN-to-LAN, distinguished from a mobile client style VPN.

Perfect Forward Secrecy (PFS)
: Some vendors have different names for PFS. It may only be a toggle which uses the same value as the Phase 1 **DH Group**, others label it with full text or the acronym, others label it **DH Group**. It may also be labeled **MODP** in some implementations.

Transform Set
On Cisco devices, a set of parameters which define Phase 2 handling such as encryption and hash algorithms.

ISAKMP Policy
On Cisco devices, a set of parameters that define Phase 1 (IKE) handling such as authentication, encryption, and hash algorithms, and others.

Proposals
On Juniper and FortiGate, sets of options that define parameters for Phase 1 (IKE) or Phase 2 (IPsec) handling.

NAT Exemption or no-NAT
On Juniper and Cisco, exceptions to NAT that must be made to ensure that traffic traversing a VPN does not have NAT applied. Not generally relevant to IPsec on pfSense software since NAT is not performed on IPsec by default.

Life bytes or Traffic Lifetime
Limits on the amount of traffic sent over a VPN before it renegotiates. Cannot be configured in the pfSense software GUI. If present on a remote device it may need to be disabled.

Compatible Devices

Nearly any device that implements standard IPsec can be connected with pfSense software. This page lists devices reported working by users, though it should not be considered complete.

Warning: Ensure firmware is up-to-date on devices before attempting to configure IPsec. Older devices and firmware may not be compatible with modern secure algorithms and standards.

- Adtran
- Cisco routers
- Cisco PIX and ASA firewalls
- Checkpoint NG
- D-Link VPN Routers
- DrayTek VPN routers
- IBM z/OS mainframes
- IPCop
- Juniper routers and firewalls
- Kerio Control
- LANCOM VPN Routers with LCOS
- Linksys VPN Routers
- m0n0wall
- MikroTik
- Nortel Contivity
- Palo Alto Networks
- SonicWall

- StoneGate Firewall/VPN
- Ubiquiti UniFi Security Gateway
- WatchGuard
- Zyxel firewalls

. . . and many more.

If a device is not listed and is known to work with pfSense software for IPsec, please *submit a documentation update*.

Consult the device documentation for IPsec configuration details.

20.8.12 Accessing Firewall Services over IPsec

With an out-of-the-box configuration it is not possible to query SNMP or other similar services on the LAN interface address of a remote firewall running pfSense® software over a tunnel mode IPsec VPN connection.

Note: Most of the advice in this document only applies to tunnel mode, not routed IPsec (VTI). See *Routed IPsec (VTI)* for details.

Fred Wright explained in a post to the m0n0wall mailing list on September 12, 2004, why this is, and it's the same reason here.

> Due to the way IPsec tunnels are kludged into the FreeBSD kernel, any traffic *initiated* by m0n0wall to go through an IPsec tunnel gets the wrong source IP (and typically doesn't go through the tunnel at all as a result). Theoretically this *shouldn't* be an issue for the *server* side of SNMP, but perhaps the server has a bug (well, deficiency, at least) where it doesn't send the response out through a socket bound to the request packet. You can fake it out by adding a bogus static route to the remote end of the tunnel via the m0n0wall's LAN IP (assuming that's within the near-end tunnel range). A good test is to see whether you can ping something at the remote end of the tunnel (e.g. the SNMP remote) *from* the m0n0wall. There's an annoying but mostly harmless side effect to this - every LAN packet to the tunnel elicits a no-change ICMP Redirect.

Most notably this is a problem for UDP services bound to all interfaces (*) and ICMP. In these cases, a reply uses the "closest" address to the client from the perspective of the system routing table. Without a route present matching the desired destination this results in using the IP address of the interface containing the default gateway.

Service Binding Workaround

Some services have options which make it possible to change the interface binding so that the daemon only binds to a specific interface or IP address (e.g. the IP address of the internal network on the local end of the VPN) on the firewall. The interface binding for SNMP, NTP, the DNS Forwarder, and several other services can be set in this way.

With the daemon bound to only that specific address, that is also the only address it can use to reply, and thus it can generate the expected replies which will properly take the IPsec path back to the other end.

Static Route Workaround

If changing the service binding is not possible, or for full connectivity between the endpoints, use static routes to work around the situation.

Note: When both endpoints must be accessed, a static route is required on each endpoint.

Setting up a static route is done by first adding a gateway pointing to the LAN IP address of the firewall (See *Gateways*) and then adding a static route using this gateway (See *Static Routes*).

Warning: Do not attempt to use this method with routed IPsec (VTI). See *Routed IPsec (VTI)* for details.

Add Gateway

First, add a gateway for the address of the firewall itself:

- Navigate to **System > Routing** on the **Gateways** tab
- Click  **Add** to add a gateway
- Configure the following settings:

  Interface
  LAN

  Name
  `IPsecGW` or another appropriate name, as desired.

  Gateway
  Enter the LAN IP address of **this** firewall.

  Disable Gateway Monitoring
  Checked

- Click **Save**
- Click **Apply Changes**

Add Route

Now add this route in the GUI:

- Navigate to the **Static Routes** tab
- Click  **Add**
- Configure the following settings:

  Destination Network
  The **remote** VPN network

  Gateway
  IPsecGW or whichever name was used when creating the gateway

Description
`Allow firewall itself to communicate over VPN`

- Click **Save**
- Click **Apply Changes**

Test

Perform a test as described in *Testing IPsec Connectivity* using the address of the far side system and the local address that system is attempting to query.

Routed IPsec (VTI)

The previous advice in this document does not generally apply to route-based IPsec (VTI mode) since it operates in a fundamentally different way than policy-based IPsec (tunnel mode).

With VTI mode IPsec the routing table contains proper routes back to the remote end using the VTI interfaces. Responses will utilize these routes to select the IPsec interface when looking up the path back across the tunnel.

This may mean the remote end should query the address of the IPsec interface directly so that the responses are received from a matching address, but whether this is necessary depends on the service and client.

The *Service Binding Workaround* can still help here if queries must be made to and from the LAN IP address directly, but the static route method must not be used as it will conflict with the routing necessary for VTI traffic to function.

See also:

- *IPsec Logs*
- *IPsec Status*
- *IPsec in Multi-WAN Environments*
- *IPsec in High Availability Environments*
- *IPsec Site-to-Site VPN Example with Pre-Shared Keys*
- *IPsec Site-to-Site VPN Example with Certificate Authentication*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS*
- *IPsec Remote Access VPN Example Using IKEv1 with Xauth*
- *IPsec Remote Access VPN Example Using IKEv1 with Pre-Shared Keys*
- *Routing Internet Traffic Through a Site-to-Site IPsec Tunnel*
- *Connecting to L2TP/IPsec from Android*
- *L2TP/IPsec Remote Access VPN Configuration Example*
- *Troubleshooting IPsec VPNs*

20.9 WireGuard

20.9.1 WireGuard Settings

WireGuard Package Settings

The WireGuard package contains the following configurable options:

Enable
Controls whether the WireGuard service itself is enabled or disabled.

Note: The WireGuard service cannot be disabled when one or more tunnels is assigned to an interface via *Interface Configuration*.

Keep Configuration
Controls whether the tunnel/peer configurations and package settings will persist when the package is removed.

Endpoint Hostname Resolve Interval
Controls how often peer endpoint hostnames are resolved and updated by the WireGuard service. By default, this is `300` seconds (5 minutes).

Track System Resolve Interval
This option overrides the **Endpoint Hostname Resolve Interval** setting and configures the WireGuard service to track and use the system *Aliases Hostnames Resolve Interval*.

Interface Group Membership
Controls which WireGuard tunnels are implicit members of the WireGuard interface group. By default, this is `All Tunnels`.

Tip: See *Rule Methodology* for more on **Interface Groups** and **Rule Processing Order**.

Hide Secrets
Controls whether secrets (private and pre-shared keys) are hidden in the user interface.

Warning: **Hide Secrets** only hides secrets in the user interface. It does not obfuscate secrets for storage in the pfSense® software configuration file, `config.xml`. For more information on password storage and protecting configuration file backups see *Password Storage Security Policies*

WireGuard Tunnel Settings

When creating or editing a WireGuard tunnel, the following options are available:

Enable
Controls whether this WireGuard tunnel is enabled or disabled.

Note: A WireGuard tunnel cannot be disabled while assigned as an interface.

Description
A short text description of this WireGuard tunnel.

Listen Port
The local port upon which this WireGuard tunnel will listen for incoming traffic from peers, and the port from which it will source outgoing packets. The default port is `51820`, additional tunnels must use a different port.

Note: The GUI will automatically suggest the next highest available port.

Interface Keys
The private and public key pair for this WireGuard tunnel. The public key is derived from the private key and does not need to be entered separately. The GUI will display the public key automatically when possible. When entering a new private key manually, the public key will be available after saving the tunnel.

The private key will stay only on this firewall, the public key will be copied to peers.

A new set of keys can be generated by the  **Generate** button.

Tip: Click **Copy** under the public key to copy it to the clipboard.

Interface Addresses
A list of IPv4 and/or IPv6 addresses which will be assigned to this WireGuard tunnel.

Note: Interface addresses are configured here only for WireGuard tunnels that are not assigned to an interface via *Interface Configuration*.

WireGuard Peer Settings

When creating or editing a WireGuard peer, the following options are available:

Enable
Controls whether this WireGuard peer is enabled or disabled.

Tunnel
Controls which WireGuard tunnel to associate with this peer. The default is *Unassigned*.

Tip: Peers can easily be staged or moved between tunnels using this option.

Description
A short text description of this peer.

Dynamic Endpoint
This option controls whether a WireGuard peer should be considered dynamic. Uncheck this option for a peer that has a fixed, static endpoint address or hostname.

Endpoint
The IP address or hostname of the remote WireGuard peer, from which the peer will connect to this firewall, and to which this WireGuard instance will send traffic destined for this peer.

This can be left empty if the peer endpoint is unknown, such as for dynamic remote access clients. When empty, the tunnel will track the endpoint dynamically based on the key used by the peer. Additionally, when empty, this firewall cannot initiate traffic on the tunnel to the peer until the remote peer sends traffic.

Endpoint Port
The port used by the peer for WireGuard traffic. The default port is `51820` if left empty.

Note: If the **Endpoint** is empty, this value is ignored.

Keep Alive
An interval, in seconds, at which an empty packet is sent to the peer to keep the session active. This can improve handling through stateful firewalls. Disabled by default.

Public Key
The public key of this peer.

Pre-Shared Key
An optional pre-shared key which provides an additional layer of symmetric-key cryptography on top of the public key cryptography for post-quantum resistance.

A new pre-shared key can be generated by the  **Generate** button.

Tip: Click **Copy** under the public key to copy it to the clipboard.

Allowed IPs
List of networks **on the peer side** which the firewall can reach through this peer. For example, on a site-to-site VPN this would be the tunnel address of the peer and any LAN segments reachable via this peer.

When a tunnel has multiple peers this list allows WireGuard to determine which peer will receive traffic for destinations routed through the WireGuard interface.

The networks listed here are transformed into proper subnet start boundaries prior to validating and saving.

> **Warning:** These networks cannot be duplicated between multiple peers on the same tunnel, they must be unique. Otherwise, only the last peer in the list will be configured properly.

Note: All traffic may be associated with a peer by using `0.0.0.0/0` for IPv4 or `::/0` for IPv6, but this won't work for a tunnel with multiple peers. Only the last peer in the list will be configured properly.

Note: Routes are not automatically created in the system routing table. Routes for networks other than the tunnel network itself must be configured separately using *static* or *dynamic* routes.

Tip: For those familiar with OpenVPN, the internal routing used by WireGuard is similar to `iroute`

statements which associate remote networks with specific clients.

20.9.2 Design Considerations

One of the main considerations when choosing a WireGuard implementation layout is whether to use one tunnel with many peers, or one tunnel per peer.

Routing to WireGuard Peers

WireGuard uses what it calls "Cryptokey Routing" to map traffic inside WireGuard to a specific peer which is then encrypted using the public key for that peer. In practice, this means that when multiple peers are defined on a WireGuard instance, it must have all networks which will be routed to each peer defined on the peer. This can make managing networks and routes cumbersome.

When there is only one peer on a WireGuard interface, it can instead assume that the one peer is the correct destination for all traffic which crosses the interface (e.g. **Allowed IPs** set to `0.0.0.0/0` or `::/0`). And in that case, a routing protocol such as BGP or OSPF can manage the operating system routing to the neighbor instead of static routes.

Design Style

WireGuard does not have a concept of "Client" and "Server" per se, but depending on the configuration the firewall can behave in a manner similar to a "Client" (initiates locally, remote never initiates) or "server" (never initiates, remotes always initiate).

Technically every WireGuard tunnel is a peer to peer connection, but there are three main ways a WireGuard tunnel can be configured depending on whether a peer endpoint is known or defined:

- Site-to-Site (peer endpoint filled in on both sides)
- Remote Access "Server" (endpoint only filled in on remote peers)
- Remote Access "Client" (endpoint only filled in locally, not on the "server" peer)

Any of those roles can technically be configured no matter how the peer endpoint settings are defined, but not defining an endpoint on one side or the other limits the capacity in which a peer can operate.

In the case of remote access style setups, the peer endpoint address is typically unknown and can change at any time. In this case, the peer endpoint can be left blank and WireGuard will accept connections from any remote address, validating the key instead.

Note: WireGuard handles roaming automatically, and can detect when a peer changes IP addresses. WireGuard will recognize that authenticated data is coming from a new address and update itself accordingly.

20.9.3 WireGuard Limitations

There are a few limitations of the WireGuard implementation in FreeBSD which must be taken into consideration when deciding if WireGuard fits the needs of a use case.

Response Sourcing

WireGuard does not bind itself to an interface or a specific address on the firewall, but instead can accept traffic on any local IP address. This makes it very flexible, but can cause problems with functionality which requires traffic to use a specific address. When a WireGuard peer contacts the firewall, the firewall will respond from the address the peer used to contact it, if possible.

There are certain cases where this may not function as expected, such as when a peer was communicating with a CARP VIP which changed status.

Remote Peer Endpoint Requirements

When this firewall initiates a packet to a remote WireGuard endpoint, the source will be based on the main IP address on the WAN interface with the default gateway. If a remote endpoint expects to communicate with a different address, it is unlikely to work.

In this case, the problem can be worked around by updating the remote peer to use the correct endpoint address.

If this firewall is the only peer to initiate traffic in the tunnel, an outbound NAT rule could be crafted which matches the incorrect source address and translates it to the expected address.

Another alternative is to remove the remote endpoint from the peer configuration on this firewall, and set the remote endpoint to use the expected peer address for this firewall. If the remote peer always initiates traffic and sends to the expected address, replies will use that address.

High Availability

The address selection behavior mentioned earlier can cause some unexpected behavior with CARP.

For example, if a CARP node goes into maintenance mode during ongoing communication with a peer, the CARP VIP will no longer be available for use by WireGuard and the CARP node may respond to the client using its interface address instead. This can cause the remote peer to continue attempting communication with that specific node, and not the node holding MASTER status for the CARP VIP. This isn't a problem when a CARP node is powered off or reboots, as in those cases the CARP node cannot issue the unexpected responses.

When initiating traffic to a remote endpoint it can't be properly sourced from a CARP VIP. Outbound NAT cannot work around this limitation as the state created by outbound NAT is tied to a specific HA node. If that node fails, the other node cannot utilize the same state as it doesn't match its own local address. Allowing the remote peer to always initiate can work around this limitation.

Multi-WAN

If a client contacts the firewall via its WAN2 address, the firewall will respond from its WAN2 address as expected. However, if this firewall initiates, the traffic will always leave via the interface with the default route unless the routing table sends the traffic by another path.

If the only concern is connectivity, this may be acceptable. However, it is not acceptable if the intent is to cause WireGuard to use WAN2 for its initiated traffic.

For remote peers with static addresses, this can be worked around by adding a static route for the remote endpoint address using the correct WAN gateway.

20.9.4 Configure a WireGuard Tunnel

To configure a WireGuard Tunnel:

- Navigate to **VPN > WireGuard > Tunnels**
- Click  **Add Tunnel**
- Fill in the WireGuard Tunnel settings as described in *WireGuard Package Settings*
- Click **Save Tunnel**
- Add firewall rules on **Firewall > Rules**, **WAN** tab to allow UDP traffic to the port for this WireGuard tunnel (*WireGuard and Rules / NAT*)
- Add firewall rules on the common **Firewall > Rules**, **WireGuard** tab to pass traffic inside the VPN (*WireGuard and Rules / NAT*)

Configure a WireGuard Peer

To configure a WireGuard peer:

- Navigate to **VPN > WireGuard > Peers**
- Click  **Add Peer**
- Fill in the WireGuard Peer settings as described in *WireGuard Peer Settings*
- Click **Save Peer**
- Repeat the add/configure steps if there are multiple peers

Additional Configuration Steps

After configuring the WireGuard tunnel, there are a few more optional steps depending on the requirements of the use case:

- Navigate to **System > Routing**
- Set the **Default gateway** options to a specific gateway or group, as long as they are not left at *Automatic* (*Managing the Default Gateway*)

> **Warning:** If the default gateway remains set to *Automatic* the firewall may end up using the WireGuard interface as the default gateway, which is unlikely to be the desired outcome.

- Assign the WireGuard interface as a new OPTx interface (*Assign a WireGuard Interface*)
- Add firewall rules specific to this tunnel on **Firewall > Rules**, **OPTx** tab to pass traffic inside the VPN (*WireGuard and Rules / NAT*)
- Setup one of the alternate routing methods as described in *WireGuard Routing*, if needed.

20.9.5 Assign a WireGuard Interface

Some functionality for WireGuard interfaces depends upon them being assigned as their own interfaces on the firewall. Benefits of assignment include:

- Adds a firewall tab under **Firewall > Rules**
- Allows the interface to be selected for use with NAT rules
- Allows the interface to be selected throughout the GUI and packages for various purposes
- Rules on assigned interface tabs get `reply-to` which ensures return routing will exit back the expected interface for inbound connections.

Assignment Procedure

To assign the interface:

- Navigate to **System > Routing**
- Set the **Default gateway** options to a specific gateway or group, as long as they are not left at *Automatic* (*Managing the Default Gateway*)

> **Warning:** If the default gateway remains set to *Automatic* the firewall may end up using the WireGuard interface as the default gateway, which is unlikely to be the desired outcome.

- Navigate to **Interfaces > Assignments**
- Select the appropriate `tun_wg<number>` interface in the **Available network ports** list

  The description of the tunnel is printed next to the interface name in the list.

- Click  **Add** to assign the interface as a new OPT interface (e.g. OPT1)
- Navigate to the Interface configuration page, **Interfaces > OPTx**
- Check **Enable**
- Enter an appropriate **Description** which will become the interface name (e.g. `WG_S2S`)
- Configure an appropriate **MTU** value for the WireGuard interface

  The appropriate *MTU* varies depending on the MTU of the underlying circuit. WireGuard overhead is approximately `80` Bytes for IPv6 packets and `60` Bytes for IPv4 packets.

  On WANs with `1500` byte MTUs, the MTU for WireGuard interfaces should be `1420` for VPNs carrying IPv6 packets, or `1440` for VPNs which only carry IPv4 traffic.

Other WAN types with smaller MTUs, such as PPPoE, should subtract the overhead from their actual WAN MTU. When in doubt, use a slightly lower value to avoid excess fragmentation.

- Configure interface addresses and gateways as necessary
- Click **Save**
- Click **Apply Changes**

20.9.6 WireGuard and Rules / NAT

There are multiple concerns with firewall rules for WireGuard.

External Traffic

Firewall rules must pass traffic on WAN to the WireGuard **Listen Port** for a tunnel if remote WireGuard peers will initiate connections to this firewall. The protocol is always UDP, and the default port is `51820`.

Tunneled Traffic

Firewall rules must pass traffic on WireGuard interfaces to allow traffic inside the VPN, assuming remote connections should be allowed to local internal hosts. Use rules on the WireGuard group tab or rule tabs for assigned interfaces.

Rules on the WireGuard group tab are considered first and can match traffic on any WireGuard interfaces whether they are assigned.

Assigned WireGuard interfaces get their own individual rule tabs and will only match traffic on that specific tunnel interface. Rules on assigned WireGuard interface tabs also get `reply-to` which ensures that traffic entering a specific assigned WireGuard interface exits back out the same interface. Without that, return traffic will follow the default gateway.

Warning: Rules on the WireGuard group tab are matched first, so ensure rules on the group tab are removed, disabled, or do not match traffic which requires `reply-to`.

NAT functions on WireGuard interfaces once assigned. Outbound NAT, 1:1 NAT, and port forwards all work as expected.

Note: The firewall will automatically perform Outbound NAT on traffic exiting assigned WireGuard interfaces when using the default **Automatic Outbound NAT** mode (*See Outbound NAT*).

20.9.7 WireGuard Routing

WireGuard can work with both static and dynamic routing, depending on the environment.

Static Routing

WireGuard routing can be handled manually to reach remote LAN segments in addition to the tunnel network itself. To set up static routes:

Warning: Before assigning the interface, make sure default gateway for the firewall is not set to *Automatic* or the firewall may end up using the `tun_wg<num>` interface as the default gateway, which is unlikely to be the desired outcome.

- Assign the WireGuard interface for this peer (*Assign a WireGuard Interface*)
- Create a gateway using the peer address (*Gateways*)
- Add static routes using this gateway for the WireGuard tunnel

Dynamic Routing

WireGuard can work with dynamic routing, but there are some special considerations to take into account.

Note: This has only been tested with FRR.

The primary requirement to use dynamic routing with WireGuard is that there can only be one peer per WireGuard tunnel. When more than one peer is connected to a single WireGuard tunnel, WireGuard requires **Allowed IPs** to decide where to send specific networks. In that case, having to define these networks manually negates the purpose of dynamic routing. Using a single peer allows WireGuard to send any traffic it needs across the interface, including arbitrary networks.

BGP
: BGP works without any special configuration. Define the neighbor using the WireGuard interface address of the peer.

OSPF
: OSPF works, but needs special settings because it cannot utilize multicast traffic to find neighbors.

  In the OSPF settings of FRR:

  - Set the WireGuard interface **Network Type** to *Non-Broadcast* mode
  - Add a manual entry on the **Neighbors** tab using the WireGuard interface address of the peer

Other routing protocols have not been tested. If a routing protocol relies on broadcast or multicast traffic, it is unlikely to work.

Return Routing

When allowing inbound connections from arbitrary remote networks, use rules only on assigned WireGuard interface tabs only to ensure proper return routing.

Assigned WireGuard interfaces get their own individual rule tabs and will only match traffic on that specific tunnel interface. Rules on assigned WireGuard interface tabs also get `reply-to` which ensures that traffic entering a specific assigned WireGuard interface exits back out the same interface. Without that, return traffic will follow the default gateway.

See also:

- *WireGuard Remote Access VPN Configuration Example*

- *WireGuard Site-to-Site VPN Configuration Example*
- *WireGuard VPN Client Configuration Example*

20.9.8 WireGuard Overview

WireGuard is a new VPN Layer 3 protocol designed for speed and simplicity. It performs nearly as fast as hardware-accelerated IPsec and has only a few options in its configuration.

Due to this simplicity, WireGuard lacks many of the conveniences of more complicated VPN types which can help automate large deployments. Thus, while its performance scales well, the management can become cumbersome for large numbers of peers.

WireGuard behaves unlike other traditional VPN types in several ways:

- It operates completely in the kernel
- Configuration is placed directly on the interfaces
- It has no concept of connections or sessions
- It has no facilities for user authentication
- There is minimal logging from the kernel
- It does not bind to a specific interface or address on the firewall, it accepts traffic to any address on the firewall on its specified port

WireGuard instances consist of a tunnel and one or more peer definitions which contain the necessary keys and other configuration data.

WireGuard interfaces carry Layer 3 information and above.

To use WireGuard, upgrade to the latest version of pfSense Plus or pfSense CE software then install the WireGuard package from the *Package Manager*.

See also:

- *L2TP Server Configuration*
- *Troubleshooting Cisco VPN Pass Through*

VPNs provide a means of tunneling traffic through an encrypted connection, preventing it from being seen or modified in transit. pfSense® software offers several VPN options: IPsec, OpenVPN, WireGuard and L2TP. This section provides an overview of VPN usage, the pros and cons of each type of VPN, and how to decide which is the best fit for a particular environment. Subsequent sections discuss each VPN option in detail.

L2TP is purely a tunneling protocol and does not offer any encryption of its own. It is typically combined with another method of encryption such as IPsec in transport mode. Because of this, it doesn't fit in with most of the discussion in this chapter. See *L2TP VPN* for more information on L2TP.

20.10 PPTP Warning

pfSense software does not include a PPTP server. Despite the attraction of its convenience, PPTP **must not be used** under any circumstances because it is no longer secure. This is not specific to the implementation of PPTP that was in pfSense software; Any device that utilizes PPTP is no longer secure.

PPTP relies upon MS-CHAPv2 which has been completely compromised. Intercepted traffic can be decrypted by a third party 100% of the time, so consider any traffic carried in PPTP unencrypted. Migrate to another VPN type as soon as possible. More information on the PPTP security compromise can be found at https://isc.sans.edu/diary/End+of+Days+for+MS-CHAPv2/13807 and https://www.cloudcracker.com/blog/2012/07/29/cracking-ms-chap-v2/.

CHAPTER

TWENTYONE

L2TP VPN

21.1 L2TP and Firewall Rules

By default, when the L2TP server is enabled, firewall rules **will not** be automatically added to the chosen interface to permit UDP port `1701`. A firewall rule must be added to whichever interface the L2TP traffic will be entering, typically WAN, the WAN containing the default gateway, or IPsec.

21.2 L2TP and Multi-WAN

L2TP uses UDP port `1701`. Because L2TP relies on UDP, the server may have issues using any WAN that is not the default gateway. The daemon will respond from the firewall using the closest address to the client, following the routing table, which is the WAN with the default gateway for remote clients.

21.3 L2TP Server Configuration

To use L2TP, first browse to **VPN > L2TP**. Select **Enable L2TP server**.

> **Warning:** L2TP is **not** a secure protocol by itself; it only provides tunneling, **it does not perform encryption**.

See also:

L2TP/IPsec is a way to secure L2TP traffic by sending it through an encrypted IPsec tunnel. This may be used in combination with a mobile IPsec setup to configure L2TP+IPsec; see *L2TP/IPsec Remote Access VPN Configuration Example* for details.

21.3.1 Interface

The **Interface** setting controls where the L2TP daemon will bind and listen for connections. This is typically the *WAN* interface accepting inbound connections.

21.3.2 IP Addressing

Before starting, determine which IP addresses to use for the L2TP server and clients and now many concurrent clients to allow.

Server Address
An *unused* IP address outside the **Remote Address Range**, such as `10.3.177.1` as shown in Figure *L2TP IP Addressing*.

Remote Address Range
The addresses to be assigned to clients when they connect. Usually a new and unused subnet, such as `10.3.177.128/25` (`.128` through `.255`).

Number of L2TP users
Controls how many L2TP users will be allowed to connect at the same time, in this example `13` has been selected.

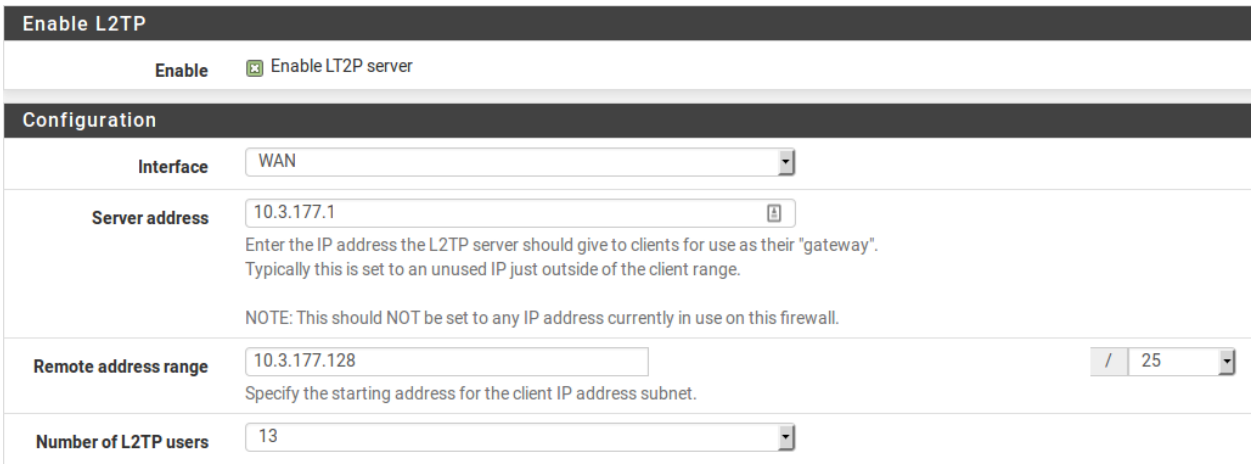


Fig. 1: L2TP IP Addressing

DNS servers can also be defined for end users when needed. Fill in the **Primary** and **Secondary L2TP DNS server** fields with the DNS server IP addresses for connecting clients.

21.3.3 Authentication

Secret
Required by some L2TP implementations, similar to a group password or pre-shared key. Compatibility varies from client to client. Leave the field blank unless it is known to be required. If required, enter and confirm the secret.

Authentication Type
Decides between *PAP*, *CHAP*, or *MS-CHAPv2* authentication for users. Compatibility varies from client to client, and it may also depend on the RADIUS server. The *CHAP* based types are more secure, but *PAP* is more widely compatible.

Users may be authenticated from the local user database, or via an external RADIUS server. This can be used to authenticate L2TP users from Microsoft Active Directory (see *Authenticating from Active Directory using RADIUS/NPS*) as well as numerous other RADIUS capable servers.

If using RADIUS, check the **Use a RADIUS server for authentication** box and fill in the RADIUS server and shared secret. For authentication using the local user database, leave that box unchecked. Users must be added manually on the

Users tab of the **VPN > L2TP** screen unless using RADIUS. See *Adding Users* below for more details on the built-in authentication system.

21.3.4 Save changes to start L2TP server

After filling in the aforementioned items, click **Save**. This will save the configuration and launch the L2TP server.

21.3.5 Configure firewall rules for L2TP clients

Browse to **Firewall > Rules** and click the **L2TP VPN** tab. These rules control traffic from L2TP clients. Until a firewall rule has been added to allow traffic, all traffic initiated from connected L2TP clients will be blocked. Traffic initiated from the LAN to L2TP clients is controlled using LAN firewall rules. Initially an allow all rule may be desired here for testing purposes as shown in Figure *L2TP VPN Firewall Rule*, and once functionality has been verified, restrict the ruleset as desired.

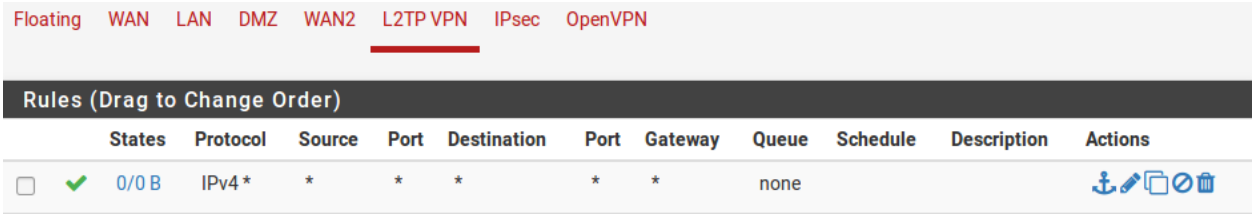


Fig. 2: L2TP VPN Firewall Rule

Note: Remember that a rule must also be added to the interface receiving the L2TP traffic, typically WAN or IPsec, to pass UDP to the firewall with a destination port of `1701`.

21.3.6 Adding Users

Adding users to the built-in L2TP users system is simple. To add local users:

- Navigate to **VPN > L2TP**, **Users** tab. The users screen as shown in Figure *L2TP Users Tab* will be presented.
- Click  **Add** to show the form used to add users.



Fig. 3: L2TP Users Tab

- Enter the **Username**, **Password** and **Confirm Password** for a user, as in Figure *Adding a L2TP User*.
- Enter a static **IP assignment** if desired.
- Click **Save**, and then the user list will return.

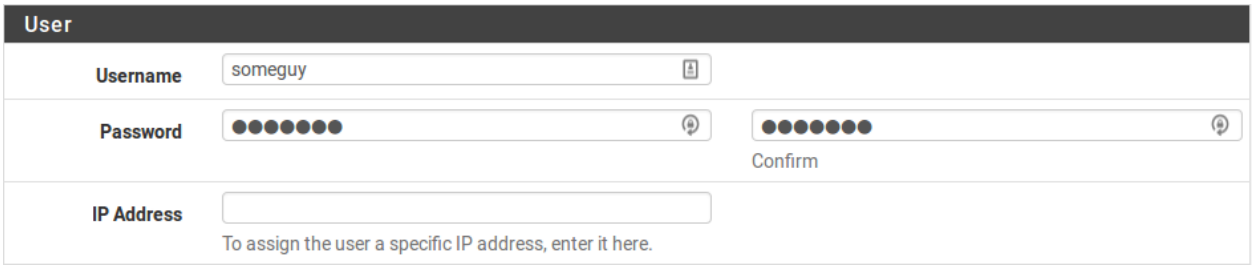


Fig. 4: Adding a L2TP User

- Repeat the process for each user to add.

To edit an existing user, click . Users may be deleted by clicking .

See also:

- *L2TP/IPsec Remote Access VPN Configuration Example*
- *Troubleshooting L2TP*
- *L2TP Logs*

pfSense® software can act as an L2TP VPN server. L2TP is purely a tunneling protocol that offers no encryption of its own, so it is typically combined with some other encryption technique, such as IPsec.

Warning: While L2TP over IPsec can be configured on pfSense software, it has severe limitations and problems compared to other types of remote access VPNs, and it should be avoided unless absolutely necessary. Current best practices include using IKEv2 IPsec, OpenVPN, or WireGuard for remote access VPNs.

Most L2TP/IPsec clients **will not** work properly in common scenarios. The most common problem scenario is Windows clients behind NAT, which is nearly all Windows clients in practice. The Windows L2TP/IPsec client and the strongSwan IPsec daemon are not fully compatible when the client is behind a NAT device, which leads to failure. In the few situations where L2TP/IPsec can function properly it still suffers from security and performance concerns compared to other types of remote access VPNs.

See also:

IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2 contains a walkthrough for configuring IKEv2, which is a much more flexible solution.

There are also recipes for other types of remote access VPNs in *pfSense® software Configuration Recipes*.

See also:

For general discussion of the various types of VPN implementations available in pfSense software and their pros and cons, see *Virtual Private Networks*.

21.4 L2TP Security Warning

L2TP on its own is not encrypted, so it is not intended for private traffic. Some devices, such as Android, offer an L2TP-only client which is capable of connecting back to pfSense software, but it should only be used for traffic that is already encrypted, or if the traffic is not considered private. For example, tunneling Internet traffic so it appears to originate from another location.

CHAPTER

TWENTYTWO

SERVICES

22.1 DHCPv4 Server

The DHCPv4 server in pfSense® software allocates addresses to IPv4 DHCP clients and automatically configures them for network access. By default, the DHCPv4 server is enabled on the LAN interface and configured to serve addresses in the LAN subnet (e.g. `192.168.1.0/24`).

To alter the behavior of the IPv4 DHCP server, navigate to **Services > DHCP Server** in the web interface. This page contains a tab for each interface capable of offering DHCPv4 service. The behavior of the IPv4 DHCP server for an interface is controlled on each tab, along with static IP address mappings and related options.

Warning: The DHCPv4 server cannot be active on any interface if the *DHCPv4 Relay* service is in use.

22.1.1 Settings Tab

When using the Kea DHCP backend there is a **Settings** tab with global options to control DHCP server behavior not specific to a given interface. The options on the **Settings** tab are covered in *Kea Settings Tab*.

22.1.2 Choosing an Interface

The DHCP configuration page contains a tab for each interface with a static IP address. Each interface has its own separate DHCP server configuration, and they may be enabled or disabled independently of one another. Before making any changes, visit the tab for the correct interface.

Note: The settings available on the page vary depending on the active DHCP backend (Kea or ISC DHCP). Differences are noted where applicable.

22.1.3 General Options

DHCP Backend
This read-only field displays the current DHCP backend, either Kea DHCP or ISC DHCP.

The backend can be changed under **System > Advanced**, **Networking** tab (*Server Backend*).

Enable
The first setting on the tab enables or disables DHCP service for the interface. To turn on DHCP for the interface, check **Enable DHCP server on [name] interface**. To disable the service, uncheck the box instead.

Deny unknown clients
Controls how the DHCP server handles requests from clients which it does not know.

Note: This option is per-pool, meaning that if unknown clients are denied in the default range, another pool of IP addresses may be defined that allows clients instead.

Can be set to one of the following values:

Allow All Clients
This is the default behavior. The DHCPv6 server will answer requests from any client requesting a lease. In most environments this is normal and acceptable behavior, but in restricted or secure environments this behavior is undesirable.

Allow known clients from any interface
With this option set, clients with static mappings defined on any interface will receive leases from this pool. This is a more secure practice but requires much more management overhead.

Allow known clients from only this interface
With this option set, clients with static mappings defined on this interface will receive leases from this pool. This is a more secure practice but requires much more management overhead.

Note: This will protect against low-knowledge users and people who casually plug in devices. Be aware, however, that a user with knowledge of the network could hardcode an IP address, subnet mask, gateway, and *DNS* which will still give them access. They could also alter/spoof their MAC address to match a valid client and still obtain a lease. Where possible, couple this setting with static ARP entries, access control in a switch that will limit MAC addresses to certain switch ports for increased security, and turn off or disable unused switch ports.

Ignore Denied Clients (ISC Only)
When checked, the ISC DHCP daemon will ignore denied clients rather than responding with a rejection message.

Note: This option is not compatible with high availability failover.

Ignore Client Identifiers
When set, the DHCP server will not record a unique identifier (UID) in client lease data if present in the client DHCP request.

This option may be useful when a client can dual boot using different client identifiers but the same hardware (MAC) address.

Note: This server behavior violates the official DHCP specification.

DNS Registration
Controls the *DNS Registration* behavior of this interface.

Track Server
Follows the default behavior for DNS Registration configured on the **Settings** tab.

Enable
Ignores the default setting and enables DNS Registration for DHCP clients on this interface.

Disable
Ignores the default setting and disables DNS Registration for DHCP clients on this interface.

Early DNS Registration
Controls the *Early DNS Registration* behavior of this interface.

Track Server
Follows the default behavior for Early DNS Registration configured on the **Settings** tab.

Enable
Ignores the default setting and enables Early DNS Registration for DHCP clients on this interface.

Disable
Ignores the default setting and disables Early DNS Registration for DHCP clients on this interface.

22.1.4 Primary Address Pool

Subnet
A read-only field with the current subnet on this interface.

Subnet Range
The range of available addresses inside the interface subnet, for reference and to help determine the desired range for DHCP clients. The network address and broadcast address are excluded, but interface addresses and Virtual IP addresses are not excluded.

Address Pool Range
This defines the DHCP address range, also referred to as the Scope or Pool. This range can be as large or as small as the network needs, but it must be wholly contained within the subnet.

Addresses between the entered values, inclusive, will be used for clients which request addresses via DHCP.

From
The starting address of the pool.

Must be lower than the **To** address.

To
The ending address of the pool.

Must be higher than the **From** address.

Note: The default LAN DHCP range is based off of the subnet for the default LAN IP address. It is `192.168.1.100` to `192.168.1.199`.

Additional Pools

The **Additional Pools** section defines extra pools of addresses inside the same subnet. These pools can be used to craft sets of IP addresses specifically for certain clients, or for overflow from a smaller original pool, or to split up the main pool into smaller chunks with a GAP of non-DHCP IP addresses in the middle of what used to be the pool. A combination of the MAC Address Control options may be used to guide clients from the same manufacturer into a specific pool, such as VoIP phones.

To add a new pool, click + **Add Address Pool** and the screen will switch to the pool editing view, which is nearly the same as the normal DHCP options, except a few options that are not currently possible in pools are omitted. The options behave the same as the others discussed in this section. Items left blank will, by default, fall through and use the options from the main DHCP range.

Note: See the MAC Address Control section below for specifics on directing clients into or away from pools.

22.1.5 Server Options

WINS Servers
Defines up to two WINS Servers (Windows Internet Name Service) which the server provides to clients.

Note: WINS is deprecated, but still active on some legacy networks.

DNS Servers
Defines up to four DNS server IP addresses which the server provides to clients. To use custom DNS Servers instead of automatic choices, fill in the DNS server IP addresses.

Tip: When using the DNS Resolver or DNS forwarder in combination with high availability clustering, specify a CARP Virtual IP address on this interface as the only DNS server.

When left empty, the firewall will automatically determine which addresses to supply to clients depending on the DNS configuration on this firewall:

- If the firewall is using the built-in DNS Resolver or DNS Forwarder to handle DNS, leave these fields blank. It will automatically assign itself as the DNS server for client devices.
- If the DNS Resolver or Forwarder is disabled and these fields are left blank, the firewall will pass on whichever DNS servers are defined under **System > General Setup**.

Tip: In networks with Windows servers, especially those employing Active Directory, the best practice is to use those servers for client DNS.

22.1.6 OMAPI (ISC Only)

Options for OMAPI service provided by the DHCP server, which allows querying and controlling the DHCP server remotely.

OMAPI Port
Set the port that OMAPI will listen on. The default port is `7911`, leave blank to disable.

Only the first OMAPI configuration is used, configurations on other interface tabs are ignored.

OMAPI Key
Enter a key matching the selected **Key Algorithm** to secure connections to the OMAPI endpoint.

Generate New Key
When checked, generates a new key based on the selected algorithm when the settings are saved.

Key Algorithm
The algorithm used to generate the OMAPI key.

22.1.7 Other DHCP Options

Gateway
This may also be left blank if this firewall is acting as the gateway for the network on this interface. If that is not the case, fill in the IP address for the gateway to be used by clients on this interface. When using CARP, fill in the CARP Virtual IP address on this interface here.

Domain Name
Specifies the domain name passed to the client to form its fully qualified hostname. If the **Domain Name** is left blank, then the domain name of the firewall it sent to the client. Otherwise, the client is sent this value.

Domain Search List
Controls the DNS search domains that are provided to the client via DHCP. If multiple domains are present and short hostnames are desired, provide a list of domain names here, separated by a semicolon. Clients will attempt to resolve hostnames by adding the domains, in turn, from this list before trying to find them externally. If left blank, the Domain Name option is used.

Note: The **Domain Search List** is provided via DHCP option `119`. Compatibility with this option varies by Operating System and version. See *Using DHCP Search Domains on Windows DHCP Clients*.

Default lease time
Controls how long a lease will last when a client does not request a specific lease length. Specified in seconds, default value is `7200` seconds (2 hours)

Maximum lease time
Limits a requested lease length to a stated maximum amount of time. Specified in seconds, default value is 86400 seconds (1 day).

Failover Peer IP (ISC Only)
If this firewall is part of a High Availability failover cluster, enter the real IP address of the other node in this subnet here.

Do not enter a CARP Virtual IP address.

Note: When **Failover Peer IP** is configured in a High Availability setup, the failover node should be available when the service is started to allow lease pool information to be synchronized; failing this, the DHCPD service will not respond to DHCPDISOVER requests.

Static ARP (ISC Only)
This checkbox works similar to denying unknown MAC addresses from obtaining leases, but takes it a step further in that it also restricts any unknown MAC address from communicating with this firewall. This stops would-be abusers from hardcoding an unused address on this subnet, circumventing DHCP restrictions.

Note: When using static ARP, all systems that need to communicate with the firewall must be listed in static mappings before activating this option, especially the system being used to connect to the firewall GUI. Also be aware that this option may prevent people from hardcoding an IP address and talking to the firewall, but it does not prevent them from reaching each other on the local network segment.

Time Format Change (ISC Only)
By default, the ISC DHCP daemon maintains lease times in UTC. When this option is checked, the times on the DHCP Leases status page are converted to the local time zone defined on the firewall.

Statistics Graphs (ISC Only)
This option, disabled by default, activates RRD graphing for monitoring the DHCP pool utilization.

Ping Check (ISC Only)
When checked, the DHCP server will attempt to ping a client address before allocation to ensure it is not in use by another device.

22.1.8 Dynamic DNS (ISC Only)

For Dynamic DNS settings, click **Display Advanced** to the right of that field, which displays the following options:

DHCP Registration
Check the box to enable registration of DHCP client names in DNS using an **external** DNS server (not on the firewall).

DDNS Domain
The domain name used for registering clients in DNS

DDNS Hostnames
When set, forces the dynamic DNS hostname to match the hostname on a static mapping instead of taking the name given by the client.

Primary DDNS Address
The DNS server used for registering clients in DNS

Secondary DDNS Address
The secondary DNS server used for registering clients in DNS

DNS Domain Key Name
The name of the encryption key used for DNS registration

Key Algorithm
The algorithm used to generate the **DDNS Domain Key Secret** value.

DDNS Domain Key Secret
The secret for the key used for DNS registration

DDNS Client Updates
How the DHCP server handles Forward entries when a client indicates it wishes to update DNS itself.

Allow
Prevents DHCP from updating Forward entries, allowing the client to make the update request itself.

Deny
Indicates that DHCP will do the updates and the client should not.

Ignore
Specifies that DHCP will do the update and the client can also attempt the update, usually using a different domain name.

DDNS Reverse
When set, attempts to add reverse DNS entries.

22.1.9 MAC Address Control

For MAC Address Control, click **Display Advanced** to show the lists of allowed and denied client MAC addresses. Each list is comma-separated and contains portions of MAC addresses. For example, a group of VoIP phones from the same manufacturer may all start with the MAC address `aa:bb:cc`. This can be leveraged to give groups of devices or users separate DHCP options.

Allow
A list of MAC Addresses to allow in this pool. If a MAC address is in this text box, then all others will be denied except the MAC address specified here.

Deny
A list of MAC Addresses to deny from this pool. If a MAC address is in this list, then all others are allowed.

The best practice is to use a combination of these lists to get the desired result, such as: In the main pool, leave **Allow** blank and set **Deny** to `aa:bb:cc`. Then in the VoIP pool, allow `aa:bb:cc`. If that extra step is not taken to allow the MAC prefix in the additional pool, then other non-VoIP phone clients could receive IP addresses from that pool, which may lead to undesired behavior.

This behavior may also be used to prevent certain devices from receiving a DHCP response. For example to prevent Example brand printers from receiving a DHCP address, if MAC addresses all start with `ee:ee:ee`, then place that in the deny list of each pool.

22.1.10 NTP Servers

To specify NTP Servers (Network Time Protocol Servers), click the **Display Advanced** button to the right of that field, and enter IP addresses for up to four NTP servers.

22.1.11 TFTP Server

Click the **Display Advanced** button next to **TFTP** to display the TFTP server option. The value in the TFTP Server box, if desired, must be an IP address or hostname of a TFTP server. This is most often used for VoIP phones, and may also be referred to as "option 66" in other documentation for VoIP and DHCP.

22.1.12 LDAP URI

Click the **Display Advanced** button next to **LDAP** to display the **LDAP Server URI** option. **LDAP Server URI** will send an LDAP server URI to the client if requested. This may also be referred to as DHCP option 95. It takes the form of a fully qualified LDAP URI, such as `ldap://ldap.example.com/dc=example,dc=com`. This option can help clients using certain kinds of systems, such as Open Directory, to find their server.

22.1.13 Network Booting

These options control how the DHCP server will direct clients to boot over the network (e.g. PXE).

Warning: Both a filename and a boot server must be configured for this to function properly. For UEFI & ARM to boot properly, all five filenames and a configured boot server are required.

Enable
Enables Network Booting options for DHCPv4.

Next Server
The IPv4 address from which boot images are available.

Default BIOS File Name
Filename to use when a client does not specify an architecture, such as for legacy BIOS booting.

UEFI 32 bit File Name
Filename to supply for 32-bit UEFI clients.

UEFI 64 bit File Name
Filename to supply for 64-bit UEFI clients.

ARM 32 bit File Name
Filename to supply for 32-bit ARM clients.

ARM 64 bit File Name
Filename to supply for 64-bit ARM clients.

UEFI HTTPBoot URL
URL to boot files for clients capable of booting using the HTTPBoot method. Must be in the format `http://<server-name>/<firmware-path>`.

Root Path
String to target a specific device as the client's root filesystem device, such as `iscsi:<server-name>:<protocol>:<port>:<LUN>:<target-name>`.

22.1.14 Custom Configuration (KEA Only)

The KEA GUI allows administrators to configure Kea directly with JSON-formatted configuration blocks. See *Custom Configuration* for details.

22.1.15 Additional BOOTP/DHCP Options (ISC Only)

Other numeric DHCP options can be sent to clients using the **Additional BOOTP/DHCP Options** controls. To view these options, click **Display Advanced** in this section. To add a new option, click  **Add Custom Option**.

Number
The DHCP option code number. IANA maintains a list of all valid DHCP options.

Type
The choices and formats for each type may be a little counter-intuitive, but the labels are used directly from the DHCP daemon.

The proper uses and formats are:

Text
Free-form text to be sent in reply, such as `http://www.example.com/wpad/wpad.dat` or `Example Company`.

String
A string of hexadecimal digits separated by a colon, such as `c0:a8:05:0c`.

Boolean
Either `true` or `false`.

Unsigned 8, 16, or 32-bit Integer
A positive Integer that will fit within the given data size, such as `86400`.

Signed 8, 16, or 32-bit Integer
A positive or negative Integer that will fit within the given data size, such as `-512`.

IP address or host
An IP address such as `192.168.1.1` or a hostname such as `www.example.com`.

Value
The value associated with this numeric option and type.

For more information on which options take a specific type or format, see the linked list above from the IANA.

Note: When using numbered custom options, be careful of the type. Some will be OK on text/string, but others are not.

For example, DHCP options for code `132` (and presumably `133`) for VLAN ID must be set for a type of unsigned integer `32`.

22.1.16 Save Settings

After making changes, click **Save** before attempting to create static mappings. Changes to settings will be lost if the browser leaves this page without saving.

22.1.17 Static Mappings

Static DHCP mappings express a preference for which IP address will be assigned to a given client based on its MAC address. In a network where unknown clients are denied, this also serves as a list of "known" clients which are allowed to receive leases or have static ARP entries. Static mappings can be added in one of two ways:

- From this screen, click  **Add Static Mapping**.
- From the DHCP leases view, click  on a lease row.

On this screen, only the **MAC address** is necessary.

MAC Address
The client MAC address which identifies a host. This can be used to deliver customized options on this page. Alternately, by entering only the MAC address it will be added to the list of known clients for use when the **Deny unknown clients** option is set.

Note: Client MAC address can be obtained from a command prompt on most platforms. On many UNIX-based or UNIX-work-alike operating systems including macOS, typing `ifconfig -a` will show the MAC address for each interface. On Linux, use `ip link`. On Windows, `ipconfig /all` will show the MAC address. The MAC address may also sometimes be found upon a sticker on the network card, or near the network jack for integrated adapters. For hosts on the same subnet, the MAC can be determined by pinging the IP address of the host and then running `arp -a`.

Client Identifier
An optional ID sent by the client to identify itself as per RFC 2132. This is used for matching, similar to the MAC address, it does not set a value for the client.

IP Address
The IP address field is needed if this will be a static IP address mapping instead of only informing the DHCP server that the client is valid.

This IP address is a **preference**, not a reservation. Assigning an IP address here will not prevent another host from using the same IP address. If the IP address is in use when this client requests a lease, the server will instead assign the client an address from the general pool. For this reason, the GUI does not allow assigning static mappings inside of pools.

See also:

Static Mappings Inside DHCP Pools

Hostname
The hostname of the client. This does not have to match the hostname set on the client. The hostname set here will be used when registering DHCP addresses in the DNS resolver.

Description
Cosmetic only, and available for use to help track any additional information about this entry. It could be the name of the person who uses the PC, its function, the reason it needed a static address, or the administrator who added the entry. It may also be left blank.

ARP Table Static Entry
If checked, this entry will receive a static ARP entry in the OS tying this IP address to this MAC address.

Note: If this option is used rather than using the global static ARP option, it does not prevent that MAC address from using other IP addresses, it only prevents other MAC addresses from using this IP address. In other words, it prevents another machine from using that IP address to reach the firewall, but it doesn't stop the user from changing their own IP address to something different.

The remaining options available to set for this client are the same in behavior to the ones found earlier in this section for the main DHCP settings.

Click **Save** to finish editing the static mapping and return to the DHCP Server configuration page.

22.2 DHCPv6 Server

The DHCPv6 server in pfSense® software allocates addresses to DHCPv6 clients and automatically configures them for network access. By default, the DHCPv6 server is enabled on the LAN interface and set to use a prefix obtained by tracking a DHCPv6 delegation from the WAN interface.

To alter the behavior of the IPv6 DHCP server, navigate to **Services > DHCPv6 Server** in the web interface. This page contains a tab for each interface capable of offering DHCPv6 service. The behavior of the IPv6 DHCP server for an interface is controlled on each tab, along with static IP address mappings and related options.

Warning: The DHCPv6 server cannot be active on any interface if the *DHCPv6 Relay* service is in use.

Note: For clients to query DHCPv6, *Router Advertisements* must also be enabled and set to either *Managed* or *Assisted* mode under **Services > Router Advertisements**.

22.2.1 Settings Tab

When using the Kea DHCP backend there is a **Settings** tab with global options to control DHCP server behavior not specific to a given interface. The options on the **Settings** tab are covered in *Kea Settings Tab*.

22.2.2 Choosing an Interface

The DHCPv6 daemon can run and be configured on interfaces with a Static IP address or interfaces which track delegated prefixes from upstream sources. If a tab for an interface is not present, check that it is enabled and configured either with a Static IPv6 address or to track a DHCPv6 type WAN which obtains a prefix delegation from an external source.

Note: The settings available on the page vary depending on the active DHCP backend (Kea or ISC DHCP). Differences are noted where applicable.

Note: DHCPv6 does not provide gateway information. *Router Advertisements* inform hosts on the network about available routers. DHCPv6 is for other host configuration such as DNS, delegation, and so on.

See also:

See the *DNS Forwarder* article for information on the default DNS server behavior.

22.2.3 General Options

DHCP Backend
This read-only field displays the current DHCP backend, either Kea DHCP or ISC DHCP.

The backend can be changed under **System > Advanced**, **Networking** tab (*Server Backend*).

Enable
The first setting on the tab enables or disables DHCPv6 service for the interface. To turn on DHCPv6 for the interface, check **Enable DHCPv6 server on [name] interface**. To disable the service, uncheck the box instead.

Deny unknown clients
Controls how the DHCP server handles requests from clients which it does not know.

Note: This option is per-pool, meaning that if unknown clients are denied in the default range, another pool of IP addresses may be defined that allows clients instead.

Can be set to one of the following values:

Allow All Clients
This is the default behavior. The DHCPv6 server will answer requests from any client requesting a lease. In most environments this is normal and acceptable behavior, but in restricted or secure environments this behavior is undesirable.

Allow known clients from any interface
With this option set, clients with static mappings defined on any interface will receive leases from this pool. This is a more secure practice but requires much more management overhead.

Allow known clients from only this interface
With this option set, clients with static mappings defined on this interface will receive leases from this pool. This is a more secure practice but requires much more management overhead.

Note: This will protect against low-knowledge users and people who casually plug in devices. Be aware, however, that a user with knowledge of the network could hardcode an IP address, subnet mask, gateway, and *DNS* which will still give them access. They could also alter/spoof their DUID to match a valid client and still obtain a lease. Where possible, couple this setting with access control in a switch that will limit access to switch ports for increased security, and turn off or disable unused switch ports.

DNS Registration
Controls the *DNS Registration* behavior of this interface.

Track Server
Follows the default behavior for DNS Registration configured on the **Settings** tab.

Enable
Ignores the default setting and enables DNS Registration for DHCP clients on this interface.

Disable
Ignores the default setting and disables DNS Registration for DHCP clients on this interface.

Early DNS Registration
Controls the *Early DNS Registration* behavior of this interface.

Track Server
Follows the default behavior for Early DNS Registration configured on the **Settings** tab.

Enable
Ignores the default setting and enables Early DNS Registration for DHCP clients on this interface.

Disable
Ignores the default setting and disables Early DNS Registration for DHCP clients on this interface.

22.2.4 Primary Address Pool

Prefix
A read-only field with the current prefix on this interface. If the prefix is delegated, this field also shows the WAN which supplied the tracking data and the subnet ID.

Prefix Range
A read-only field with the total available range of addresses in the prefix.

Address Pool Range
This field defines the start and end of an address range defining a pool from which the DHCP server will allocate addresses. This range can be as large or as small as the network needs, but it must be wholly contained within the subnet.

Addresses between the entered values, inclusive, will be used for clients which request addresses via DHCPv6.

From
The starting address of the pool. For interfaces which obtain a prefix dynamically, the prefix itself may be omitted. The value defined in this field will be added to the prefix automatically.

Must be lower than the **To** address.

To
The ending address of the pool. For interfaces which obtain a prefix dynamically, the prefix itself may be omitted. The value defined in this field will be added to the prefix automatically.

Must be higher than the **From** address.

Tip: Given the vast amount of space available inside even a /64, a good trick is to craft a range that restricts hosts to use an easy to remember or recognize range. For example, Inside a /64 such as `2001:db8:1:1::`, set the DHCPv6 range be: `2001:db8:1:1::d:0000` to

`2001:db8:1:1::d:FFFF`, using the `d` in the second to last section of the address as a sort of shorthand for "DHCP". That example range contains 2^16 (65,536) IPs, which is extremely large by today's IPv4 standards, but only a small portion of the whole `/64`.

Additional Pools

The **Additional Pools** section defines extra pools of addresses inside the same subnet. These pools can be used to craft sets of IPv6 addresses specifically for certain clients, or for overflow from a smaller original pool, or to split up the main pool into smaller chunks with a GAP of non-DHCPv6 IPv6 addresses in the middle of what used to be the pool.

To add a new pool, click  **Add Address Pool** and the screen will switch to the pool editing view, which is nearly the same as the normal DHCPv6 options, except a few options that are not currently possible in pools are omitted. The options behave the same as the others discussed in this section. Items left blank will, by default, fall through and use the options from the main DHCPv6 range.

22.2.5 DHCPv6 Prefix Delegation

Prefix delegation, covered earlier in *DHCP6 Prefix Delegation* and *Track Interface*, allows automatically dividing and allocating a block of IPv6 addresses to networks that will live behind other routers and firewalls which reside downstream from this firewall (e.g. in the LAN, DMZ, etc.). Downstream devices which request a delegation can in turn use prefixes in their delegation for their LAN, VPNs, DMZ, etc. Downstream firewalls can even further delegate their own allocation to routers behind them.

Note: Most users on networks act only in a client capacity and will not need this, so it will likely remain blank.

Prefix delegation can be used to allocate `/60` chunks of a `/48` to downstream routers automatically, or many other combinations. The downstream router obtains an IPv6 address then requests a delegation, the server delegates a prefix and dynamically routes the delegated prefix to that downstream router.

Delegated Prefix
Defines the delegation pool. This block of IPv6 addresses must be routed to this firewall by upstream routers either directly or as part of a larger allocation.

The length of this prefix must be smaller than, or equal to, the **Delegated Length**.

Warning: The deprecated ISC DHCPv6 server defined this value as a range instead of a prefix. As such, it may not be possible to set this value identically to previous configurations. However, it is much easier to specify as a prefix rather than having to calculate the start/end and check boundaries.

Delegated Length
Sets the size of the prefix delegations allocated to clients. This must be larger than, or equal to, the prefix length of the **Delegated Prefix**.

Note: In nearly all cases this should be smaller than `/64` and not larger than or equal to `/64`. A delegated prefix with a length larger than /64 could never be properly utilized by the majority of clients. While a single `/64` delegation is valid, it is of minimal use to downstream routers as it only allows room for a single proper internal network.

For example, start with a prefix of `2001:db8::ffff:f000/52` earmarked for delegation to downstream routers. This **Delegated Prefix** has a length of `/52`. Using a **Delegated length** value of `/60` accommodates up to **256** downstream router client delegations. Each of these delegations is a contiguous `/60` block of the **Delegated Prefix**, which is equivalent to **16** subnets of size `/64`.

22.2.6 Server Options

Enable
When set, the DHCP server provides DNS servers to DHCPv6 clients upon request.

> **Warning:** Unchecking this box disables the `dhcp6.name-servers` option. Use with caution as the resulting behavior may violate RFCs and lead to unintended client behavior.

DNS Servers
Defines up to four DNS server IPv6 addresses which the server provides to clients. To use custom DNS Servers instead of automatic choices, fill in the DNS server IPv6 addresses.

Tip: When using the DNS Resolver or DNS forwarder in combination with high availability clustering, specify an IPv6 CARP Virtual IP address on this interface as the only DNS server.

When left empty, the firewall will automatically determine which addresses to supply to clients depending on the DNS configuration on this firewall:

- If the firewall is using the built-in DNS Resolver or DNS Forwarder to handle DNS, leave these fields blank. It will automatically assign itself as the DNS server for client devices.
- If the DNS Resolver or Forwarder is disabled and these fields are left blank, the firewall will pass on whichever DNS servers are defined under **System > General Setup**.

Tip: In networks with Windows servers, especially those employing Active Directory, the best practice is to use those servers for client DNS.

22.2.7 Other DHCPv6 Options

Domain Name
Specifies the domain name passed to the client to form its fully qualified hostname. If the **Domain Name** is left blank, then the domain name of the firewall it sent to the client. Otherwise, the client is sent this value.

Domain Search List
Controls the DNS search domains that are provided to the client via DHCP. If multiple domains are present and short hostnames are desired, provide a list of domain names here, separated by a semicolon. Clients will attempt to resolve hostnames by adding the domains, in turn, from this list before trying to find them externally. If left blank, the Domain Name option is used.

Note: The **Domain Search List** is provided via DHCP option `119`. See *Using DHCP Search Domains on Windows DHCP Clients* for details.

Default lease time
Controls how long a lease will last when a client does not request a specific lease length. Specified in seconds, default value is `7200` seconds (2 hours)

Maximum lease time
Limits a requested lease length to a stated maximum amount of time. Specified in seconds, default value is 86400 seconds (1 day).

Time Format Change (ISC Only)
By default, the ISC DHCP daemon maintains lease times in UTC. When this option is checked, the times on the DHCP Leases status page are converted to the local time zone defined on the firewall.

22.2.8 Dynamic DNS (ISC Only)

For Dynamic DNS settings, click **Display Advanced** to the right of that field, which displays the following options:

DHCP Registration
Check the box to enable registration of DHCP client names in DNS using an **external** DNS server (not on the firewall).

DDNS Domain
The domain name used for registering clients in DNS

DDNS Hostnames
When set, forces the dynamic DNS hostname to match the hostname on a static mapping instead of taking the name given by the client.

Primary DDNS Address
The DNS server used for registering clients in DNS

Secondary DDNS Address
The secondary DNS server used for registering clients in DNS

DNS Domain Key Name
The name of the encryption key used for DNS registration

Key Algorithm
The algorithm used to generate the **DDNS Domain Key Secret** value.

DDNS Domain Key Secret
The secret for the key used for DNS registration

DDNS Client Updates
How the DHCP server handles Forward entries when a client indicates it wishes to update DNS itself.

Allow
Prevents DHCP from updating Forward entries, allowing the client to make the update request itself.

Deny
Indicates that DHCP will do the updates and the client should not.

Ignore
Specifies that DHCP will do the update and the client can also attempt the update, usually using a different domain name.

DDNS Reverse
When set, attempts to add reverse DNS entries.

22.2.9 NTP Servers

To specify NTP Servers (Network Time Protocol Servers), click the **Display Advanced** button to the right of that field, and enter IP addresses for up to four NTP servers.

22.2.10 Network Booting

Enable
Enables Network Booting options for DHCPv6

Boot File URL
URL containing boot files.

22.2.11 Custom Configuration (KEA Only)

The KEA GUI allows administrators to configure Kea directly with JSON-formatted configuration blocks. See *Custom Configuration* for details.

22.2.12 Additional BOOTP/DHCP Options (ISC Only)

Other numeric DHCP options can be sent to clients using the **Additional BOOTP/DHCP Options** controls. To view these options, click **Display Advanced** in this section. To add a new option, click  **Add Option**.

Number
The DHCP option code number. IANA maintains a list of all valid DHCP options.

Value
The value associated with this numeric option and type.

Warning: When using numbered custom options, be aware that numbered options do NOT correspond exactly to the DHCP numbered options for IPv4

For more information on DHCP option numbers and types, see https://tools.ietf.org/html/draft-ietf-dhc-v6opts-00

22.2.13 Save Settings

After making changes, click **Save** before attempting to create static mappings. Changes to settings will be lost if the browser leaves this page without saving.

22.2.14 DHCPv6 Static Mappings

Static mappings on DHCPv6 work differently than IPv4. On IPv4, mappings were matched and identified using the MAC address of a device. For IPv6, the designers decided that wasn't good enough, since the MAC address of a device could change, but still be the same device. Thus, the designers came up with the DHCP Unique Identifier (DUID) which in theory would be a unique ID per device without the same constraints as using MAC addresses.

DHCP Unique Identifier (DUID)

The **DHCP Unique Identifier**, or **DUID**. The DUID of the host is generated by the operating system of the client and, in theory, will remain unique to that specific host until such time as the user forces a new DUID or the operating system is reinstalled. The DUID can range from 12 to 20 bytes, and varies depending on its type.

This field expects a DUID for a client PC in a special format, represented by pairs of hexadecimal digits, separated by colons, such as `00:01:00:01:1b:a6:e7:ab:00:26:18:1a:86:21`.

How to obtain this DUID depends on the operating system. The easiest way is to allow the device to obtain a lease via DHCPv6, and then add an entry from the DHCPv6 Leases View (Status > DHCPv6 Leases). On Windows, it can be found as DHCPv6 Client DUID in the output of `ipconfig /all`.

Note: On Windows, the DUID is generated at install time, so if a base image is used and workstations are cloned from there, they can all end up with the same DUID, and thus all end up pulling the same IPv6 address over DHCPv6.

Clear the DUID from the registry before making an image to clone, by issuing the following command:

```
reg delete HKLM\SYSTEM\CurrentControlSet\Services\Tcpip6\Parameters /f /v Dhcpv6DUID
```

That command may also be run on a working system to reset its DUID if needed.

DUID Format

The DUID format is listed on the page, but it roughly follows the format:

```
DUID-LLT - ETH -- TIME --- ---- address ----
```

DUID-LLT is link-layer plus time, which means it uses the link type of network interface on the system (Generally `00:01` to indicate the format, plus `00:01`, or `00:06` for Ethernet), plus the timestamp at which the DUID was generated in hex, plus the MAC address of the first NIC. It may be difficult or impossible to predict the DUID for a device. Unless the operating system has a way to look it up, it may be best to allow the client to obtain a dynamic lease and then copy the DUID from the leases view.

The DUID may be entered in colon-separated or dash-separated format.

Static Mapping Settings

DHCPv6 Static Mappings have the following settings:

DHCP Unique Identifier
: The DUID, as discussed in *DHCP Unique Identifier (DUID)*, using the format specified in *DUID Format*.

IPv6 Address
: The IPv6 address field is needed if this will be a static IPv6 address mapping instead of only informing the DHCP server that the client is valid.

  This IPv6 address is a **preference**, not a reservation. Assigning an IPv6 address here will not prevent another host from using the same IPv6 address. If the IP address is in use when this client requests a lease, the server will instead assign the client an address from the general pool. For this reason, the GUI does not allow assigning static mappings inside of pools.

  See also:

  Static Mappings Inside DHCP Pools

Hostname
The hostname of the client. This does not have to match the hostname set on the client. The hostname set here will be used when registering DHCP addresses in the DNS resolver.

Description
Cosmetic only, and available for use to help track any additional information about this entry. It could be the name of the person who uses the PC, its function, the reason it needed a static address, or the administrator who added the entry. It may also be left blank.

22.3 IPv6 Router Advertisements

Automatic address and router assignment for IPv6 works much differently than IPv4. Even so, most of the DHCP options are similar, but there are notable differences in behavior in how things are assigned and also how items like the gateway are handed off to clients. Unless otherwise noted, options of the same name work the same for DHCP and DHCPv6. DHCPv6 is configured under **Services > DHCPv6 Server** and Router Advertisements (RA) are configured under **Services > Router Advertisement**.

22.3.1 DHCPv6 vs Stateless Address Autoconfiguration

There are a few clients that cannot use DHCPv6. Some clients can only obtain IPv6 addresses using Stateless Address Autoconfiguration (SLAAC). There is no way for the firewall to have direct knowledge of a list of hosts on the segment using SLAAC addresses, so for some environments it is much less desirable because of the lack of control and reporting of addresses. Consider address tracking and operating system requirements when deciding how to allocate IPv6 addresses to clients on the network.

Many operating systems such as Windows, macOS, FreeBSD, Linux, and their cousins contain DHCPv6 clients that are capable of obtaining addresses as expected via DHCPv6. Some lightweight or mobile operating systems such as Android do not contain a DHCPv6 client and will only function on a local segment with IPv6 using SLAAC.

22.3.2 Router Advertisements (Or: "Where is the DHCPv6 gateway option?")

In IPv6, hosts locate a router through Router Advertisement (RA) messages sent from routers instead of by DHCP; IPv6-enabled routers capable of performing dynamic address assignment are expected to announce themselves on the network to all clients. As such, DHCPv6 does not include any gateway information. So clients can obtain their addresses from DHCPv6 or SLAAC, but unless they are statically configured, they always locate their next hop by using RA packets sent from available gateways.

To enable the RA service:

- Navigate to **Services > Router Advertisement**
- Click the interface tab for the interface being configured
- Select a mode other than *Disabled* from the **Router Mode** drop-down list
- Click **Save**

The other options to control RA behavior may be set as needed for the network:

Router Mode
The modes for the RA daemon control the services offered by pfSense® software, announce the firewall as an IPv6 router on the network, and direct clients on how to obtain addresses.

Disabled
The RA daemon is disabled and will not run. IPv6 gateways must be entered manually on any client hosts.

Router Only
This firewall will send out RA packets that advertise itself as an IPv6 router. DHCPv6 and SLAAC are disabled in this mode.

Unmanaged
The firewall will send out RA packets and clients are directed to assign themselves IP addresses within the interface subnet using SLAAC. DHCPv6 is disabled in this mode.

Managed
The firewall will send out RA packets and addresses will **only** be assigned to clients using DHCPv6.

Assisted
The firewall will send out RA packets and addresses can be assigned to clients by DHCPv6 or SLAAC.

Stateless DHCP
The firewall will send out RA packets and addresses can be assigned to clients by SLAAC while providing additional information such as DNS and NTP from DHCPv6.

Router Priority
If multiple IPv6 routers exist on the same network segment, they can indicate to clients in which order they should be used. If a high priority router becomes unavailable, clients will try a normal priority router, and finally a low priority router. Select either *Low*, *Normal*, or *High* from the list. If there is only one router on the network, use *Normal*.

Note: Clients may lose a few packets or experience a delay when attempting to determine which router to use when relying on this type of failover between routers.

RA Interface
Chooses a specific address to advertise as the router instead of the link-local address on this interface.

This is used when using IPv6 in a High Availability cluster (*High Availability*). Advertising a link-local CARP VIP allows cluster members to direct clients to send traffic to whichever node is active, rather than trying to make each client determine the router on its own.

This option is only offered when the firewall has CARP VIPs configured.

Note: While it is possible to advertise a GUA IPv6 address, using a link-local address is more appropriate. Some client operating systems may not behave properly when given a non-link-local address for a router.

Valid Lifetime
Length of time, specified in seconds, that the advertised prefix will be valid. The default value is `86400` seconds (one day).

Preferred Lifetime
Length of time, specified in seconds, that the client addresses generated in this prefix using SLAAC are valid. The default value is `14400` seconds (four hours).

Minimum & Maximum RA interval
The router sends advertisements on each interface configured to transmit messages. Advertisements include route information and indicate to network hosts that the router is operational. The router sends these unsolicited multicast router advertisements periodically, with a time range defined by minimum and maximum values in seconds.

The maximum interval for sending RA messages should be less than or equal to the router lifetime in RA messages.

Router lifetime
The lifetime associated with the default router in seconds. A value of `0` indicates that the router is not a default router and that associated default routes should be discarded.

The default is three times the maximum RA interval seconds.

NAT64 Prefix
Enables PREF64 and configures a prefix for *NAT64*. This allows IPv6-only clients to automatically discover NAT64 prefixes which they can use to reach IPv4-only hosts through this firewall.

In most cases this value should be `64:ff9b::/96`. NAT64 must also be configured.

See also:

- *NAT64*
- *Configuring NAT64 for IPv6-only Clients*
- RFC 8781

NAT64 Prefix Lifetime
The length of time in seconds (relative to the time the packet is sent) that the prefix is valid for the purpose of NAT64 existence determination. The default is `3 * Maximum RA Interval` seconds.

RA Subnets
This section allows defining a list of subnets for which this firewall will send RA packets. Enter as many subnets as needed, each with an appropriate prefix (typically */64*.). To create an additional row for another subnet, click  **Add RA Subnet**.

DNS Configuration
Not all clients are capable of obtaining DNS information from RA messages. For clients that can, using SLAAC to give an IP address and DNS from RA can do away with the need for DHCPv6 entirely.

Enable
Enables providing DNS configuration using the RA daemon.

Mirror DHCPv6
Uses the DNS values configured in the DHCPv6 settings.

DNS Servers
Enter up to four IPv6 addresses for DNS Servers, or leave the fields blank to use the system default DNS servers or DNS Resolver/DNS forwarder if enabled.

Domain Search List
Operates identically to the DHCP option of the same name.

22.4 DHCPv4 & DHCPv6 Relay

DHCP requests are broadcast traffic. Broadcast traffic is limited to the broadcast domain where it is initiated. To provide DHCP service on a network segment without a DHCP server, use the DHCP relay to forward those requests to a defined server on another segment.

> **Warning:** It is not possible to run both a DHCP server and a DHCP Relay at the same time. To enable the DHCP relay, first disable the DHCP server on all interfaces.

22.4.1 DHCP Relay Options

The DHCP Relay service has the following options:

Enable DHCP Relay
Checked to enable the service.

Downstream Interfaces
Select one or more local interfaces containing clients for which the service will relay requests.

CARP Status VIP
A CARP type VIP which will be used to indicate whether the relay service will be started or stopped based on its status. For example, in an HA cluster, the service will only run when the chosen VIP is in CARP MASTER status, not when it is in BACKUP status. This ensures the HA nodes do not send conflicting relay messages.

Append circuit ID and agent ID to requests
Check this to add a circuit ID (interface number on the firewall) and the agent ID to the DHCP request. This may be required by the DHCP server on the other side, and can help distinguish where the requests originated.

Upstream Servers

A list of target DHCP server(s). If there is more than one upstream server, click  **Add** to create more entries.

22.4.2 DHCP Relay Configuration

To configure the DHCP Relay:

- Disable the DHCP Server on each interface where the Relay will run
- Navigate to **Services > DHCP Relay**
- Click the tab for the interface to use with DHCP Relay
- Configure the options as described in *DHCP Relay Options*.
- Click **Save**

The DHCPv6 Relay function works identically to the DHCP Relay function for IPv4.

22.5 DNS Resolver

The DNS Resolver in pfSense® software utilizes unbound, which is a validating, recursive, caching DNS resolver capable of using DNSSEC, DNS over TLS, and a wide variety of options. It can act in either a DNS resolver or forwarder role.

Note: The DNS Resolver is enabled in resolver mode by default in current versions of pfSense software.

22.5.1 DNS Resolver Mode

The DNS Resolver can act in either a DNS resolver or forwarder role. These roles are described in detail on *DNS Resolution Process*.

Resolver mode

In resolver mode (default) the DNS Resolver contacts root DNS servers and other authoritative servers directly in search of answers to queries submitted by clients. This eliminates issues typically encountered by users with missing or incorrect local DNS configuration since it does not require forwarding DNS servers to operate. Resolver mode also enables the use of Domain Name System Security Extensions (DNSSEC) which makes the DNS results more trustworthy and verifiable.

Note: Some ISPs block or rate limit these types of DNS queries and instead prefer users to contact forwarders. If resolver mode does not work, use forwarding mode.

As this mode contacts servers which cannot be known beforehand, it must utilize the default route on the firewall to make outbound connections. This may not be optimal with multiple WANs, but there are ways around this limitation such as configuring failover for the default gateway. See *Interface and DNS Configuration*.

Forwarding mode

In forwarding mode the DNS Resolver will forward DNS queries to the list of servers configured under **System > General Setup** or those obtained automatically from a dynamic WAN.

Tip: For increased privacy these forwarded queries can be made using *DNS over TLS*.

This method tends to work better with multiple WANs as each forwarding DNS server may be configured to use a different WAN, allowing queries to be sent over whichever WAN is available at a given moment. See *Interface and DNS Configuration*.

While in forwarding mode the DNS Resolver monitors response timing from all available DNS servers in its infrastructure cache. The daemon will direct queries to servers based on their current status so it can avoid using servers which are slow or unavailable. This data is available in the GUI at **Status > DNS Resolver** (*DNS Resolver Status*).

22.5.2 DNS Resolver Configuration

To configure the DNS Resolver, navigate to **Services > DNS Resolver**

DNS Resolver Options

Enable
Controls whether the DNS Resolver is enabled. Check the box to enable the DNS Resolver service, uncheck to disable the service.

Two DNS services cannot both be active at the same time on the same ports. This includes, but is not limited to, the DNS Resolver, the DNS Forwarder, and the BIND package. Ensure other services are disabled or moved to different ports before attempting to enable the DNS Resolver.

Listen Port
The TCP and UDP port on which the DNS Resolver will listen for queries from clients. By default, this is port `53`. This is the normal port for any DNS server, as it is the port expected by clients.

Certain use cases may involve moving the DNS Resolver to another Listen Port, such as `5353` or `54`, and then specific sources may be forwarded there via port forwards.

Enable SSL/TLS Service
Configures the DNS Resolver to act as a DNS over TLS server which can answer queries from DNS over TLS **clients**.

Note: Activating this option disables automatic interface response routing behavior, thus it works best with specific interface bindings.

SSL/TLS Certificate
The server certificate to use when acting as an SSL/TLS server.

For clients to properly validate the server, they must trust this certificate. One way to accomplish that easily is to use a certificate generated by the *ACME package*.

SSL/TLS Listen Port
The TCP and UDP port on which the DNS Resolver will listen for queries from DNS over TLS clients. By default, this is port `853`.

Network Interfaces
The network interface(s) to which the DNS Resolver will bind when listening for queries from clients.

By default, the DNS Resolver listens on every available interface and IPv4 and IPv6 address. This option limits the interfaces where the DNS Resolver will accept and answer queries. This can be used to increase security in addition to firewall rules.

If specific interfaces are selected, both the IPv4 and IPv6 addresses on those interfaces will be used for answering queries. Additionally, The unbound daemon will only bind to the selected interfaces. Queries sent to other IP addresses on the firewall will be silently discarded.

Outgoing Network Interfaces
Controls which interfaces the firewall will utilize when sending its own queries to other DNS servers.

By default, the DNS Resolver utilizes all interfaces for outbound queries so it will source the query from whichever interface and IP address is closest to the target server from a routing perspective. Selecting specific interfaces will limit the choices to only specific interfaces that may be used as a source of queries.

System Domain Local Zone Type
This option determines the type of `local-zone` configured in `unbound` for the system domain. The zone type governs the type of response given to clients when there is no match in local data such as Host Overrides, DHCP hosts, etc. In each case, if there is a local match, the query is answered normally. The available types to govern non-matching responses are:

Deny
Drops the query and does not answer the client.

Refuse
Notifies the client that the query was refused (Using `rcode REFUSED`).

Static
Returns a `NODATA` or `NXDOMAIN` response to the client.

Transparent
This is the default behavior. If the query is for a name that does not exist locally, it is resolved as usual. If the name has a local match, but the type is different, a `NOERROR`, `NODATA` response is sent to the client

Type Transparent
Similar to *Transparent* but it also passes through queries where the name matches, but the type does not. For example, if a client queries for an AAAA record but only an A record exists, the AAAA query is passed on rather than resulting in a negative response.

Redirect
Handles queries from local data and redirects queries for zones underneath the local zone (e.g. subdomains). This can be used to control queries for all subdomains under the given domain.

Inform
Answers normally, but logs the client query.

Inform Deny
Denies and logs the query.

No default
Disables any default content for the zone without affecting query behavior.

DNSSEC
Enables Domain Name System Security Extensions (DNSSEC), which allows clients to trust the origin and content of DNS responses. This is enabled by default.

DNSSEC protects against manipulation of DNS responses, such as DNS cache poisoning or other query interception, but it does not make the contents of responses secret.

DNSSEC works best when using the root servers directly, unless the forwarding servers also perform DNSSEC validation. Even if the forwarding DNS servers use DNSSEC, the response cannot be fully validated.

DNSSEC may need to be disabled if upstream DNS servers do not perform DNSSEC validation in forwarding mode or with domain overrides, DNS queries are known to be intercepted upstream, or clients have issues with large DNS responses.

Python Module
Enables the DNS Resolver Python module. This feature utilizes a Python script to act on queries or results. For example, a script could prevent certain domains or record type combinations from being resolved.

Python Module Order
Controls the position of the Python module in the DNS resolution process. If DNSSEC

is disabled, this option has no effect.

Pre Validator
The script is run before DNSSEC validation.

Post Validator
The script is run after DNSSEC validation.

Python Module Script
The python script file to execute. The script must be uploaded to the firewall in `/var/unbound/`. The filename must end in `.py`.

DNS Query Forwarding
Controls whether unbound uses resolver mode (unchecked) or forwarding mode (checked). See *DNS Resolver Mode* for an explanation of the modes.

The default is resolver mode (unchecked).

When checked, unbound will use the system *DNS Servers* from **System > General Setup** or those received from a dynamic WAN, rather than using the root servers directly.

Certain situation require or work better with forwarding mode, such as when utilizing *DNS over TLS* for outgoing queries or for *optimal multi-WAN configurations*.

Use SSL/TLS for outgoing DNS Queries to Forwarding Servers
Sends queries to all upstream forwarding DNS servers using SSL/TLS on the default port of `853`. Requires **DNS Query Forwarding** to be checked.

See *Configuring DNS over TLS* for detailed instructions.

> **Warning:** **All** upstream forwarding servers **must** respond to SSL/TLS queries on port `853`.

DHCP Registration
Controls whether internal machine names for DHCP clients are registered in the DNS Resolver. The domain name from **System > General Setup** is used as the domain name on the hosts.

This feature allows systems using the DNS Resolver as their DNS server to resolve these names using DNS.

Note: This only works for clients that specify a hostname in their DHCP requests.

> **Warning:** The DNS Resolver is reloaded when updating hostnames it learns from DHCP lease data. On busy networks with many DHCP clients, this can result in temporary DNS outages as unbound reloads. In most cases this is only a factor when using add-on packages which increase the burden on the DNS Resolver or which make it take longer than usual to reload.

Static DHCP
This works the same as **Register DHCP leases in DNS resolver**, except that it registers the DHCP static mapping addresses.

OpenVPN Client
Controls whether OpenVPN client names are registered in the DNS Resolver.

If this option is set, then the common name (CN) of connected OpenVPN clients will be registered in the DNS Resolver along with the client address inside the VPN. The domain in **System > General Setup** is used as the domain name on these entries.

Note: This option requires an OpenVPN server to be operating in *Remote Access SSL/TLS* mode or in *User Auth* mode with **Username as Common Name** active.

Custom Options
A text area for advanced `unbound` directives not implemented in the GUI.

Tip: If `unbound` does not start correctly after entering custom options, add `server:` on a line at the top of the custom options text area.

22.5.3 Host Overrides

Custom DNS entries can be created in the **Host Overrides** section of the DNS Resolver configuration.

Host overrides define new records or override existing records so that local clients receive the configured responses instead of responses from upstream DNS servers.

This is useful for split DNS configurations (see *Split DNS*) and as a semi-effective means of *blocking access to certain specific websites*.

Warning: Do not use DNS override functionality as the only means of blocking access to sites.

Blocking via DNS requires that local clients utilize the firewall as their only DNS source. See *Redirecting Client DNS Requests* and *Blocking External Client DNS Queries* for suggestions on ensuring clients get their DNS responses from the firewall. It will stop non-technical users, but it is easy to circumvent for those with more technical aptitude.

Multiple records may be defined for the same hostname and all IP addresses will be included in the result. This can be used to supply both an IPv4 (A) and IPv6 (AAAA) result for a single hostname.

Host
This field defines the hostname portion of the DNS override record (without the domain), e.g. `www`.

This may be left blank to make an override record for the domain itself, similar to an `@` record.

Domain
Defines the domain name portion of the DNS override record, e.g. `example.com`.

This field is required.

IP Address
The IP address (either IPv4 or IPv6) to return as the result for a DNS lookup of this entry. Can be a single address or a comma-separated list of multiple addresses.

Description
A text description used to identify or give more information about this entry.

Additional Names for This Host
Defines additional hostnames for the same IP address to keep them in a single override entry.

22.5.4 Domain Overrides

Domain overrides are found at the bottom of the DNS Resolver configuration. These entries specify an alternate DNS server to use for resolving hosts in a specific domain.

A common use of domain overrides is to resolve internal DNS domains at remote sites using a DNS server at the main site accessible over VPN. In such environments all DNS queries are typically resolved at the central site for centralized control over DNS, however some organizations prefer to let Internet DNS resolve with a local caching resolver at each site, and only forward queries for internal domains to the central DNS server (e.g. for *Split DNS*).

Note: A static route may be necessary for this to function over IPsec. See *Accessing Firewall Services over IPsec* for more information.

This can also be leveraged as a semi-effective means of *blocking access to certain specific websites*.

Warning: Do not use DNS override functionality as the only means of blocking access to sites.

Blocking via DNS requires that local clients utilize the firewall as their only DNS source. See *Redirecting Client DNS Requests* and *Blocking External Client DNS Queries* for suggestions on ensuring clients get their DNS responses from the firewall. It will stop non-technical users, but it is easy to circumvent for those with more technical aptitude.

Domain
The domain name that will be resolved using this entry.

This does not have to be a valid TLD, it can be anything (e.g. `local`, `test`, `lab`), or it can be an actual domain name (`example.com`).

IP Address
Specifies the IP Address of the DNS server to which the queries for hostnames in Domain are sent. If the target DNS server is running on a port other than `53`, add the port number after the IP address with an `@` separating the values, for example: `192.0.2.3@5353`

TLS Queries
Controls whether **all** queries for this domain going to this server are sent using SSL/TLS.

TLS Hostname
An optional hostname used to validate the SSL/TLS server certificate.

Description
A text description used to identify or give more information about this entry.

22.5.5 DNS Resolver Advanced Options

pfSense® software provides a GUI to configure some of the more common advanced options available in the DNS Resolver (Unbound).

See also:

The options below are documented as found in the unbound.conf man page.

Advanced Privacy Options

Hide Identity
Controls whether Unbound will allow queries for the server identity. This offers extra privacy.

When set, Unbound rejects queries for `id.server` and `hostname.bind`.

Hide Version
Controls whether Unbound will allow queries for the server version. This offers extra privacy.

When set, Unbound rejects queries for `version.server` and `version.bind`.

Query Name Minimization
Controls whether Unbound attempts to minimize the amount of data sent with a query for extra privacy. Default is unchecked.

When set, Unbound only sends the minimum required labels of the `QNAME` and sets `QTYPE` to `A` when possible.

Note: This is the best effort approach; Unbound sends the full `QNAME` and original `QTYPE` when an upstream server replies with an `RCODE` other than `NOERROR`, except when receiving `NXDOMAIN` from a `DNSSEC` signed zone.

See also:

Refer to RFC 7816 for in-depth information on Query Name Minimization.

Strict Query Name Minimization
Controls whether Unbound performs Query Name Minimization in a strict manner for even stronger privacy at the expense of potentially failing to resolve many queries. Default is unchecked.

When set, enables `QNAME` minimization in strict mode, which does not fall back to sending the full `QNAME` and original `QTYPE` to potentially broken name servers.

This option requires Query Name Minimization.

> **Warning:** Use with extreme caution. A significant number of domains will fail to resolve when this option in enabled.

Advanced Resolver Options

Prefetch Support
Controls whether Unbound prefetches message cache elements before they expire to help keep the cache up to date.

This option can cause an increase of around 10% more DNS traffic and load on the server, but frequently requested items will not expire from the cache.

Prefetch DNS Key Support
Controls whether Unbound fetches DNSKEYs earlier in the validation process when a Delegation Signer record is encountered.

This helps lower the latency of requests but utilizes a little more CPU, and requires the cache to be set above zero.

Harden DNSSEC Data
Controls whether Unbound requires DNSSEC data for trust-anchored zones.

When checked (default), if DNSSEC data is absent in a response for a trust-anchored zone, the zone becomes bogus.

If unchecked and Unbound does not receive DNSSEC data then Unbound behaves as if the zone had no trust anchor; the zone is marked insecure, but the data is still used. This can work around issues with receiving DNSSEC data, but also opens up the results to a potential downgrade attack.

Serve Expired
Controls whether Unbound will serve cache records with a TTL of `0`.

When enabled, allows Unbound to serve a query even with a TTL of `0`. If the TTL is `0` then a new record will be requested in the background when the cache is served to ensure that the cache is updated without adding latency to the client DNS request.

Drop Old UDP Queries
Timeout in seconds before dropping UDP queries waiting in the socket buffer. Queries that have waited for a long time don't need to be processed and can be dropped. A value of 3 should be safe in most setups.

Aggressive NSEC
Controls how aggressively Unbound attempts to predict negative responses based on the contents of the DNSSEC NSEC chain.

When enabled, Unbound uses the DNSSEC NSEC chain to synthesize `NXDOMAIN` and other denials, using information from previous `NXDOMAIN` answers. This helps to reduce the query rate towards targets with high nonexistent name lookup rates.

Message Cache Size
Controls the amount of memory used to cache DNS response codes and validation statuses. The default is *4 MB*.

Note: The resource record set (RRSet) cache will automatically be set to twice this amount. The RRSet cache contains the actual resource record data.

Outgoing TCP Buffers
The number of outgoing TCP buffers to allocate per thread. The default value is *10*.

If set to *0*, TCP queries will not be sent to authoritative servers.

Incoming TCP Buffers
The number of incoming TCP buffers to allocate per thread. The default value is *10*.

If set to *0*, TCP queries will not be accepted from clients.

EDNS Buffer Size
Number of bytes size to advertise as the EDNS reassembly buffer size. This value is placed in UDP datagrams sent to peers.

The default is *Automatic* and is calculated based on the MTU values of active interfaces. A variety of other common values are provided in a drop-down list.

Automatic mode sets optimal buffer size by using the smallest MTU of active interfaces and subtracting the IPv4/IPv6 header size. If fragmentation reassembly problems occur, usually seen as timeouts, then try a value of `1432`.

The `512`, `1220`, and `1232` values bypass most IPv4 and IPv6 MTU path problems but can generate an excessive amount of TCP fallback.

Number of Queries per Thread
The number of queries that every Unbound thread will service simultaneously. The default value is *512*.

If additional queries arrive that need to be serviced, and no queries can be jostled out, the new queries are dropped

Jostle Timeout
Timeout in milliseconds used when the server is very busy. This protects against denial of service by slow queries or high query rates. The default value is *200* milliseconds.

Set to a value that approximates the round-trip time to the authority servers. As new queries arrive, 50% are allowed to run and 50% are replaced by new queries if they are older than the stated timeout.

Maximum TTL for RRsets and Messages
The Maximum Time to Live (TTL) for RRsets and messages in the cache, specified in seconds. The default is `86400` seconds (1 day).

When the internal TTL expires the cache item is expired. This can be configured to force the resolver to query for data more often and not trust very large TTL values

Minimum TTL for RRsets and Messages
The Minimum Time to Live for RRsets and messages in the cache, specified in seconds. The default is `0` seconds.

If a record has a TTL lower than the configured minimum value, data can be cached for longer than the domain owner intended, and thus fewer queries are made to look up the data. The `0` value ensures the data in the cache is not kept longer than the domain owner intended.

> **Warning:** High values can lead to trouble as the data in the cache may not match up with the actual data if it changes.

TTL for Host Cache Entries
Time to Live, in minutes, for entries in the infrastructure host cache. The default value is *15 minutes*.

The infrastructure host cache contains round trip timing, lameness, and EDNS information for DNS servers.

Number of Hosts to Cache
Number of infrastructure hosts for which information is cached. The default is *10,000*.

Unwanted Reply Threshold
Controls whether Unbound tracks the total number of unwanted replies in every thread. The default is disabled

When the threshold is reached, a defensive action is taken and a warning is printed to the log file. The defensive action is to clear the RRSet and message caches, hopefully flushing away any poison.

If enabled, a value of 10 million is the best starting value.

Log Level
Controls the verbosity of data logged by Unbound. Default is *Level 1*.

Level 0
No logging, only errors.

Level 1
Basic operational information.

Level 2
Detailed operational information.

Level 3
Query level information, output per query.

Level 4
Algorithm level information.

Level 5
Logs client identification for cache misses.

Disable Auto-added Access Control
Controls whether Unbound uses automatic access control entries. The default is unchecked.

The automatic access control entries permit queries from IPv4 and IPv6 networks residing on internal interfaces of this firewall, and on certain known subnets used for VPNs.

When checked, networks from which queries are allowed must be manually configured on the **Access Lists** tab.

Disable Auto-added Host Entries
Controls whether Unbound registers primary IPv4 and IPv6 addresses of this firewall as records for the system domain as configured in **System > General Setup**.

When checked, these automatic entries are omitted from the configuration.

Experimental Bit 0x20 Support
Use `0x20` encoded random bits in the DNS query to foil spoofing attempts. See the implementation draft dns-0x20 for more information:

DNS64 Support
Controls whether Unbound will provide DNS64 (RFC 6147) records in responses. This is typically used with features such as *NAT64* to enable communication from IPv6-only clients to IPv4-only servers.

See also:

- *NAT64*
- *Configuring NAT64 for IPv6-only Clients*

22.5.6 DNS Resolver Access Lists

Unbound requires access lists (ACLs) to control which clients are allowed to submit queries. By default, IPv4 and IPv6 networks residing on internal interfaces of this firewall are permitted. Additional networks must be allowed manually.

Note: The automatic ACLs may be disabled using the **Disable Auto-added Access Control** option on the **Advanced Settings** tab.

To manage access lists for the DNS Resolver, navigate to **Services > DNS Resolver**, **Access Lists** tab. This page has controls to add new entries as well as edit or delete existing entries.

When adding or editing an entry, the following options are available:

Access List Name
The name for the access list, which appears as a comment in the access list configuration file.

Action
Controls how Unbound will handle queries for networks contained in this access list.

Deny
Stops queries from clients in the configured networks

Refuse
Stops queries from clients in the configured networks and sends back a REFUSED response code

Allow
Allows queries from clients in the configured networks

Allow Snoop
Allows recursive and non-recursive queries from clients in the configured networks, used for cache snooping, and typically only configured on administrative hosts.

Deny Nonlocal
Allow only authoritative local-data queries from hosts within the network on this ACL. Unbound will drop disallowed messages.

Refuse Nonlocal
Allow only authoritative local-data queries from hosts within the network on this ACL. Unbound sends back a REFUSED response code for disallowed messages.

Description
A longer text field for reference notes about this entry.

Networks
A list of IPv4 or IPv6 networks governed by this access list entry.

22.5.7 CLI Commands

Unbound provides various command line utilities to manage the DNS Cache server. The following control commands are currently not available in the GUI but can be executed from the command line.

Note: Unbound does not use the default configuration file location; Use the `-c` flag to tell Unbound the configuration file location:

```
unbound-control -c /var/unbound/unbound.conf <unbound-command-to-run>
```

Remove <name> from the cache, all record types which include A, AAAA, NS SOA, CNAME, DNAME, MX, PTR, SRV and NAPTR records:

```
unbound-control -c /var/unbound/unbound.conf flush <name>
```

Remove the <name> and <type> from the cache where <type> is a particular record type:

```
unbound-control -c /var/unbound/unbound.conf flush_type <name> <type>
```

Remove all information at or below the <name> from cache. For example, `.com` will remove all entries below `.com`. Note this process is slow as the entire cache must be inspected:

```
unbound-control -c /var/unbound/unbound.conf flush_zone <name>
```

Reload Unbound and clear the entire cache:

```
unbound-control -c /var/unbound/unbound.conf reload
```

Determine the name servers Unbound will Query to look up a zone:

```
unbound-control -c /var/unbound/unbound.conf lookup <name>
```

See also:

- *DNS Resolver Status*

22.6 DNS Forwarder

The DNS Forwarder in pfSense® software utilizes the `dnsmasq` daemon, which is a caching DNS forwarder.

Unlike the DNS Resolver, the DNS Forwarder can only act in a forwarding role. It is not capable of acting as a resolver.

The DNS Forwarder uses *DNS Servers* configured at **System > General Setup** and those obtained automatically from an ISP for dynamically configured WAN interfaces (DHCP, PPPoE, etc.).

See also:

- *DNS Resolution Process*
- *DNS Rebinding Protections*

Note: This service is disabled by default. The *DNS Resolver* (`unbound`) is the default DNS service.

The DNS Forwarder remains enabled on upgraded installations where it was active before the upgrade.

22.6.1 DNS Forwarder Configuration

To configure the DNS Forwarder, navigate to **Services > DNS Forwarder**

The available options for the DNS Forwarder are:

Enable
: Controls whether the DNS Forwarder service is enabled.

  Checking this box turns on the DNS Forwarder, or uncheck to disable this service.

  Two DNS services cannot both be active at the same time on the same ports. This includes, but is not limited to, the DNS Resolver, the DNS Forwarder, and the BIND package. Ensure other services are disabled or moved to different ports before attempting to enable the DNS Forwarder.

DHCP Registration
: Controls whether internal machine names for DHCP clients are registered in the DNS Forwarder. The domain name from **System > General Setup** is used as the domain name on the hosts.

  This feature allows systems using the DNS Forwarder as their DNS server to resolve these names using DNS.

  Note: This only works for clients that specify a hostname in their DHCP requests.

Static DHCP
: This works the same as **Register DHCP leases in DNS forwarder**, except that it registers the DHCP static mapping addresses.

Prefer DHCP
Controls whether DNS records from DHCP sources are returned before host overrides if both use the same name.

When one IP address has multiple hostnames, doing a reverse lookup may give an unexpected result if one of the hostname is in host overrides and the system uses another hostname over DHCP. Checking this option will place the DHCP obtained hostnames above the static mappings in the hosts file on the firewall, causing them to be consulted first.

This only affects reverse lookups (PTR), since they only return the first result and not multiple. For example, this would yield a result of `labserver01.example.com`, a test server DHCP obtained IP address, rather than a host override name of `testwww.example.com` that would be returned otherwise.

Query DNS servers sequentially
Controls whether the DNS Forwarder queries all DNS servers at the same time, or in sequence.

By default, the firewall queries all DNS servers simultaneously and uses the fastest result. This is not always desirable, especially if there is a local DNS server with custom hostnames that could be bypassed if a faster public DNS server replies first.

Checking this option causes queries to be made to each DNS server in sequence from the top down, and the firewall waits for a timeout before moving on to the next DNS server in the list. This results in more predictable responses but may be considerably slower if a server high in the list is unreachable.

Require domain
Controls whether the DNS Forwarder requires a domain name on hostnames to be forwarded to upstream DNS servers.

When checked, hosts without a domain name will still be checked against host overrides and DHCP results, but they will not be queried against name servers. If a short hostname does not exist locally, an NXDOMAIN result ("Not Found") is returned to the client.

Do not forward private reverse lookups
Controls whether the DNS Forwarder will make reverse DNS (PTR Record) lookups for RFC1918 private IP addresses to upstream name servers.

The DNS Forwarder will still return results from local entries in either case.

Tip: Use a domain override entry for the reverse lookup zone, e.g. `1.168.192.in-addr.arpa`, to make the DNS Forwarder send queries for a specific subnet to a DNS server.

Listen Port
The TCP and UDP port on which the DNS Forwarder will listen for queries from clients. By default, this is port `53`. This is the normal port for any DNS server, as it is the port expected by clients.

Certain use cases may involve moving the DNS Forwarder to another Listen Port, such as `5353` or `54`, and then specific sources may be forwarded there via port forwards.

Interfaces
The network interface(s) to which the DNS Forwarder will bind when listening for queries from clients.

By default, the DNS Forwarder listens on every available interface and IPv4 and IPv6 address. This option limits the interfaces where the DNS Forwarder will accept and answer queries. This can be used to increase security in addition to firewall rules.

If specific interfaces are selected, both the IPv4 and IPv6 addresses on those interfaces will be used for answering queries. Queries sent to other IP addresses on the firewall will be silently discarded.

Strict Interface Binding
Controls how the `dnsmasq` daemon binds to interfaces when deciding how to handle queries.

When set, the DNS forwarder will only bind to the interfaces containing the IP addresses selected in the **Interface** control, rather than binding to all interfaces and discarding queries to other addresses.

This can be used similarly to the **Listen Port** for controlling the way that the service binds so that it can coexist with other DNS services that have similar options.

Note: This option is not compatible with IPv6. If this is checked, the `dnsmasq` daemon will not bind to any IPv6 addresses.

Advanced Options

Custom `dnsmasq` configuration parameters that are not configurable in the GUI can be placed in **Advanced Options**. Separate each command by either a space or a newline.

For example, to set a lower TTL for DNS records:

```
max-ttl=30
```

To craft a wildcard DNS record resolving `*.lab.example.com` to `192.2.5.6`:

```
address=/lab.example.com/192.2.5.6
```

See also:

For more information on the possible parameters that may be used, consult the dnsmasq documentation.

22.6.2 Host Overrides

Host override entries provide a means to configure customized DNS entries. The configuration is identical to *Host Overrides* in the DNS Resolver, refer there for details. The main difference is that overrides in the DNS Forwarder can only contain a single address per entry.

22.6.3 Domain Overrides

Domain overrides configure an alternate DNS server to use for resolving a specific domain. The configuration is similar to *Domain Overrides* in the DNS Resolver, but there are a few differences:

IP Address
This field can be used in one of three ways to control how the DNS Forwarder handles queries for the given **Domain**:

- To specify the IP Address of a DNS server to which the DNS Forwarder will send queries for hostnames in the **Domain**.
- To override another entry by entering `#`.

  For example, to forward `example.com` to `192.2.66.2`, but have `lab.example.com` forward on to the standard name servers, enter a `#` in this field.
- To prevent non-local lookups by entering a `!`.

If host override entries exist for `www.example.org` and `mail.example.org`, but other lookups for hosts under *example.org* must not be forwarded on to remote DNS servers, enter a `!` in this field.

Source IP
The source IP address that the DNS Forwarder will use when sending queries to the DNS server in this domain override.

This field is optional and primarily used to contact a DNS server across a VPN where the VPN requires queries to have a specific source.

22.6.4 DNS Forwarder Behavior

By default, the DNS Forwarder queries all DNS servers at once, and it uses and caches only the first response it receives. This results in much faster DNS service from a client perspective, and can help smooth over problems that stem from DNS servers which are intermittently slow or have high latency, especially in Multi-WAN environments. This behavior can be disabled by activating the **Query DNS servers sequentially** option.

See also:

- *Interface and DNS Configuration*

22.7 Dynamic DNS

The Dynamic DNS client built into pfSense® software registers the IP address of a WAN interface with a variety of dynamic DNS service providers. This is used to remotely access services on hosts that have WANs with dynamic IP addresses, most commonly VPNs, web servers, and so on.

Any number of Dynamic DNS clients may be configured using any of over 20 different Dynamic DNS providers, or even custom Dynamic DNS providers. Dynamic DNS clients can use any WAN, and can even register the real public IP address in environments where the firewall receives a private IP address for its WAN and is behind upstream NAT.

In addition to the typical HTTP/HTTPS-based Dynamic DNS providers, pfSense software can also update records using RFC 2136 style Dynamic DNS updates.

22.7.1 Configuring a Dynamic DNS Client

pfSense® software can use Dynamic DNS to automatically update DNS providers when an interface address changes. This allows remote clients to reference a constant hostname instead of a dynamic IP address which could change over time.

This service is located in the GUI at **Services > Dynamic DNS**.

Choosing a Dynamic DNS Provider

pfSense software allows registration with many different dynamic DNS providers. The available providers may be viewed by clicking the **Service Type** selector. More information about the providers may be found by searching for their name to find their website. Several providers offer a basic level service at no cost, and some offer additional premium services at a cost. There is also a *Custom* option that allows for a custom URL to accommodate a provider not present in the GUI.

Select a provider, visit their website, register for an account, and set up a hostname. The procedures for this vary with each provider, but they all have instructions on their websites. After configuring a hostname with a provider, configure the firewall with matching settings.

Dynamic DNS Settings

Most providers have the same, or similar options. There are a few types with custom options that will be covered later in this section.

Disable
Check to disable the entry, or leave unchecked so it will be active.

Service Type
Select the dynamic DNS provider here.

Interface to Monitor
Select the interface that has the IP address to keep updated, such as WAN, or an OPTx interface. Selecting a gateway group for the interface allows the Dynamic DNS entry to switch between WANs so it can allow inbound Multi-WAN failover of services on this hostname.

Hostname
Enter the hostname created at the dynamic DNS provider. This is typically the complete fully qualified domain name, such as `myhost.example.com`, except for Namecheap where this is only the host portion of the address.

Domain Name
For Namecheap hosts, this box must be set to the domain part of the full hostname.

MX
An MX (Mail Exchanger) record is how Internet mail servers know where to deliver mail for a domain. Some dynamic DNS providers will let MX records be configured via the dynamic DNS client. If the chosen provider allows this, enter the host name of the mail server that will receive Internet mail for the dynamic DNS domain.

Wildcards
When wildcard DNS is enabled on a dynamic DNS name, all host name queries under the given domain will resolve to the IP address of the dynamic DNS host name. For example, if the host name is `example.dyndns.org`, enabling wildcard will make `*.example.dyndns.org` (`a.example.dyndns.org`, `b.example.dyndns.org`, etc.) resolve the same as `example.dyndns.org`.

Verbose Logging
Check this option to increase the logging for the Dynamic DNS update process, which is useful for troubleshooting update problems.

Verify SSL Peer
When checked, the SSL certificate of the DynDNS provider server will be validated. Some servers with self-signed certificates, or those using a less common CA, may require this to be set.

Username
Enter the username for the dynamic DNS provider. Provider-specific requirements:

Namecheap, FreeDNS
Leave blank

Route 53
Enter the **Access Key ID**

GleSYS
Enter the **API user**

Custom
The username is used with basic HTTP authentication and may be left blank.

Password
Enter the password for the dynamic DNS provider. Provider-specific requirements:

Namecheap, FreeDNS
This is the **Authentication Token**

Route 53
Enter the **Secret Access Key**

GleSYS
Enter the **API Key**

DNSimple
Enter the **API Token**

Description
A text field for reference.

Providers with Extra or Different Settings

Some providers have special settings or certain fields that need to be set in a specific way that may not be obvious. The differences are outlined in this section.

Namecheap

As mentioned in the settings, Namecheap requires that the fully qualified domain name be split into the hostname part and domain name part in separate fields.

When setting up Dynamic DNS for a *Namecheap* domain, an authentication token is given by Namecheap. This goes in the **Password** field, and the **Username** field is left blank.

HE.net Tunnelbroker

The *HE.net Tunnelbroker* choice updates an IPv6 tunnel endpoint IP address when the WAN IP changes. The **Hostname** in this case is the **Tunnel ID** from HE.net.

Route 53

When using an Amazon *Route 53* type, the Username is the **Access Key ID** provided by Amazon.

The following additional options are available when using *Route 53*:

Verify SSL Peer
Enable to verify the server certificate when using HTTPS.

Zone ID
Received when creating the domain in Route 53.

This field is required.

TTL
Time to Live for the DNS record.

Custom

The *Custom* Dynamic DNS type configures options that allow for updating DNS services not included in the GUI. When using the custom Dynamic DNS type, the **Username** and **Password** fields are sent using HTTP basic authentication.

The following additional options are available when using *Custom*:

Interface to send update from
: Almost always the same as the Interface, but can be changed as needed.

Force IPv4 Resolving
: When checked, the update host will only be resolved using IPv4

Verify SSL Peer
: Enable to verify the server certificate when using HTTPS

Update URL
: The URL given by the Dynamic DNS provider for updates. If the IP address must appear in the URL, enter it as `%IP%` and the real value will be substituted as needed.

Result Match
: Defines expected output from the Dynamic DNS query. If it succeeds and matches the output given, then the firewall will know that the update was successful. If it does not match exactly, then it is assumed that the update failed. Leave empty to disable result checking.

DNSSimple

Verify SSL Peer
: Enable to verify the server certificate when using HTTPS.

Zone ID
: Received when creating the domain.

TTL
: Time to Live for the DNS record.

Configuring a Dynamic DNS Entry

To configure a Dynamic DNS client:

- Navigate to **Services > Dynamic DNS**
- Click  **Add** to add a new entry
- Configure the Dynamic DNS entry with general and provider-specific settings
- Click **Save**

22.7.2 Configuring RFC 2136 Dynamic DNS updates

RFC 2136 Dynamic DNS registers a hostname on any DNS server capable of receiving RFC 2136 style updates. This can be used to update DNS records on BIND and Windows Server DNS servers, amongst others.

RFC 2136 Dynamic DNS entries may be used at the same time as regular style Dynamic DNS service providers, and like those, any number of entries can be created. RFC 2136 will update the A record, and the AAAA record if IPv6 is configured on the monitored interface.

See also:

Configuring the server infrastructure for RFC 2136 Dynamic DNS hosting is beyond the scope of this documentation, but there is a basic how-to in the recipes section: *Configuring BIND as an RFC 2136 Dynamic DNS Server*.

RFC 2136 Settings

Enable
Controls whether the entry is active. If it is unchecked, updates will not be performed for this entry.

Interface
The IP address on the chosen interface will be sent when performing the DNS update.

Hostname
The fully qualified domain name (FQDN) of the dynamic DNS entry to update. For example, `myhost.example.com`.

Zone
The hostname of the zone to update (optional).

TTL
The Time To Live for the DNS entry, in seconds. Higher values will be cached longer by other name servers, so lower values are better to be sure that DNS updates are picked up in a timely manner by other servers. Usually a value between `30` and `180` seconds is reasonable, depending on how often the IP address changes.

Key Name
The name of the key as specified in the DNS server configuration. For Host keys, this is typically the FQDN, so it would be identical to the value in the **Hostname** field. For Zone keys this would be the name of the DNS zone.

Key Algorithm
The algorithm used for the key.

Key
Secret TSIG domain key. Contains the actual text of the key, e.g. `/0/4bxF9A08n/zke/vANyQ==`. This value is generated by the DNS server or administrator.

Server
The IP address or hostname of the DNS server to which updates are sent.

Protocol
When unchecked, the DNS update is sent over UDP, when checked it uses TCP instead.

Use Public IP
By default, the interface IP address is always sent to the name server for the DNS update. If this box is checked, when a private IP address is detected on the selected **Interface**, a check is done to determine what the actual public IP address is, and then that IP address is used for the DNS update.

Update Source
Interface or address from which the firewall will send the DNS update request.

Update Source Family
Address family to use for sourcing updates (IPv4 or IPv6)

Record Type
Determines which record(s) will be updated for this entry. For the IPv4 address, use *A*, for IPv6, use *AAAA*, or choose *Both*.

Description
A free-text description of the entry for reference.

Configuring an RFC 2136 Client

To configure an RFC 2136 Dynamic DNS client:

- Navigate to **Services > Dynamic DNS**
- Click the **RFC 2136** tab
- Click  **Add** to add a new entry
- Configure the options
- Click **Save**

As with the other Dynamic DNS types, RFC 2136 updates are performed only when an IP address change is detected, or once every 25 days.

22.7.3 Configuring IP Address Check Services for Dynamic DNS

pfSense® software can utilize custom IP address check services. These services are used by Dynamic DNS clients to determine the public IP address of the firewall when a WAN interface is behind an upstream NAT device.

To create or edit one of these services, navigate to **Services > Dynamic DNS** on the **Check IP Services** tab.

Settings

The following settings are available for a service entry:

Enable
Allow this service to be used by Dynamic DNS clients.

Name
A short name to identify this service.

URL
The full URL to the IP address check page.

Username/Password
Optional authentication to use when accessing the URL.

Verify SSL Peer
Check this box if the server has a self-signed SSL certificate or a certificate from a CA that is not trusted by the firewall.

Description
A longer description of this service.

Once a service is defined, it may be selected on individual Dynamic DNS service entries.

Server-Side Configuration Examples

Hosting one of these services is very simple. The server page need only print the requesting client IP address in the expected format:

```
Current IP Address: x.x.x.x
```

nginx (internal/native)

```
location /ip {
    default_type text/html;
    return 200 "<html><head><title>Current IP Check</title></head><body>Current IP␣
↪Address: $remote_addr</body></html>";
}
```

nginx (internal with LUA)

```
location = /ip {
    default_type text/html;
    content_by_lua '
        ngx.say("<html><head><title>Current IP Check</title></head><body>Current IP␣
↪Address: ")
        ngx.say(ngx.var.remote_addr)
        ngx.say("</body></html>")
    ';
}
```

PHP

```
<html>
     <head>
             <title>Current IP Check</title>
     </head>
     <body>
             Current IP Address: <?=$_SERVER['REMOTE_ADDR']?>
     </body>
</html>
```

22.8 SNMP

The Simple Network Management Protocol (SNMP) daemon enables remote monitoring of certain pfSense® software parameters. The SNMP daemon facilitates monitoring of network traffic, network flows, PF queues, and general system information such as CPU, memory, and disk usage.

The SNMP implementation is bsnmpd, which by default only has the most basic management information bases (MIBs) available, and is extended by loadable modules. In addition to acting as an SNMP daemon, it can also send traps to

an SNMP server for certain events. These vary based on the modules loaded. For example, network link state changes will generate a trap if the MIB II module is loaded.

The SNMP service can be configured by navigating to **Services > SNMP**.

The easiest way to see the available data is to run `snmpwalk` against the firewall from another host with `net-snmp` or an equivalent package installed. The full contents of the MIBs available are beyond the scope of this documentation, but there are plenty of print and online resources for SNMP, and some of the MIB trees are covered in RFCs. For example, the Host Resources MIB is defined by RFC 2790.

See also:

The Hangouts Archive contains a video which covers monitoring via SNMP.

22.8.1 SNMP Daemon

These options dictate if, and how, the SNMP daemon will run.

Enable
Controls whether the SNMP daemon will run.

Polling Port
SNMP connections are made using only UDP, and SNMP clients default to using UDP port 161. This setting controls which port the SNMP daemon uses when listening for client queries.

SNMP clients and/or polling agents must be set to match this value.

System location
A string to return when an SNMP client requests the system location.

Any text may be used here. For some devices a city or state may be close enough, while others may need more specific detail such as which rack and position in which the system resides.

System contact
A string defining contact information for the system. It can be a name, an e-mail address, a phone number, or whatever is needed.

Read Community String
With SNMP, the community string acts as a kind of username and password in one. SNMP clients will need to use this community string when polling.

Tip: The default value of `public` is common, so the best practice is to use a different value in addition to restricting access to the SNMP service with firewall rules.

22.8.2 SNMP Traps

Controls SNMP Trap behavior.

Enable
When set, the SNMP daemon will generate SNMP traps. Additionally, when set, the GUI displays options to control SNMP trap behavior.

Trap server
The hostname or IP address to which the SNMP daemon will forward SNMP traps.

Trap server port
The port on which the trap server is listening for traps.

By default, SNMP traps are set on UDP port `162`. If the SNMP trap server is set for a different port, adjust this setting to match.

SNMP trap string
The SNMP daemon sends this string along with any SNMP trap.

22.8.3 Modules

Loadable modules allow the SNMP daemon to understand and respond to queries for additional system information. Each loaded module consumes additional resources. As such, ensure that only required modules are loaded.

MibII
This module provides information specified in the standard MIB II tree, which covers networking information and interfaces. Having this module loaded will provide network interface information including status, hardware and IP addresses, the amount of data transmitted and received, and much more.

Netgraph
The netgraph module provides netgraph-related information such as netgraph node names and statuses, hook peers, and errors.

PF
The PF module provides a wealth of information about the PF packet filter. The MIB tree covers aspects of the ruleset, states, interfaces, tables, and ALTQ queues.

Host Resources
This module provides information about the host itself. This includes uptime, load average and processes, storage types and usage, attached system devices, and even installed software.

Note: This module requires MibII. If MibII is unchecked when this option is checked, MibII will be checked automatically.

UCD
This module provides various system information knows as the `ucdavis` MIB, or UCD-SNMP-MIB. It provides information about memory usage, disk usage, running programs, and more.

Regex
The Regex module is reserved for future use or use by users customizing the code to their needs. It allows creating SNMP counters from log files or other text files.

22.8.4 Interface Binding

Binding to a specific local interface can ease communication over VPN tunnels as it eliminates the need for *workarounds like static routes*. It also provides extra security by not exposing the service to other interfaces. It can also improve communication over multiple local interfaces, since the SNMP daemon will reply from the "closest" address to a source IP address and not the IP address to which a client sent its query.

Internet Protocol
This controls whether the SNMP daemon will listen for queries on IPv4, IPv6, or both.

Bind Interfaces
This option configures the SNMP daemon to listen only on the chosen interface or virtual IP address. All interfaces with IP addresses, CARP VIPs, and IP Alias VIPs are displayed in the drop-down list.

22.9 UPnP IGD & PCP

Universal Plug and Play Internet Gateway Device (UPnP IGD) and *Port Control Protocol* (PCP) are network protocols which allow local software and devices to configure each other when attaching to a network. This includes autonomously creating dynamic NAT rules to redirect and pass incoming connections from remote hosts.

Note: PCP is the successor to NAT Port Mapping Protocol (NAT-PMP) and is compatible with clients using NAT-PMP.

The UPnP IGD & PCP service, located at **Services > UPnP IGD & PCP**, enables client devices such as computers and game consoles to autonomously allow required inbound traffic and can account for outbound NAT on traffic using the same ports. There are many popular programs and platforms which support UPnP IGD & PCP, such as Steam/Steam Deck, Nintendo consoles, PlayStation consoles, Xbox consoles, video conferencing apps, torrent clients, and more. PCP is supported primarily by Apple products but can also be found in other applications and devices.

See also:

For advice on specific consoles and games, see *Configuring pfSense Software for Online Gaming*.

22.9.1 UPnP IGD & PCP Service Ports

The UPnP IGD daemon used by pfSense® software, `miniupnpd`, uses several ports to communicate with clients. UPnP IGD employs the Simple Service Discovery Protocol (SSDP) for network discovery, which uses UDP port `1900`. PCP uses UDP port `5351`. The daemon also uses TCP port `2189` for HTTP and SOAP queries.

Table 1: UPnP IGD & PCP Ports

Protocol	Port	Service
UDP	1900	SSDP
TCP	2189	HTTP/SOAP
UDP	5351	PCP

When using a strict LAN firewall ruleset, manually add rules to allow access to these services, especially if the default LAN-to-any rule has been removed, or in bridge configurations.

22.9.2 UPnP IGD & PCP with IPv6

As of this writing, the UPnP IGD & PCP service on current versions of pfSense software is compatible with IPv6, but IPv6-capable client implementations are still rare.

Note: IPv6 UPnP IGD & PCP client traffic will almost always require manual rules to pass. Clients are likely to use a link-local source going to a multicast destination, which is not covered by the default interface rules.

22.9.3 Security Concerns

UPnP IGD & PCP are a classic example of the "Security vs. Convenience" trade-off. By their very nature, these services are insecure. Any program on the network can allow in and forward any traffic – a potential security nightmare. On the other side, it can be a chore to enter and maintain NAT port forwards and their associated rules, especially when it comes to game consoles. There is a lot of guesswork and research involved to find the proper ports and settings, but UPnP IGD & PCP *just works* and requires little administrative effort. Manual port forwards to accommodate these scenarios tend to be overly permissive, potentially exposing services that should not be open from the Internet. Manual port forwards are also always enabled, where UPnP IGD & PCP rules may be temporary.

Access controls in the UPnP IGD & PCP service configuration can lock down which devices are allowed to make port mappings and for which ports. Over and above the built-in access controls, further control may be exerted with firewall rules. When properly controlled, UPnP IGD & PCP can also be a little more secure by allowing programs to pick and listen on random ports, instead of always having the same port open and forwarded.

22.9.4 Configuration Options

Service Settings

Enable port mapping service
: Master control for the entire UPnP IGD & PCP daemon. Also requires enabling one or both of the UPnP IGD or PCP/NAT-PMP protocols.

  When unchecked, all the services on this page are disabled.

Allow UPnP IGD Port Mapping
: When checked, clients can make requests using Universal Plug and Play Internet Gateway Device (UPnP IGD).

Allow PCP/NAT-PMP Port Mapping
: When checked, clients can make requests using Port Control Protocol (PCP) and client requests using PCP-compatible protocols, such as the older NAT-PMP standard.

External Interface
: The interface for **outgoing** traffic. This must be set to the single WAN containing the default gateway.

Internal Interfaces
: The **local** interfaces where clients allowed to use UPnP IGD & PCP reside. Multiple interfaces may be selected.

Note: When using a bridge configuration, only select the bridge interface with an IP address.

External Address Settings

Port mapping with the UPnP IGD & PCP service requires a routable public IP address capable of receiving incoming connections from remote hosts. If this device is behind NAT, port forwarding may still be possible, provided the upstream router can forward traffic using an additional layer of NAT.

As the public address is key to its operation, the UPnP IGD & PCP port mapping service **must** be able to locate its routable public IP address, or it will refuse to allow port mapping.

If the **External Interface** has a public IP address, the settings in this section may be left empty or disabled (default) since the service can recognize that the interface address is public. If the interface has a private address, the service will refuse to map any ports without additional configuration.

If the **External Interface** is behind unrestricted NAT (e.g. 1:1 NAT), and the upstream router forwards incoming traffic to this device without any filtering, this service can still function so long as it can locate the correct public IP address.

The service can learn the public IP address and NAT type using an external server via the STUN protocol or if the address is static it can be hard-coded using the **Override WAN Address** setting.

Enable STUN
Enables retrieving the external IP address from a remote STUN server. When checked, the service will query an external STUN server to locate the routable public IP address for the external interface. This also tests inbound NAT connectivity to ensure forwarded traffic is delivered to this device.

This is useful for devices behind 1:1 NAT with a dynamic public address.

> **Warning:** The service will disable port mapping if NAT testing determines inbound connections are not being forwarded to this device.

STUN Server
The hostname or IP address of a remote STUN server.

There are public STUN servers available, including:

- `stun.3cx.com`
- `stun.counterpath.com`

STUN Port
The UDP port on which the STUN server is listening for client connections. The default port number is `3478`.

Override WAN Address
This option manually configures a public routable IP address this device can use to accept inbound connections.

This address can be an upstream static routable IP address with traffic forwarded to this device via NAT, or it can be an alternate routable IP address on this device, such as a virtual IP address.

Advanced Settings

Download Speed
Value to report when clients query the maximum link download speed, specified in kbit/s.

The default value is the link speed of the interface.

Upload Speed
Value to report when clients query the maximum link upload speed, specified in kbit/s.

The default value is the link speed of the interface.

Traffic Shaping
The name of an ALTQ (not Limiter) traffic shaping queue in which the firewall will place traffic passed by rules created via UPnP IGD & PCP.

Note: Exercise caution when selecting this queue. UPnP IGD & PCP can be used by traffic such as game consoles, which need high priority, and also by file transfer clients which may need low priority.

Custom Presentation URL
A custom URL this daemon presents to UPnP IGD & PCP clients who click this device when listing devices on the local network. For example, when browsing the network in Windows Explorer.

When left blank, the daemon uses the URL of the GUI on this device.

Custom Model Number
A custom model number presented to clients who click this device when listing devices on the local network. For example, when browsing the network in Windows Explorer.

When left blank, the daemon uses the current firmware version of this device.

Firewall Logs
When checked, rule generated by UPnP IGD & PCP will be set to log, so each connection matching these rules will appear in the firewall logs.

Note: The firewall logs can be viewed at **Status > System Logs**, on the **Firewall** tab.

Uptime
By default, the UPnP IGD & PCP daemon reports the service uptime when queried rather than the system uptime. Checking this option will cause it to report the system uptime (since last boot) instead.

Service Access Control Lists

Default Deny
When checked, UPnP IGD & PCP only allows access to client requests matching configured access control lists. This is a more secure method of controlling the service, but as discussed above, is also less convenient.

ACL Entries
ACL entries grant or deny access to the port mapping service based on several criteria.

> **Warning:** These entries only control access for IPv4 clients, they do not apply to IPv6 clients.

Rules are formulated using the following format:

```
<[allow|deny]> <[external port|range]> <[internal IP|IP/CIDR]> <[internal␣
↪port|range]>
```

Click  **Add** to create additional rules.

Note: If the **Default Deny** option is enabled, rules must be set to allow access.

UPnP IGD & PCP User Permission Examples

Deny access to external port `80` forwarding from everything on the LAN, `192.168.1.0`, with a `/24` subnet, to local port `80`:

```
deny 80 192.168.1.0/24 80
```

Allow `192.168.1.10` to forward any unprivileged port:

```
allow 1024-65535 192.168.1.10 1024-65535
```

22.9.5 Configuration Procedure

To configure UPnP IGD & PCP:

- Navigate to **Services > UPnP IGD & PCP**
- Configure the options as needed
- Click **Save**

The UPnP IGD & PCP service will be started automatically.

22.9.6 UPnP IGD & PCP Status

To view a list of currently forwarded ports and clients, navigate to **Status > UPnP IGD & PCP**. The output will be similar to *UPnP IGD & PCP Status Screen Showing Client PCs With Forwarded Ports*.

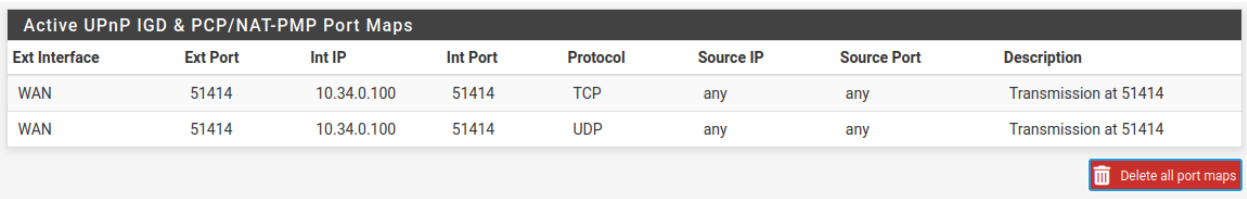


Fig. 1: UPnP IGD & PCP Status Screen Showing Client PCs With Forwarded Ports

View the status of the UPnP IGD & PCP daemon at **Status > Services**. The Service Status page shows if the daemon is running or stopped, and allows the service to be stopped, started or restarted. Under normal circumstances, manually managing the daemon is not necessary.

22.9.7 Troubleshooting

Most issues with UPnP IGD & PCP tend to involve clients not being able to communicate with the service. It is important to have firewall rules allow UPnP IGD traffic on UDP port `1900`. Since UPnP IGD uses multicast traffic, the destination will be the broadcast address for the subnet, or in some cases a destination of *any* will be necessary.

Note: For IPv6, it's important to note that the traffic will almost always require manual rules to pass the traffic. Clients are likely to use a link-local source going to a multicast destination. That source is not covered by the default interface rules.

Consult the firewall logs at **Status > System Logs**, on the **Firewall** tab to see if traffic is being blocked. Pay particular attention to the destination address, as it may be different from expected.

Further trouble with game consoles may also be alleviated by switching to manual outbound NAT and enabling Static Port. See *Static Port* for more details.

If the WAN does not have a public address, or if the STUN NAT test fails, the service can fail to add NAT mappings when clients make requests. The routing log will contain messages such as `Failed to add`. The routing log also contains other messages from the daemon at startup indicating these test results.

See also:

For advice on specific consoles and games, see *Configuring pfSense Software for Online Gaming*.

22.10 NTPD

The NTP service is a Network Time Protocol (NTP) daemon which will listen for requests from clients and allow them to synchronize their clock with that of a firewall running pfSense® software. By running a local NTP server and using it for local clients, it reduces the load on the lower-stratum servers and can ensure that local systems can always reach a time server. Before delegating this task to a firewall running pfSense, the best practice is to ensure that the firewall has an accurate clock and keeps time reasonably.

22.10.1 Clock Bootstrap Behavior

The firewall bootstraps its clock at boot in two ways and the firewall performs both of these actions once per boot before it starts the NTP daemon.

- The firewall checks a few commonly modified files on the filesystem and sets the clock to whichever value is latest if it's also later than the clock. This action is taken early in the boot process just after the firewall has mounted its filesystems.

  This ensures that the clock is set to a reasonably close value to current so long as the firewall has been active recently. This works even if there is no network connection available at boot time. The longer it has been since the firewall was last active, the less accurate this method becomes.

- The firewall performs a one-time sync to multiple NTP servers with static IP addresses from Google Public NTP. This avoids a chicken-and-egg problem where the firewall cannot resolve NTP servers because DNSSEC, which is enabled by default, cannot function when the clock is inaccurate. This happens much later in the boot process because it cannot be performed until after the firewall has configured its interfaces and routing.

  This gets the clock as close to accurate as possible without a persistent NTP daemon. There is a hard timeout of 30 seconds in case the upstream servers are unreachable.

Changing Clock Bootstrap Behavior

The NTP clock bootstrap behavior can easily be disabled or changed if an administrator does not want a firewall to contact the default list of servers.

To disable the bootstrap, create the file `/conf/ntp-boot-time-servers` as an empty file. If the file exists and is empty the firewall will skip the initial sync.

To use alternate servers, create the file `/conf/ntp-boot-time-servers` and add in one or more **IP addresses** separated by a single space each. If this file contains a list of space-separated IP addresses, the firewall will use those for the bootstrap sync instead.

22.10.2 NTP Server Configuration

The NTP server is located in the GUI at **Services > NTP**.

NTP Server Settings

The NTP server has the following options:

Interface
Select the interface(s) to use for NTP. The NTP daemon binds to all interfaces by default to receive replies properly. This may be minimized by selecting at least one interface to bind, but that interface will also be used to source the NTP queries sent out to remote servers, not only to serve clients. Deselecting all interfaces is the equivalent of selecting all interfaces.

Time Servers
A list of servers to query in order to keep the clock of this firewall synchronized. This list is initially pulled from the entries under **System > General Setup**. For best results, the best practice is to use at least three servers, but no more than five. Click  **Add** to configured additional time servers.

Prefer
When checked, this NTP server entry is favored by the NTP daemon over others.

No Select
When checked, this NTP server is not used for time synchronization, but only to display statistics.

Orphan Mode
Orphan mode uses the system clock when no other clocks are available, otherwise clients will not receive a response when other servers are unreachable. The value entered here is the stratum used for **Orphan Mode**, and is typically set high enough that live servers are preferred. The default value is `12`.

NTP Graphs
Check to enable RRD graphs for NTP server statistics.

Logging
When logging options are active, NTP logs are written using syslog and may be found under **Status > System Logs**, on the **NTP** tab.

Log Peer Messages
When checked, NTP will log messages about peer events, information, and status.

Log System Messages
When checked, NTP will log messages about system events, information, and status.

Statistics Logging

Click  **Show Advanced** to view these options. When enabled, NTP will create persistent daily log files in `/var/log/ntp` to keep statistics data. The format of the statistics records in the log files can be found in the ntp.conf man page

Log reference clock statistics
When checked, NTP records clock driver statistics on each update.

Log clock discipline statistics
When checked, NTP records loop filter statistics on each update of the local clock.

Log NTP Peer Statistics
When checked, NTP records statistics for all peers of the NTP daemon, along with special signals.

Leap Seconds

Click  **Show Advanced** to view these options. Defines the contents of the Leap Second file, used by NTP to announce upcoming leap seconds to clients. This is typically used only by stratum 1 servers. The exact format of the file may be found on the IETF leap second list

---

Note: IANA updates the leap second data file every 6-12 months and the data file expires on a semi-annual basis in June and December of a given year. Update the content periodically to match the latest available data.

Access Restrictions

Access restrictions (ACLs) are configured on the **ACL** tab under **Services > NTP**. These ACLs control how NTP interacts with clients.

Default Access Restrictions
Control behavior for all clients by default.

Kiss-o'-Death
When set, NTP will send a KoD packet when an access violation occurs. Such packets are rate limited and no more than one per second will be sent.

Modifications
When set, `ntpq` and `ntpdc` queries that attempt to change the configuration of the server are denied, but informational queries are returned.

Queries
When set, all queries from `ntpq` and `ntpdc` are denied.

> **Warning:** Setting this will effectively disable the NTP status page, which relies on `ntpq`.

Service
When set, NTP will deny all packets except queries from `ntpq` and `ntpdc`.

Peer Association
When set, NTP denies packets that would result in a new peer association, including broadcast and symmetric active packets for peers without an existing association.

Trap Service
When set, NTP will not provide mode 6 control message trap service, used for remote event logging.

Custom Access Restrictions

Defines the behavior for specific client addresses or subnets. Click  **Add** to add a new network definition.

Network/mask
The subnet and mask to define the client controlled by the restrictions in this entry.

Restrictions
The option names are abbreviated versions of those in the default list, in the same order.

22.10.3 Serial GPS

If this firewall has an available serial port, a Serial GPS may be used to provide a reference clock for the firewall. If the GPS also generates a Pulse Per Second (PPS) signal, that may also be used as a PPS clock reference.

Warning: USB GPS units may function, they are not desirable time sources due to USB timing issues. The overhead of USB makes its unreliable as a clock or timing source.

For best results, the best practice is to configure at least three NTP servers under **System > General Setup** or **Services > NTP** to avoid loss of sync if the GPS data is not valid over time. Otherwise, the NTP daemon may only use values from the unsynchronized local clock when providing time to clients.

Serial GPS Settings

GPS Type
Select the make and model of the GPS unit. If the model is unknown, use the *Default* choice. If the model is known but not listed, use *Custom*.

Serial Port
All serial ports detected on the firewall are listed. Select the port with the GPS attached. On-board hardware serial ports start with `cuau`, USB serial ports are prefixed with `cuaU`.

Baud Rate
Enter the serial speed for the GPS, typically a low value such as `4800`

NMEA Sentences
By default, NTP will listen for all NMEA sentences it is capable of interpreting. To limit this to specific types, select them from the list.

Fudge Time 1
Specifies a constant to be added to the GPS PPS signal as an offset.

Fudge Time 2
Specifies a constant to be added to the GPS time as an offset.

Stratum
Used to configure the stratum of the GPS clock. The default value is `0` so the GPS is preferred over all others. If another clock must be preferred instead, set the stratum value higher than the stratum of the preferred clock.

Flags
These options provide additional tweaks to fine-tune the GPS behavior:

Prefer this clock
Marks the reference clock as preferred by NTP.

Do not use this clock
Prevents the clock from being used by NTP for time synchronization, it is only displayed for reference.

PPS signal processing
Enables processing of the Pulse Per Second (PPS) signal in the GPS driver. Only enable this if the GPS is known to output a usable PPS signal.

Falling edge PPS signal processing
When set, the falling edge of the PPS signal is used for timing, rather than the rising edge.

Kernel PPS clock discipline
When set, the OS Kernel will use PPS directly for timing.

Obscure location in timestamp
Obscures the GPS data, so the location of the clock cannot be determined.

Log the sub-second fraction of the received time stamp
When checked, this can rapidly fill the log, but can be useful for fine-tuning of **Fudge Time 2**.

Clock ID
A 1-4 character identifier used to change the GPS Clock ID. The default value is `GPS`.

GPS Initialization
Contains the initialization string sent to the GPS at start up to configure its behavior. When using the *Custom* GPS type, a proper initialization string for the GPS must be entered manually.

NMEA Checksum Calculator
Calculates a checksum for use when crafting new **GPS Initialization** values or adjusting existing values.

Configuring a Serial GPS

To configure a GPS for use by NTP:

- Navigate to **Services > NTP**
- Click the **Serial GPS** tab
- Configure the settings
- Click **Save**

22.10.4 PPS Source (Non-GPS)

A non-GPS PPS Source, such as a radio, may also be used for clock timing. It cannot be used for synchronization since there is no time data, but it can be used to ensure a clock ticks accurately.

PPS Source Settings

Serial Port
All serial ports detected on the firewall are listed. Select the port with the GPS attached. On-board hardware serial ports start with `cuau`, USB serial ports are prefixed with `cuaU`.

Fudge Time 1
Specifies a constant to be added to the PPS signal as an offset, to account for delay between the transmitter and receiver.

Stratum
Used to configure the stratum of the PPS source. The default value is `0` so the PPS source is preferred over all others. If another clock must be preferred instead, set the stratum value higher than the stratum of the preferred clock.

Flags

Falling edge PPS signal processing
When set, the falling edge of the PPS signal is used for timing, rather than the rising edge.

Kernel PPS clock discipline
When set, the OS Kernel will use PPS directly for timing.

Record a timestamp
Record a timestamp once for each second, which is useful for constructing Allan deviation plots.

Clock ID
A 1-4 character identifier used to change the PPS Clock ID. The default value is PPS.

Configure a PPS Source

To configure a Non-GPS PPS source:

- Navigate to **Services > NTP**
- Click the **PPS** tab
- Configure the settings
- Click **Save**

See also:

- *NTP Daemon Status*

22.11 Wake on LAN

The Wake on LAN (WOL) page at **Services > Wake on LAN** can wake up computers from a powered-off state by sending special "Magic Packets".

The network interface card in the client computer that is to be woken up must implement WOL, and it must be configured properly. Typically, there is a BIOS setting to enable WOL, and non-integrated adapters often require a WOL cable connected between the NIC and a WOL header on the motherboard.

WOL has many potential uses. Typically, workstations and servers are kept running because of services they provide, files or printers they share, or for convenience. Using WOL would allow these to remain in a sleep state to conserve power. When a service is required, the system can be woken up when needed. Another example would be if someone needs remote access to a system, but the user shut it down before leaving the office. Using WOL the target system can be awoken, and it may then be accessed once it has booted.

Warning: WOL offers no inherent security. Any system on the same layer 2 network may transmit a WOL packet, and the packet will be accepted and obeyed. It is best to only configure WOL in the BIOS for machines that need it, and disable it in all others.

22.11.1 Wake Up a Single Machine

To wake up a single machine:

- Navigate to **Services > Wake on LAN**
- Set the options as follows:

  Interface
  The interface through which the firewall can reach the target host.

  MAC Address
  The **MAC address** of the target host in the format `xx:xx:xx:xx:xx:xx`.

- Click **Send**

pfSense® software will transmit a WOL Magic Packet out the chosen interface, and if everything went as planned, the system will power on and start to boot. Keep in mind that systems will take some time to boot. It may be several minutes before the target system is available.

22.11.2 Storing MAC Addresses

To store a MAC address for convenience:

- Navigate to **Services > Wake on LAN**
- Click  **Add** under the list of stored MAC addresses to add a new entry
- Configure the entry as follows:

  Interface
  The interface through which the firewall can reach the target host.

  MAC Address
  The **MAC address** of the target host in the format `xx:xx:xx:xx:xx:xx`.

  Description
  Text describing the entry, such as the target system hostname, owner, or location. For example: "Pat's PC" or "Sue's Server"

- Click **Save**

Once saved, the entry will be available on the list at **Services > Wake on LAN**.

Maintaining the entries is similar to other tasks in pfSense: Click  to edit an existing entry, and click  to remove an entry.

22.11.3 Wake a Single Stored Machine

To send a WOL Magic Packet to a system that has been previously stored:

- Navigate to **Services > Wake on LAN**
- Locate the desired entry in the list
- Click the **MAC address** or the  icon in the **Actions** column

The WOL page will reload, and the Magic Packet will be sent. The status of the WOL attempt will also be displayed.

22.11.4 Wake All Stored Machines

To send a WOL Magic Packet to all stored systems at once:

- Navigate to **Services > Wake on LAN**
- Click  **Wake All Devices** under the list of stored addresses.

22.11.5 Wake from DHCP Leases View

To send a WOL Magic Packet from the DHCP Leases view:

- Navigate to **Status > DHCP Leases**
- Locate the desired system in the list
- Click  at the end of the lease row to send a WOL Magic Packet

---

Note: The WOL function is only available for systems marked **offline**, meaning they are not in the ARP table on the firewall. If a system was very recently powered off, it can take a few minutes for the ARP entry to expire before it will be marked offline.

If a system has been powered off for quite some time, clicking  **Show all configured leases** might be required to see the previous lease.

When the link is clicked, the browser will return to the WOL page, and the Magic Packet will be sent.

22.11.6 Save from DHCP Leases View

A MAC address and hostname may be copied to a new WOL mapping entry while viewing the DHCP leases.

- Navigate to **Status > DHCP Leases**
- Locate the desired system in the list
- Click  at the end of lease entry
- Confirm the values on the page, and enter any missing information.
- Click **Save**

22.12 PPPoE Server

pfSense® software can act as a PPPoE server, accepting and authenticating connections from PPPoE clients on a local interface, in the role of an access concentrator (LAC). This feature can be used to force users to authenticate before gaining network access, or otherwise control their login behavior.

The PPPoE Server is located at **Services > PPPoE Server**.

22.12.1 PPPoE Server Settings

The PPPoE Server page has several options, which fall into multiple categories.

Server Settings

These options control the general behavior of the PPPoE Server.

Enable
When checked, this PPPoE Server instance will be active.

Interface
The single interface upon which PPPoE service will be available.

Total User Count
Determines how many clients in total are allowed to connect to this instance.

User Max Logins
Determines how many times a single client may log in concurrently.

Server Address
The IP address which the firewall will send to the PPPoE clients to use as their gateway.

> **Warning:** This IP address **must not** be an IP address currently in use on the firewall.

Remote Address Range
The IP address for the start of the PPPoE client subnet. Together with the Subnet Mask it defines the network used by the PPPoE clients.

Subnet Mask
Defines the CIDR mask assigned to PPPoE clients.

Description
Optional explanatory text for this server instance.

DNS Servers
Optional fields used to send specific DNS servers to the PPPoE clients, otherwise the firewall IP address will be sent to the client for DNS if the DNS Forwarder or DNS Resolver are enabled. If the DNS Forwarder and DNS Resolver are both disabled, then the DNS servers configured on the firewall will be sent instead.

RADIUS Settings

These options configure RADIUS authentication for the server.

Use RADIUS Authentication
Check to configure the PPPoE server to use at least one RADIUS server for Authentication instead of local users.

Use RADIUS Accounting
Optional, sends RADIUS accounting data to the RADIUS server to note items such as login and logout times, and bandwidth used.

Use a Backup RADIUS Authentication Server
A second RADIUS server to use if the primary RADIUS server fails.

NAS IP Address
Optional, sends a specific IP address to the RADIUS server for the NAS-IP-Address attribute.

RADIUS Accounting Update
The interval at which accounting data is sent to the RADIUS server, in seconds.

RADIUS Issued IP Addresses
When checked, IP addresses can be assigned to users via RADIUS reply attributes.

Primary RADIUS Server
The preferred RADIUS server to use for Authentication.

IP Address
The IP address of the RADIUS server.

Authentication Port
The port used for authentication (typically `1812`).

Accounting Port
The port used for accounting data (typically `1813`).

Primary RADIUS Server Shared Secret
The shared secret configured for this firewall on the RADIUS server. The same value must be entered in the Confirm box.

Secondary RADIUS Server
Same type of settings as the primary, but defines the secondary RADIUS server.

Users

The user list defines account credentials the server will allow when not using RADIUS authentication.

Username
The username for the user account.

Password
The password for the user account.

IP Address
An optional static IP address to assign the user at login.

22.12.2 PPPoE Server Configuration

Multiple PPPoE servers may be configured on separate interfaces. Each of the available options are covered above.

To begin setting up a PPPoE server:

- Navigate to **Services > PPPoE Server**
- Click  **Add** to add a new server entry
- Configure the PPPoE Server settings
- Choose an authentication source, either RADIUS or manually defined users
  - Configure **RADIUS** if that will be utilized for user authentication
  - Add users to the server to utilize local authentication if not using RADIUS
    - Click  **Add User**
    - Fill in the credentials and settings for the user.
    - Repeat as needed
- Click **Save**

22.13 IGMP Proxy

The Internet Group Management Protocol (IGMP) Proxy provides a means to proxy multicast traffic between network segments.

The IGMP Proxy service can be found at **Services > IGMP Proxy**.

22.13.1 IGMP Proxy Settings

The IGMP Proxy service has the following settings:

Interface
: The interface to be used for this instance

Description
: Optional text to describe this instance

Type
: The type of network interface defined by this instance

  Upstream Interface
  : The outgoing interface which is responsible for communicating to available multicast data **sources**. There can only be **one** upstream interface.

  Downstream Interface
  : The distribution interfaces to the **destination** networks, where multicast **clients** can join groups and receive multicast data. One or more downstream interfaces must be configured.

Note: For a working IGMP Proxy configuration, one upstream and at least one downstream interface must be defined.

Threshold
The TTL threshold for forwarded data on an interface, to prevent looping from occurring. Packets with a TTL lower than the value in this field will be ignored. The default TTL is 1 if the field is left blank.

Networks
A list of CIDR-masked Network entries to control what subnets are allowed to have their multicast data proxied. Click  **Add Network** to enter additional networks.

22.13.2 IGMP Proxy Configuration

To configure the IGMP Proxy:

- Navigate to **Services > IGMP Proxy**
- Click  **Add** to create a new interface instance
- Configure the instance
- Click **Save**

IGMP requires a firewall rule on the **Downstream** side (e.g. *LAN*) to pass its multicast traffic. In the *Advanced Options* of the firewall rule, **Allow packets with IP Options** must be enabled.

The base installation of pfSense® software includes services which add fundamental functionality and flexibility to the firewall. The topics in this chapter discuss services in the base installation that the firewall provides for other hosts on the network. These services include allocating IPv4 and IPv6 addresses via DHCP, DNS resolution and Dynamic DNS, SNMP, UPnP and more. Additional services can also be added with packages.

CHAPTER

TWENTYTHREE

DHCP

Dynamic Host Configuration Protocol (DHCP), allows a device such as pfSense® software to dynamically allocate IP addresses to clients from predefined pools of addresses. DHCP also sends configuration information to clients such as a gateway, DNS servers, domain name, and other useful settings.

There are currently two available DHCP backends: Kea DHCP and ISC DHCP. Kea is more modern and under active development, ISC DHCP is deprecated and will be removed in future versions of pfSense software. The backend can be changed under **System > Advanced**, **Networking** tab (*Server Backend*).

23.1 Kea Settings Tab

When using the Kea DHCP backend (*Server Backend*) there is a **Settings** tab with global options to control DHCP server behavior not specific to a given interface.

Note: There are separate pages for DHCPv4 and DHCPv6, and they operate independently. However, the options on each page are identical otherwise. Differences between DHCPv4 and DHCPv6, if any, are noted when appropriate.

23.1.1 General Settings

DHCP Client DNS Registration with the DNS Resolver

The DNS Registration options control the default Kea behavior for registering DHCP client hostnames with the *DNS Resolver* so that other clients using this firewall for DNS resolution can resolve these hostnames.

This implementation with Kea works with both DHCPv4 and DHCPv6 client lease data.

Note: When using the Kea DHCP daemon, pfSense software dynamically updates these hostnames with the DNS Resolver without restarting the daemon. DNS updates are seamless and not disruptive, unlike the previous implementation with the ISC DHCP daemon.

Kea DNS Registration also respects the domain name configured in DHCP settings for an interface or static mapping. If no domain is set in the DHCP lease, it falls back to checking for a search domain or, as a last resort, the domain name configured on the firewall.

DNS updates are also kept up-to-date between High Availability failover peers.

DNS Registration
: Controls the default DNS Registration behavior on all interfaces with DHCP enabled. When checked,

Kea will automatically register hostnames from DHCP leases on all interfaces with the DNS Resolver by default.

The setting can be overridden on a per-interface basis in either direction. This allows selectively enabling or disabling DNS Registration for specific interfaces as needed.

Early DNS Registration
Controls whether Kea will register hostnames from DHCP static mappings with the DNS Resolver at startup by default. The setting can be overridden on a per-interface basis in either direction.

When unchecked, Kea does not register the hostname until the client requests a DHCP lease.

With this setting enabled a client does not need to have DHCP enabled for its hostname to be available via the DNS Resolver. This was the default behavior of the ISC DHCP daemon when it performed DNS registration.

23.1.2 High Availability

The **High Availability** section of the page configures the daemon for high availability failover with a peer. This operates in "Hot Standby" mode where one node hands out addresses and they both share lease data while monitoring the other peer. If the active primary node becomes unresponsive, the standby node will take over handing out leases to clients. As the two nodes share lease data, failover is seamless and clients can obtain addresses so long as one of the two nodes is online and functional.

Note: Most of these settings will synchronize via XMLRPC configuration synchronization, except for the TLS Transport options which are per-node options and do not synchronize.

The **High Availability** section contains the following options:

Enable
When checked, enables hot-standby high availability (HA) for Kea DHCP services.

Node Role
Chooses the role for this node when performing high availability.

Primary
This node will serve leases for clients when both nodes are online and functioning properly.

Standby
This node will only serve leases when the primary node is down.

Local Name
The name of this device, such as its hostname.

Local Address
The IP address of this node upon which it will listen for failover data from the peer.

Note: This is typically the local IP address on the Sync interface.

The address family of this IP address does not need to match the address family of the DHCP service. Both DHCPv4 and DHCPv6 can exchange HA data with IPv4 or IPv6 peers.

Remote Name
The name of the remote peer, such as its hostname.

Remote Address
The IP address where the remote peer is listening for failover data from this node.

Note: This is typically the IP address of the Sync interface on the peer.

The address family of this IP address does not need to match the address family of the DHCP service. Both DHCPv4 and DHCPv6 can exchange HA data with IPv4 or IPv6 peers.

When these settings synchronize via XMLRPC configuration synchronization, the values are swapped – the secondary node is configured with the opposite settings so it will have appropriate settings from its frame of reference.

Advanced Options

Heartbeat Delay
Specifies a duration in milliseconds between sending the last heartbeat and the next heartbeat. Default value is `10000` (10 seconds).

The heartbeats are sent periodically to gather the status of the partner and to verify whether the partner is still operating.

Max Response Delay
Specifies a duration in milliseconds since the last successful communication with the partner, after which the server assumes that communication with the partner is interrupted. Default value is `60000` (60 seconds).

> **Warning:** This duration should be greater than the heartbeat delay.

Max Ack Delay
Specifies the maximum time in milliseconds for the client to try to communicate with the DHCP server, after which this server assumes that the client failed to communicate with the DHCP server (is "unacked"). Default value is `60000` (60 seconds).

Max Unacked Clients
Specifies how many unacknowledged clients are allowed before this server assumes that the partner is offline and transitions to the `partner-down` state. Default value is `0` (takes over immediately).

> **Warning:** Do not set this higher than the typical number of clients which would make a DHCP request in a short time frame. Otherwise, clients can be left without an address until additional clients also fail to obtain addresses, which could be a significant amount of time on a small or quiet network. See *Secondary Does Not Enter partner-down State* for more details.

Max Rejected Updates
Specifies how many lease updates for distinct clients can fail, due to a conflict between the lease and the partner configuration or state, before the server transitions to the terminated state. Default value is `10`.

These settings synchronize as-is over XMLRPC configuration synchronization.

TLS Transport

Configures optional encryption and certificate-based authentication to protect the DHCP HA data exchanged between failover peers.

These settings **do not** synchronize over XMLRPC configuration synchronization.

Enable
When checked, enables TLS encryption for DHCP HA data requests. This encrypts the lease data and heartbeat traffic exchanged between the peers.

Server Certificate
The certificate to use for the local DHCP HA server.

Mutual TLS
When checked, configures the Kea HA client to send a client certificate which allows the peer to verify that it is communicating with the expected client.

This ensures that only authorized peers can communicate with the DHCP HA services.

Client Certificate
The client certificate which identifies this node.

23.1.3 Custom Configuration

The KEA GUI allows administrators to configure Kea directly with JSON-formatted configuration snippets. The Kea GUI in pfSense software does not implement some features or options which are possible directly in Kea. As the Kea GUI develops further, the GUI will be capable of configuring more of these features. These snippets allow administrators more control in the meantime.

Tip: These snippets can be used to add functionality such as custom DHCP options, dynamic DNS registration, and fine-tune lease behavior.

These fields are present on each Kea page and the system adds the contents of each configuration snippet to a different section of the Kea configuration file based on the Kea configuration page:

Table 1: Kea Configuration Settings Sections by Page

Kea GUI Page	Kea Configuration Section
DHCPv4 Settings	`Dhcp4`
DHCPv6 Settings	`Dhcp6`
DHCPv4/v6 interfaces	`subnet`
DHCPv4/v6 Pools	`pool`
DHCPv4/v6 static mappings	`reservation`

Warning: These configuration fields must contain well-formed JSON objects and must not include the name of the section itself.

The system adds these snippets to the Kea configuration after the settings from the GUI for each section. So not only can these snippets enable new features, but they also allow administrators to modify or override the base configuration.

These configuration snippets are held in `config.xml`, so they are backed up and restored along with other Kea settings.

The system tests the configuration before attempting to start Kea. If the configuration test fails with the custom snippets, the system files a notice alerting administrators to the problem, then it attempts to start Kea without including the custom configuration snippets.

See also:

See the following forum thread for configuration snippet examples and discussion: https://forum.netgate.com/topic/196513/adding-custom-configuration-in-kea-dhcp-server-with-pfsense-25-03

23.2 Using DHCP Search Domains on Windows DHCP Clients

The DNS Search Domain functionality present in the DHCP Server settings in pfSense® software is only implemented in some DHCP clients; pfSense software uses the standard DHCP option 119 mechanism to deliver the search domains to clients which request them.

Microsoft finally added the ability to utilize option 119 in Windows 10 version 1803 released in April 2018. Older versions of the Microsoft Windows DHCP client **do not** request option 119, so no matter which DHCP server is used, clients running older versions of Microsoft Windows can never receive or use a search domain list from DHCP. Upgrade clients to a current Windows release to use this functionality.

Tip: If older clients must use these settings, and they cannot be upgraded, the setting can be pushed via GPO instead of DHCP.

Sources:

- https://docs.microsoft.com/en-us/windows-server/networking/technologies/dhcp/what-s-new-in-dhcp#new-dhcp-client-side-features-in-the-windows-10-april-2018-update
- http://social.technet.microsoft.com/Forums/en-US/winserverNIS/thread/9ba77f86-4708-42ca-a193-2a01b813ec27/
- http://social.technet.microsoft.com/Forums/en-US/winserverNIS/thread/7ba59619-3484-43fa-8585-a2d69ccd00df/
- http://technet.microsoft.com/en-us/library/dd572752%28v=office.13%29.aspx (See comments)
- http://serverfault.com/questions/37417/which-dhcp-client-os-support-dhcp-option-119-domain-suffix-search

23.3 Static Mappings Inside DHCP Pools

While the ISC DHCP daemon will allow a static mapping to be defined inside the DHCP range/pool in its configuration, doing so can result in unexpected behavior.

A static mapping entry in the ISC DHCP daemon **is not** a reservation, and it **does not** remove that IP address from the pool. The daemon only checks via ICMP ping to ensure that an IP address is not actively in use when making assignments. The static mapping only represents a *preference* for IP address assignment, and it **does not** prevent the daemon from assigning the IP address to other client devices when it is not actively in use by the intended device defined in the static mapping.

Example: The DHCP configuration contains a pool from `192.168.0.10` to `192.168.0.250` with a static mapping defined for `192.168.0.25`. If the device which normally has `192.168.0.25` is ever offline another device could be assigned `192.168.0.25` in its absence. When the original client device reconnects it will not be able to get `192.168.0.25` because it is currently in use, and will instead receive another random address from the pool.

Due to this behavior the best practice is to **only** make assignments **outside** the range/pool and the GUI enforces this practice.

If a use case requires assignments inside the pool and the administrators do not care about the risks involved and want to do so anyway, the input validation check can be removed from the PHP file that drives the DHCP editor page. The details of this change will not be covered in documentation.

See also:

- *DHCPv4 Status*
- *DHCPv6 Status*
- *DHCP Logs*
- *Troubleshooting DHCPv6 Client XID Mismatches*
- *Troubleshooting Offline DHCP Leases*

CHAPTER

TWENTYFOUR

DNS

DNS, or Domain Name System, is the mechanism by which a network device resolves a name like `www.example.com` to an IP address such as `198.51.100.25`, or vice versa. Clients must have functional DNS if they are to reach other devices such as servers using their hostnames or fully qualified domain names.

24.1 DNS Resolver/Forwarder

These topics cover using pfSense® software to handle DNS requests from local clients as either a caching DNS resolver or forwarder. When acting as a resolver or forwarder, pfSense software will perform DNS resolution directly or hand off queries to an upstream DNS forwarding server.

24.1.1 DNS Resolution Process

Every DNS query must be resolved. Depending on which DNS service is in use on the firewall and its configuration, this resolution may happen locally, or it may happen on an upstream forwarding server. The *DNS Resolver* can act in either a resolver or forwarder role, while the *DNS Forwarder* can only act as a forwarder.

Note: Some of the concepts and processes in this document are simplified to make them easier to understand.

Terms

The terms used in this section have specific meanings, though some terms are frequently used interchangeably or in ambiguous ways. The following definitions cover roles involved in DNS resolution and how they process queries.

Client
: A device that wants to translate a hostname into an IP address. This could be a PC, handheld device, a server, etc.

  A client will typically be configured with one or more forwarding DNS servers to which it will send DNS queries.

DNS Server
: A server involved in handling DNS queries. The term "DNS server" is ambiguous because a server involved in DNS can act in one or more specific roles which differ significantly. Used generally, this may refer to servers handling unknown or multiple types of DNS queries.

  The term "DNS Server" can also vary based on context. When talking about clients it usually refers to upstream forwarders or resolvers. When talking about resolvers or DNS infrastructure it usually refers to authoritative DNS servers. Where possible, it is best to use a more specific term.

Forwarder
A forwarder is a type of DNS server which accepts recursive queries from clients and makes its own queries to an upstream DNS server, which could be another forwarder or a resolver. A forwarder will typically cache results, so clients get faster responses for frequently resolved queries.

Note: A forwarder does not perform DNS resolution itself, it passes a query along to another forwarder or resolver unless it has the answer in its cache.

Resolver
A DNS server acting as a resolver will accept recursive queries from clients and/or forwarders and perform DNS resolution by making iterative queries to root and authoritative DNS servers in search of an answer. Like a forwarder it will typically cache results for increased performance.

A resolver maintains a list of root servers in its "root hints" list. Resolvers will typically ship with a stock list but update it periodically from a trusted source.

Root Server
Authoritative name servers which serve queries for the DNS root zone (`.`). Though it is held as a short list of servers (`a.root-servers.net` through `m.root-servers.net`) there are hundreds of servers worldwide handling the requests.

These servers are a starting point for queries and will direct resolvers to authoritative DNS servers for zones which can answer their queries.

Glue Records
Entries in domain name registries which associate specific authoritative DNS servers with a domain name. These are used by the DNS servers for a TLD to determine which servers are authorized to answer queries for a given domain name.

Authoritative Server
A name server which has direct knowledge of hosts in a zone and the proper responses for queries. This could be for one or more domain names, subdomains, and so on. It could be run by a company operating a domain directly, a dedicated DNS provider, or a hosting company on their behalf.

An authoritative server will not make queries of its own to locate an answer. Instead, it will either respond with a direct answer to the query or point the resolver to another authoritative server.

Recursive Query
A query for which the answer is not local to the forwarder or resolver. A forwarder or resolver will in turn make a query to another server to obtain an answer.

It is recursive in that if a forwarding server does not know the answer the *forwarding server* makes a recursive query to another upstream forwarder or resolver to obtain the answer, which is then passed back to a client. The client and every forwarder involved each only make one query upstream until the query eventually reaches a resolver.

Recursive queries can be handled by forwarders or resolvers.

Iterative Query
A query where a resolver asks a root or authoritative server for either a direct answer or for information about how to find the answer.

If the response points the resolver to another authoritative server, the *resolver* will send an additional query to that server. This process is repeated until the resolver obtains a final answer from an authoritative server. The answer is then passed back to the host which queried the resolver.

This process is iterative, rather than recursive, as each query is performed by the resolver and the *resolver* takes further action depending upon the result. The authoritative servers do not make queries

of their own, they only give answers they know locally which may be the answer to a query or a pointer to another source.

DNS Resolution Steps

This is a somewhat simplified version of the process that takes place when a client attempts to resolve a hostname.

- The client checks its own DNS configuration to see how it should handle queries. This typically involves a local `hosts` file and a list of servers which it will contact to resolve queries.

  Note: Some operating systems use additional methods for resolving names which are not a part of this process, such as mDNS or NBNS. These methods are omitted from this document but may occur before clients contact remote DNS servers. Check the operating system documentation for details.

  Local DNS configuration on the client may also include a default domain name which it will append to hostnames to form a fully qualified domain name (FQDN) for a DNS query.

- The client looks in its own `hosts` file to see if a local answer exists. If it does, the answer is used. If it does not, it proceeds to the next step.
- The client contacts the first server in its list and sends a query. If it does not receive a response, it will query the next server in the list.

  The client does not need to know or care if this server is a forwarder or resolver. Either way it gets an answer to its query.
- If the server contacted by the client is a forwarder, the forwarder will make a query to one of its own configured upstream DNS servers.

  As with a client these upstream servers could be additional forwarders or resolvers. If the upstream server is a forwarder, this step is repeated recursively until the query reaches a resolver.
- Once the query reaches a resolver, the resolver attempts to track down the answer by making iterative queries to servers in an attempt to track down a definitive answer.
- The resolver consults its list of root DNS servers in the hints file and contacts one to locate information on how to proceed.
- The resolver asks a root DNS server for information about the top level domain (TLD) in the requested FQDN (e.g. `.com`).
- The root DNS server returns a list of authoritative servers which have information about the TLD.
- The resolver queries one of the TLD servers from the response to ask about the domain name.
- The TLD servers consult glue records under the TLD to locate one or more authoritative DNS servers which can answer queries for the domain.
- The TLD server returns the list of authoritative DNS servers to the resolver.
- The resolver contacts one of the authoritative DNS servers for a domain and asks for a response to the original query.

  If the query is for an FQDN in a subdomain this may result in more steps as the authoritative server may tell the resolver to contact a different authoritative server which has knowledge of the subdomain. This process will be repeated until the resolver reaches the authoritative server with the answer to the query.
- The authoritative DNS server for the domain or subdomain responds to the resolver with the answer to the query.
- The resolver passes the response back to the host which submitted the request. The resolver may opt to cache the response.

- If the request came from a forwarder, it is passed back again to the host which submitted the request. The forwarder may opt to cache the response.

  If multiple forwarding servers were involved in the query this step is repeated until it reaches the original source of the query.

- The response to the query arrives back at the client.

All of this can happen quite fast. Even with multiple servers involved the entire process can be completed in fractions of a second.

24.1.2 DNS Rebinding Protections

pfSense® software includes built in methods of protection against DNS rebinding attacks.

A DNS rebinding attack is when someone with control over DNS responses for a domain feeds a client an address on the local network of the client – or even the client computer itself – as a response for a hostname in the domain controlled by the attacker. This would happen when the client requests a page in the malicious domain. Because the server run by the attacker and the hostname pointing to the client network are in the same domain from the perspective of the browser, the browser may allow scripts from the malicious server to run and access the other host. This can trigger the client to unintentionally exploit a device that would otherwise be unreachable from the Internet directly.

DNS rebinding attack protection is active by default. This behavior is controlled by the **DNS Rebind Check** option under **System > Advanced**, **Admin Access** tab.

DNS protection

When active, this protection causes the DNS resolver and forwarder to strip addresses from DNS responses for local and private IP addresses which should not normally be received from public DNS servers.

Tip: This is the safest and best practice as responses to DNS queries made through *public* DNS servers should never include *private* IP addresses.

For a list of addresses including in this protection, see the following table:

Table 1: Addresses included in DNS Rebinding Protection

Address	Description
127.0.0.0/8	RFC 1122 Loopback Addresses (Localhost)
10.0.0.0/8	RFC 1918 Private Addresses
::ffff:a00:0/104	IPv6 Representation of 10.0.0.0/8
172.16.0.0/12	RFC 1918 Private Addresses
::ffff:ac10:0/108	IPv6 Representation of 172.16.0.0/12
192.168.0.0/16	RFC 1918 Private Addresses
::ffff:a9fe:0/112	IPv6 Representation of 192.168.0.0/16
169.254.0.0/16	RFC 3927 IPv4 Link Local Addresses
::ffff:c0a8:0/112	IPv6 Representation of 169.254.0.0/16
fd00::/8	RFC 4193 IPv6 Unique Local Unicast Addresses (ULA)
fe80::/10	RFC 4291 IPv6 Link Local Addresses

There are some cases when public DNS servers give responses containing private IP addresses in replies. This may be the case for private internal hostnames under domains owned by an organization that does not use split DNS. It is

also common in DNS-based block lists such as those for e-mail spam prevention (DNSBL, RBL, etc.). In these cases overrides can be set for individual domains. The exact method depends on which DNS service is active.

Note: This behavior is automatically overridden for domains in the DNS Resolver and DNS Forwarder domain override lists as the most common usage of that functionality is to resolve internal DNS hostnames.

DNS Resolver

When DNS rebinding attack protection is active the *DNS Resolver* strips private addresses from DNS responses. Additionally, the DNSSEC validator may mark the answers as bogus. This is handled automatically using a list of `private-address` directives maintained by the firewall.

To exclude a domain from DNS rebinding protection, use the **Custom Options** box in the DNS resolver settings. Enter one domain per line in the following format, preceded by the `server:` line.

```
server:
private-domain: "example.com"
private-domain: "dnsbl.example"
```

DNS forwarder

The *DNS Forwarder* uses the option `--stop-dns-rebind` by default, which rejects and logs addresses from upstream name servers which are in private address ranges.

To exclude a domain from DNS rebinding protection, use the DNS forwarder **Advanced Settings** box as follows:

```
rebind-domain-ok=/example.com/
rebind-domain-ok=/dnsbl.example/
```

Additionally, it is possible to exclude the loopback range (`127.0.0.0/8`) from protection using the DNS forwarder **Advanced Settings** box as follows:

```
rebind-localhost-ok
```

Note: Rather than exclude the entire loopback range, it's generally better to allow such responses on a per-domain basis instead.

GUI protection

For those not using the DNS resolver or forwarder, and as an additional layer of checks, the GUI will block access attempts using unknown hostnames. In this case the GUI will deny access and display "Potential DNS Rebind Attack Detected".

By default, the GUI only accepts the hostname and domain configured under **System > General Setup**. For instance if `firewall.example.com` is configured as the firewall hostname, and the GUI is loaded in a browser using `fw1.example.com`, the GUI will reject that attempt. Define additional hostnames under **System > Advanced**, **Admin Access** tab in the **Alternate Hostnames** field.

Tip: If a user encounters this error they can log into the GUI using the IP address of the firewall rather than the hostname.

If a client encounters this message when attempting to access a forwarded service (Port forward, 1:1 NAT, etc.) it indicates that the request did not match any NAT rules. From the inside of the network, this would require NAT reflection or split DNS to accomplish.

See also:

Accessing Port Forwards from Local Networks

24.1.3 Creating Wildcard Records in DNS Forwarder/Resolver

A wildcard DNS record resolves `<anything>.example.com` to a single IP address, which can be useful in certain cases.

DNS Resolver (Unbound)

To create a wildcard entry the DNS Resolver (Unbound), use the following directives in the custom options box:

```
server:
local-zone: "example.com" redirect
local-data: "example.com 86400 IN A 192.168.1.54"
```

That makes any host under `example.com` resolve to `192.168.1.54`. For example, `www.example.com`, `thissitedoesnotexist.example.com`, `mystuff.example.com`, and so on.

If there are existing **Host Override** or **Domain Override** entries for the same domain, these custom options may not function as expected. When overrides are present, the zone will already be defined but with a different zone type set. For domains associated with host overrides, the default behavior of the local zones can be altered with the **System Domain Local Zone Type** setting in the *DNS Resolver Configuration*.

DNS Forwarder (dnsmasq)

To create a wildcard entry in the DNS Forwarder, use the following directives in the advanced options:

```
address=/example.com/192.168.1.54
```

If a specific host override is set for example:

```
specific.example.com 192.168.1.100
knownhost.example.com 192.168.1.101
```

Then those would be returned when doing a query for those hosts, only when no specific host has been specified in the host overrides would the advanced wildcard entry be used.

To resolve the domain to an IP address:

```
example.com 192.168.1.45
```

Leave the host field blank in the host overrides. So if the query is now for `example.com` the forwarder will return `192.168.1.45`. If a client requests `knownhost.example.com` then `192.168.1.101` would be returned instead.

If a blank hostname `example.com` host override entry has not been created, then a query for `example.com` would return the wildcard IP address set in the advanced option.

If a client queries for `madeupname.example.com` then since no specific host record for `madeupname` exists in the host overrides. The forwarder will return the wildcard entry of `192.168.1.54`.

See also:

- *DNS Lookup*
- *Interface and DNS Configuration*
- *Troubleshooting DNS Resolution Issues*
- *Troubleshooting the DNS Cache*
- *Troubleshooting DNS Queries*

24.2 DNS Guides

How to perform various tasks related to DNS.

- *Configuring DNS over TLS*
- *Blocking External Client DNS Queries*
- *Redirecting Client DNS Requests*

24.3 Dynamic DNS

Dynamic DNS updates an external DNS server with an interface IP address when it changes. This enables a firewall with a dynamic WAN such as DHCP or PPPoE to host public services even when its IP address changes periodically.

CHAPTER

TWENTYFIVE

TRAFFIC SHAPER

25.1 What the Traffic Shaper can do for a Network

The basic idea of traffic shaping is raising and lowering the priorities of packets or keeping them under a certain speed. This concept seems simple, however, the number of ways in which this concept can be applied is vast. These are but a few common examples that have proven popular with users of pfSense® software.

25.1.1 Keep Browsing Smooth

Asymmetric links, where the download speed differs from the upload speed, are commonplace, especially with DSL. Some links are so out of balance that the maximum download speed is almost unattainable because it is difficult for a firewall to send out enough ACK (acknowledgement) packets to keep traffic flowing. ACK packets are transmitted back to the sender by the receiving host to indicate that data was successfully received, and to signal that it is OK to send more. If the sender does not receive ACKs in a timely manner, congestion control mechanisms in TCP will kick in and slow down the connection.

This type of situation is common: When uploading a file over a link that has asymmetric throughput capability, browsing and downloading slows to a crawl or stalls. This happens because the uploading portion of the circuit is full from the file upload and there is little room to send ACK packets which allow downloads keep flowing. By using the shaper to prioritize ACK packets, the firewall can enable faster, more stable download speeds on asymmetric links.

This is not as important on symmetric links where the upload and download speed are the same, but may still be desirable if the available outgoing bandwidth is heavily utilized.

25.1.2 Keep VoIP Calls Clear

If Voice over IP calls use the same circuit as data, then uploads and downloads may degrade call quality. pfSense software can prioritize the call traffic above other protocols, and ensure that the calls make it through clearly without breaking up, even while streaming hi-def video from Netflix at the same time. Instead of the call breaking up, the shaper reduces speed of the other transfers to leave room for the calls.

25.1.3 Reduce Gaming Lag

The shaper also has options to give priority to the traffic associated with network gaming. Similar to prioritizing VoIP calls, the effect is that even if users on the network are downloading while playing, the response time of the game should still be nearly as fast as if the rest of the connection were idle.

25.1.4 Keep P2P Applications In Check

By lowering the priority of traffic associated with known peer-to-peer ports, administrators can rest easier knowing that even if those programs are in use, they won't hinder other traffic on the network. Due to its lower priority, other protocols will be favored over P2P traffic, which will be limited when any other services need the bandwidth.

25.1.5 Enforce Bandwidth Limits

Limiters can apply a bandwidth limit to a group of devices, such as all traffic on an interface, or masking on limiters can apply them on a per-IP address or per-network basis. This way the firewall can ensure that no one person can consume all available bandwidth.

25.2 Hardware Limitations

One mechanism pfSense® software can use for traffic shaping is **ALTQ**. Unfortunately, only a subset of all network cards are compatible with ALTQ traffic shaping because it must be implemented in the hardware drivers.

The following network cards are compatible with ALTQ traffic shaping:

> `ae(4)`, `age(4)`, `alc(4)`, `ale(4)`, `an(4)`, `aue(4)`, `axe(4)`, `bce(4)`, `bfe(4)`, `bge(4)`, `bnxt(4)`, `bridge(4)`, `bxe(4)`, `cas(4)`, `cpsw(4)`, `dc(4)`, `de(4)`, `ed(4)`, `em(4)`, `ep(4)`, `epair(4)`, `et(4)`, `fxp(4)`, `gem(4)`, `hme(4)`, `hn(4)`, `igb(4)`, `igc(4)`, `ix(4)`, `ixv(4)`, `jme(4)`, `l2tp(4)`, `le(4)`, `lem(4)`, `msk(4)`, `mxge(4)`, `my(4)`, `ndis(4)`, `nfe(4)`, `ng(4)`, `nge(4)`, `npe(4)`, `nve(4)`, `ql(4)`, `ovpnc(4)`, `ovpns(4)`, `ppp(4)`, `pppoe(4)`, `pptp(4)`, `re(4)`, `rl(4)`, `sf(4)`, `sge(4)`, `sis(4)`, `sk(4)`, `ste(4)`, `stge(4)`, `ti(4)`, `tun(4)`, `txp(4)`, `udav(4)`, `ural(4)`, `vge(4)`, `vlan(4)`, `vmx(4)`, `vr(4)`, `vte(4)`, `vtnet(4)`, `wlan(4)`, `xl(4)`

Note: This list is based on the contents of the `is_altq_capable()` function in `interfaces.inc`. If a driver is not in the list above, it is possible that it was added to the source in a later version. Check the source on Github for the most accurate and up-to-date list of drivers compatible with ALTQ.

Another type of traffic shaping on pfSense software is **Limiters**. Limiters use a different backend, operating through `dummynet` pipes and not **ALTQ**. Limiters do not have the same limitation as ALTQ, any network card is capable of using Limiters.

25.3 Network Interface Drivers with ALTQ Traffic Shaping Compatibility

This section provides information regarding drivers in FreeBSD compatible with ALTQ traffic shaping.

25.3.1 Information

The ALTQ framework is used for queuing/traffic shaping. In pfSense® software, this is utilized by the Shaper Wizard and the Queues/Interfaces tabs under **Firewall > Traffic Shaper**.

See the altq(4) manual page or the altq(9) manual page.

On those pages, select the *version of FreeBSD that corresponds to the pfSense software version being run*.

In addition to the drivers listed as being compatible with ALTQ in FreeBSD, pfSense software also includes compatibility for ALTQ on vlan(4) and IPsec enc(4) interfaces.

If a driver is not compatible with ALTQ, traffic shaping can be performed by *Limiters* instead.

25.4 ALTQ Scheduler Types

pfSense® software contains several ALTQ scheduler types to cover a large range of shaping scenarios. The options for ALTQ are:

Priority Queuing (PRIQ)
: Manages prioritization of connections

Class-Based Queuing (CBQ)
: Performs bandwidth sharing between queues and bandwidth limits

Hierarchical Fair Service Curve (HFSC)
: Enforces real-time bandwidth guarantees along with a hierarchical tree of nested queues.

Controlled Delay (CoDel)
: Attempts to combat bufferbloat.

Fair Queuing (FAIRQ)
: Attempts to fairly distribute bandwidth among all connections.

PRIQ, CBQ, and HFSC are selectable in the shaper wizards and the wizards will show the proper options and create the queues based on the chosen ALTQ discipline.

25.4.1 Performance Caveats

Enabling ALTQ traffic shaping places an extra burden on the hardware, and there will be an overall potential network performance loss. On systems that have horsepower to spare, this may not be noticeable. On systems that operate close to their specification limits the firewall may see a degradation of performance. Whether the loss is worse than working without shaping depends on the individual workload.

25.4.2 Priority Queuing (PRIQ)

PRIQ is one of the easiest disciplines to configure and understand. The queues are all directly under the root queue, there is no structure to have queues under other queues with PRIQ as there is with HFSC and CBQ. It does not care about bandwidth on interfaces, only the priority of the queues. The values for priority go from `0` to `15`, and the higher the priority number, the more likely the queue is to have its packets processed.

PRIQ can be harsh to lesser queues, starving them when the higher priority queues need the bandwidth. In extreme cases, it is possible for a lower priority queue to have little or no packets handled if the higher priority queues are consuming all available resources.

25.4.3 Hierarchical Fair Service Curve (HFSC)

The HFSC traffic shaping discipline is very powerful. It is useful for services such as VoIP and video to deliver a minimum guaranteed amount of bandwidth.

Queues in HFSC are arranged in a hierarchy, or a tree, with root queues for each interface, parent queues underneath, and child queues nested under the parent queues (etc.). Each queue can have a set bandwidth and related options.

HFSC Queue Options

HFSC has a few queue options that are not implemented by other disciplines. It is through these options that it achieves guaranteed real-time processing and link sharing.

The Service Curve (sc) is where bandwidth requirements for this queue are tuned.

m1
: Burstable bandwidth limit

d
: Time limit for bandwidth burst, specified in milliseconds. (e.g. 1000 = 1 second)

m2
: Normal bandwidth limit

For example, a connection needs **m1** bandwidth within **d** time, but a normal maximum of **m2**. Within the initial time set by **d**, **m2** is not checked, only **m1**. After **d** has expired, if the traffic is still above **m2**, it will be shaped. Most commonly, **m1** and **d** are left blank, so that only **m2** is checked.

Each of these values may be set for the following uses:

Upper Limit
: Maximum bandwidth allowed for the queue. Will do hard bandwidth limiting. The **m1** parameter here can also be used to limit bursting. In the time frame **d** a connection will not get more than **m1** bandwidth.

Real Time
: Minimum bandwidth guarantee for the queue. This is only valid for child queues. The **m1** parameter will always be satisfied in time frame **d**, and **m2** is the maximum that this discipline will allow to be used. Note The value for **m2** cannot exceed 30% of the available bandwidth from the parent queue.

Link Share
: The bandwidth share of a backlogged queue. Will share bandwidth between classes if the **Real Time** guarantees have been satisfied. The **m2** value for **Link Share** will override the **Bandwidth** setting for the queue. These two settings are the same, but if both are set, **m2** from **Link Share** is used.

By combining these factors, a queue will get the bandwidth specified by the **Real Time** factors, plus those from **Link Share**, up to a maximum of **Upper Limit**. It can take a lot of trial and error, and perhaps a lot of arithmetic, but it

may be worth it to ensure that network traffic is governed properly. For more information on **m1**, **d**, and **m2** values for different scenarios, visit the Traffic Shaping forum category.

25.4.4 Class-Based Queuing (CBQ)

Class-Based Queuing, or CBQ, is similar to HFSC in that is can have a tree of queues nested under other queues. It enforces bandwidth limits (not guarantees like HFSC), manages priorities for queues, and it has the ability to allow queues to borrow bandwidth from their parent. Because of the simpler queue configuration, it can be a good alternative to HFSC especially if the firewall does not need to guarantee minimum bandwidths.

With CBQ, queue priorities range from `0` to `7` with higher numbers indicating higher priority. Queues of an equal priority are processed in a round-robin fashion.

Note: Though child queues can borrow from their parent queue, the sum of the bandwidth of the child queues cannot exceed the bandwidth of the parent. Therefore, CBQ is not an alternative to limiters for individual (e.g. per-IP address) bandwidth limits.

CBQ Queue Options

The CBQ discipline implements the concept of *borrow*, meaning that if the **Borrow from other queues when available** checkbox on the queue is enabled, then the queue will be able to borrow other available bandwidth from its parent queue. This will only allow a child queue to obtain up to the bandwidth of its *immediate* parent, if available, it will not borrow from other parent queues.

25.4.5 CoDel Active Queue Management

The CoDel Active Queue Management (AQM) discipline is short for Controlled Delay and is pronounced “coddle”. It was designed to combat problems associated with bufferbloat in networking infrastructure.

See also:

Bufferbloat is described in detail at http://www.bufferbloat.net/projects/bloat/wiki/Introduction.

Put simply, traffic can pile up and go in chunks rather than a smooth stream due to the size of buffers in network equipment. By controlling the delay of the traffic this effect can be lessened.

CoDel has no specific configuration controls or options. When activated for a queue, it will automatically attempt to manage traffic as described in the CoDel wiki at http://www.bufferbloat.net/projects/codel/wiki. It attempts to keep traffic delays low but does permit bursting, it controls delays, but it does not pay attention to round-trip delay, load, or link speed, and it can automatically adjust if the link speed changes.

The target for CoDel is mid-range networking. It does not work well at very low bandwidth (1Mbit/s or less) and it does not gracefully handle large numbers of simultaneous flows or datacenter-grade traffic loads.

CoDel is not configurable using the wizard, but it does not require complex setup:

- Navigate to **Firewall > Traffic Shaper**, **By Interface** tab
- Select an interface (e.g. **WAN**)
- Set the **Scheduler Type** to *CODEL*
- Set an appropriate value for **Bandwidth**
- Click **Save**

- Repeat as needed for all other active WAN-type interface(s)

25.4.6 Fair Queuing (FAIRQ)

In FAIRQ, queues are monitored from highest to lowest priority, but the scheduler attempts to fairly distribute bandwidth among all connections.

When there is no contention for bandwidth, FAIRQ will send all waiting packets. When there is contention for bandwidth FAIRQ will start looking for queues that are not exceeding their limits, first starting with high priority queues and working toward lower queues. A packet in a full high priority queue is processed *after* a packet from a lower priority queue which is not full. If all queues are full, then FAIRQ will send a packet from the highest priority queue.

FAIRQ allows connections to exceed queue bandwidth, but will maintain an average consumption equal to the defined queue bandwidth.

The traffic shaper wizard cannot configure FAIRQ, it requires a manual configuration.

25.5 Advanced Customization

The rules and queues generated by the shaper wizard may not be an exact fit for a network. Network devices may use services that need shaped which are not listed in the wizard, games that use different ports, or other protocols that need limiting.

After the basic rules have been created by the wizard, it is relatively easy to edit or copy those rules to make adjustments for other protocols.

25.5.1 Editing Shaper Queues

Queues are where bandwidth and priorities are allocated by the shaper. Each queue has settings specific to the scheduler that was chosen in the wizard (*ALTQ Scheduler Types*). Queues can also be assigned other attributes that control how they behave. Queues may be managed at **Firewall > Traffic Shaper**. Click on a queue name in the list or tree shown on the **By Interface** or **By Queue** tabs, as seen in Figure *Traffic Shaper Queues List*

> **Warning:** Creating or editing queues is for advanced users only. It is a complex task with powerful results, but without thorough understanding of the settings involved the best practice is to stick with queues generated by the wizard rather than trying to make new queues.

To edit a queue, click its name in the list/tree.

To delete a queue, click it once to edit the queue, then click  **Delete This Queue**. Do not delete a queue if it is still being referenced by a firewall rule.

To add a new queue, click the interface or parent queue under which the new queue will be placed, and then click  **Add New Queue**.

When editing a queue, each of the options must be carefully considered. For more information about these settings than is mentioned here, visit the PF Packet Queuing and Prioritization FAQ or read *The OpenBSD PF Packet Filter* book.

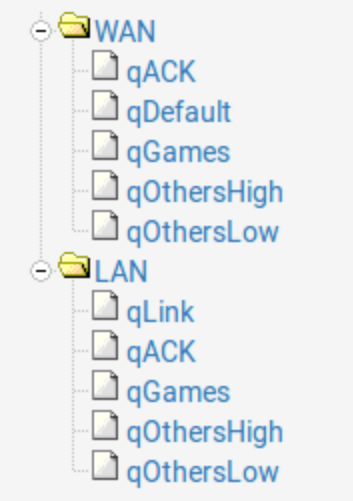


Fig. 1: Traffic Shaper Queues List

Name
The queue name must be between 1-15 characters and cannot contain spaces. The most common convention is to start the name of a queue with the letter "q" so that it may be more readily identified in the ruleset.

Priority
The priority of the queue. Can be any number from 0-7 for CBQ and 0-15 for PRIQ. Though HFSC can use priorities, the current code does not honor them when performing shaping. Queues with higher numbers are preferred by the shaper when there is an overload, so situate queues accordingly. For example, VoIP traffic is the highest priority, so it would be set to a `7` on CBQ or `15` on PRIQ. Peer-to-peer network traffic, which can be delayed in favor of other protocols, would be set at `1`.

Bandwidth (root queues)
The amount of bandwidth available on this interface in the **outbound** direction. For example, WAN-type interface root queues list upload speed. LAN type interfaces list the sum total of all WAN interface download bandwidth.

Queue Limit
The number of packets that can be held in a queue waiting to be transmitted by the shaper. The default size is `50`.

Scheduler Options
There are five different Scheduler Options that may be set for a given queue:

Default Queue
Selects this queue as the default, the one which will handle all unmatched packets on an interface. Each interface must have one and only one default queue.

Random Early Detection (RED)
A method to avoid congestion on a link. When set, the shaper will actively attempt to ensure that the queue does not get full. If the bandwidth is above the maximum given for the queue, drops will occur. Also, drops may occur if the average queue size approaches the maximum. Dropped packets are chosen at random, so connections using more bandwidth are more likely to see drops. The net effect is that the bandwidth is limited fairly, encouraging a balance. RED should only be used with TCP connections since TCP is capable of handling lost packets, and hosts can resend TCP packets when needed.

Random Early Detection In and Out (RIO)
Enables RED with in/out, which results in having queue averages being maintained and checked against incoming and outgoing packets.

Explicit Congestion Notification (ECN)
Along with RED, it allows sending of control messages that will throttle connections if both ends respect ECN. Instead of dropping the packets as RED will normally do, it will set a flag in the packet indicating network congestion. If the other side sees and obeys the flag, the speed of the ongoing transfer will be reduced.

Codel Active Queue
A flag to mark this queue as being the active queue for the Codel shaper discipline.

Description
Optional text describing the purpose of the queue.

Bandwidth (Service Curve/Scheduler)
The Bandwidth setting should be a fraction of the available bandwidth in the parent queue, but it must also be set with an awareness of the other neighboring queues. When using percentages, the total of all queues under a given parent cannot exceed 100%. When using absolute limits, the totals cannot exceed the bandwidth available in the parent queue.

Scheduler-specific Options
Next are scheduler-specific options. They change depending on whether a queue is using HFSC, CBQ, or PRIQ. They are all described in *ALTQ Scheduler Types*.

Click **Save** to save the queue settings and return to the queue list, then click **Apply Changes** to reload the queues and activate the changes.

25.5.2 Editing Shaper Rules

Traffic shaping rules control how traffic is assigned into queues. If a new connection matches a traffic shaper rule, the firewall will assign packets for that connection into the queue specified by that rule.

Packet matching is handled by firewall rules, notably on the **Floating** tab. To edit the shaper rules:

- Navigate to **Firewall > Rules**
- Click the **Floating** Tab
- Find the rule to edit in the list, as shown in Figure *Traffic Shaper Rules List*
- Click  to edit an existing rule or  to create a copy of a rule
- Make any required adjustments to match different connections
- Save and Apply Changes as usual when editing firewall rules

Queues may be applied using *pass* rules on interface tabs, but the wizard only creates rules on the **Floating** tab using the *match* action that does not affect whether a connection is passed or blocked; it only queues traffic. Because these rules operate the same as any other rules, any criteria used to match connections may be used to queue.

See also:

For more information on floating rules, see *Floating Rules* and *Configuring Firewall Rules* for information on firewall rules in general.

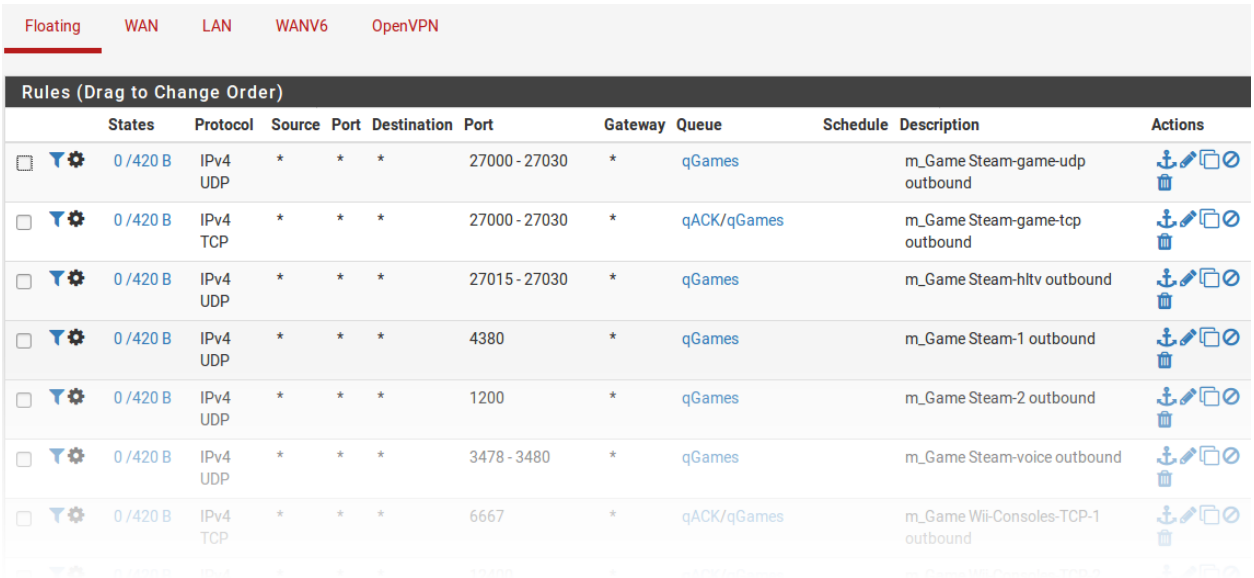


Fig. 2: Traffic Shaper Rules List

Shaper Rule Matching Tips

Connections can be tricky to match properly due to several factors, including:

- NAT applies before outbound firewall rules can match connections, so for connections that have outbound NAT applies as they leave a WAN-type interface, the private IP address source is hidden by NAT and cannot be matched by a rule.
- Some protocols such as BitTorrent will use random ports or the same ports as other services.
- Multiple protocols using the same port cannot be distinguished by the firewall.
- A protocol may use a range of ports so wide that it cannot be distinguished from other traffic.

While many of these cannot be solved by the firewall directly, there are ways to work around these limitations in a few cases.

To match by a private address source outbound in WAN floating rules, first tag the traffic as it passes in on a local interface. For example, match inbound on LAN and use the advanced **Tag** field to set a value, and then use the **Tagged** field on the WAN-side floating rule to match the same connection as it exits the firewall. Alternately, queue the traffic as it enters the LAN with a pass rule instead of when it exits a WAN.

Match by address instead of port/protocol where possible to sort out ambiguous protocols. In these cases, either the local source or the remote destination may be a single address or a small set of addresses. For example, matching VoIP traffic is much simpler if the firewall can match the remote SIP trunk or PBX rather than attempting to match a wide range of ports for RTP (e.g. `10000-20000`).

If BitTorrent is allowed on a network but must be shaped, then dedicate a specific local device that is allowed to use BitTorrent and then shape all connections to/from that device as Peer-to-Peer traffic.

25.5.3 Removing Traffic Shaper Settings

To remove all traffic shaper queues and rules created by the wizard:

- Navigate to **Firewall > Traffic Shaper**
- Click the **By Interface** tab
- Click  **Remove Shaper**
- Click **OK** on the confirmation prompt

25.6 Traffic Shaping with Differentiated Services (DiffServ) Identifiers

pfSense® software can use **Differentiated services** (**DiffServ**) values for traffic filtering or queue assignments. DiffServ takes the place of the outdated **Type of service** (**TOS**). DiffServ uses the upper six bits of the TOS field in the IP header (the six bits being called the *DiffServ Code Point field*), while the lower two bits are reserved for Explicit Congestion Notification (ECN).

Unless appropriately configured, pfSense software ignores the content of the DiffServ Code Point (DSCP) field. To prioritize traffic, the *Traffic Shaper* needs to be set up accordingly.

> **Warning:** pfSense software can **only** match DiffServ values, it *does not* implement setting or changing values.

25.6.1 DiffServ Code Point Values

Note that the interpretations of the DSCP values, as provided by the various RFCs, are only given as a reference. How the DSCP values are interpreted in any specific setup is entirely up to the user or end nodes.

The **Assured Forwarding** (AF) Behavior Group is recommended in RFC 2597.

Table 1: Assured Forwarding (AF) Behavior Group values

Precedence	Class 1 (lowest)	Class 2	Class 3	Class 4 (highest)
Low Drop	AF11 (10/0x0a)	AF21 (18/0x12)	AF31 (26/0x1a)	AF41 (34/0x22)
Med Drop	AF12 (12/0x0c)	AF22 (20/0x14)	AF32 (28/0x1c)	AF42 (36/0x24)
High Drop	AF13 (14/0x0e)	AF23 (22/0x16)	AF33 (30/0x1e)	AF43 (38/0x26)

For low-drop/low-latency traffic, use EF and VA DSCP values.

Table 2: Expedited Forwarding (EF) and Voice Admit (VA) values

PHB	DSCP Value	RFC
Expedited Forwarding (EF)	46/0x2e	RFC 3246
Voice Admit (VA)	44/0x2c	RFC 5865

The Class Selector (CS) PHB group has been retained from TOS.

Table 3: Class Selector (CS) values

Class Selector	DSCP Value
CS1	8/0x08
CS2	16/0x10
CS3	24/0x18
CS4	32/0x20
CS5	40/0x28
CS6	48/0x30
CS7	56/0x38

To provide limited backward comparability to TOS, pfSense software also recognizes the following DSCP/TOS values.

Table 4: TOS Compatibility values

TOS	DSCP Value	TOS value
reliability	1/0x01	4/0x04
throughput	2/0x02	8/0x08
lowdelay	4/0x04	16/0x10

pfSense software only matches exact values. All six bit in the DSCP field must match.

25.6.2 Caveats

By default, pfSense software only matches the **first packet** of a connection, which is the packet that creates an entry in the state table. If a connection starts with a different DSCP value, has no DSCP value in the starting packet, or otherwise changes DSCP values during the connection, the traffic will not be classified as expected.

Tip: This can be worked around by using "no state" rules, but crafting these rules securely is difficult, so it is not a viable workaround for most environments.

25.6.3 RFCs

- RFC 2474 — Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- RFC 2475 — An Architecture for Differentiated Services
- RFC 2597 — Assured Forwarding PHB Group
- RFC 2983 — Differentiated Services and Tunnels
- RFC 3086 — Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification
- RFC 3140 — Per Hop Behavior Identification Codes (replaces RFC 2836)
- RFC 3246 — An Expedited Forwarding PHB (Per-Hop Behavior) (obsoletes RFC 2598)
- RFC 3247 — Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior)
- RFC 3260 — New Terminology and Clarifications for Diffserv (updates
- RFC 2474, RFC 2475 and RFC 2597)

- RFC 4594 — Configuration Guidelines for DiffServ Service Classes
- RFC 5865 — A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic (updates RFC 4542 and RFC 4594)
- RFC 3289 — Management Information Base for the Differentiated Services Architecture
- RFC 3290 — An Informal Management Model for Diffserv Routers
- RFC 3317 — Differentiated Services Quality of Service Policy Information Base

25.7 Limiters

Limiters are an alternate method of traffic shaping. They use dummynet(4) to enact bandwidth limits and perform other prioritization tasks, not ALTQ. Limiters are currently the only way to achieve per-IP address or per-network bandwidth rate limiting using pfSense® software. They are used internally by Captive Portal for per-user bandwidth limits.

Manage Limiters in the GUI at **Firewall > Traffic Shaper** on the **Limiters** tab.

Like HFSC and CBQ, Limiters allow nesting with queues inside other queues. Root-level limiters ("pipes") may have bandwidth limits and delays, while child limiters ("queues"), may have priorities ("weights"). Bandwidth limits can be optionally masked by either the source or destination IP address, so that the limits can be applied on a per-IP address or network basis instead of as a general group.

Limiters are nearly always used in pairs: One for incoming traffic and one for outgoing traffic.

According to its man page the dummynet(4) system was originally designed as a means to test TCP congestion control, and it grew up from there. Due to this purpose, a unique feature of limiters is that they can be used to induce artificial packet loss and delay into network traffic. That is primarily used in troubleshooting and testing (or being evil and playing a prank on someone), and not often found in production.

25.7.1 Uses for Limiters

The primary use for limiters is to apply bandwidth limits for users or specific protocols, e.g. "Maximum of 1Mbit/s for SMTP", or "Joe's PC only can use 5Mbit/s". Limiters can apply a per-IP address or per-network limit, such as "All Users in 192.168.50.0/24 can use a maximum of 3Mbit/s each" or "The guest network and public network can use 1Mbit/s for each segment".

Limiters are the only type of shaper available in pfSense software which is capable of oversubscription in this manner. The ALTQ shaper requires all child queues to sum up to no more than the speed of the parent queue, but masked limiters allow a set limit to as many IP addresses as can be funneled through the limiter by firewall rules.

Conceptually, consider a limiter as a bucket of bandwidth. All traffic flowing through an unmasked limiter draws bandwidth from the same bucket. Masking a limiter effectively sets up multiple buckets of the same size, one per masked group. Whether that is a single host or an entire network depends on the mask value.

Limiters can also allow for reserved bandwidth by limiting everything *except* a specific protocol which can then consume all remaining bandwidth. In this type of setup on a 10Mbit/s link the firewall would pass traffic from, for example, a SIP server with no limiter. Then the firewall would use a pass rule for all other traffic with a limit of 8Mbit/s. This would let the SIP server use all the bandwidth it wanted, but it would always have a minimum of 2Mbit/s to itself.

Limiters can also help with issues such as Bufferbloat by controlling the delay of certain packets, using the CoDel algorithm similar to the one available in ALTQ (*CoDel Active Queue Management*).

See also:

- *Configuring CoDel Limiters for Bufferbloat*

25.7.2 How Limiters Work

Limiters, like ALTQ, hold traffic to a certain point by dropping or delaying packets to achieve a specific line rate. Usually taking advantage of built-in mechanisms from protocols that detect the loss and back off to a sustainable speed.

In situations where packets are queued under the same parent pipe, the firewall considers their weights when ordering the packets before it sends them. Unlike priorities in CBQ and PRIQ, the weight of a queue in a limiter will never starve it for bandwidth.

25.7.3 Limiters and IPv6

Limiters work with IPv6, though it requires separate IPv4 and IPv6 rules to apply limiters properly.

25.7.4 Limitations

Limiter pipes do not have a concept of borrowing bandwidth from other pipes. A limit is always a hard upper limit.

Limiters use `dummynet` pipes, so there will be additional (though small) overhead from the extra packet processing involved.

Limiters cannot effectively guarantee a minimum bandwidth amount for a pipe or queue, only a maximum.

Child queues cannot have bandwidth values, so a pipe cannot be split into smaller pipes by queues. Child queues can only use weights to prioritize packets inside a pipe.

The overhead from delaying and queuing packets can cause increased mbuf usage. For more information on increasing the amount of available mbufs, see *Hardware Tuning and Troubleshooting*.

25.7.5 Limiters and Multi-WAN

When using limiters with Multi-WAN, limits for non-default gateways must be applied using floating rules set for the *out* direction and configured with the appropriate gateway.

25.7.6 Creating Limiters

Limiters are managed under **Firewall > Traffic Shaper** on the **Limiters** tab.

To create a new root-level limiter (pipe), click  **New Limiter**.

To create a child limiter (queue), click an existing limiter under which it can be created, and click  **Add New Queue**.

Tip: In nearly all cases, limiters exist in pairs at the same level (e.g. two pipes, or two queues): One for inbound traffic and one for outbound traffic. When creating new limiters or queues, create one for each direction.

Enable
: Check the box to enable this limiter. If the limiter is disabled, it will not be available for use by firewall rules.

Name
This defines the name of the limiter, as it will appear for selection on firewall rules.

The name must be alphanumeric, and may also include - and _.

Tip: When choosing a name, avoid using `In` and `Out` since the same limiter, if used on both WAN and LAN, would be used in the *In* direction on one interface and the *Out* direction on another. The best practice is to use `Down` or `Download` and `Up` or `Upload`.

Bandwidth (Pipes)
This section defines a bandwidth value for the pipe, or multiple bandwidths if schedules are involved. This option does not appear when editing a child limiter (queue).

Bandwidth
The numerical part of the bandwidth for the pipe, e.g. `3` or `500`.

Bw Type
The units for the **Bandwidth** field, such as *Mbit/s*, *Kbit/s*, or *Bit/s*.

Schedule
If the firewall has schedules defined (*Time Based Rules*), the firewall offers them in this list. When schedules are in use by the firewall, the limiter can have a bandwidth value for each potential schedule. Define these by clicking  **Add Schedule** to add another bandwidth definition.

If a limiter contains multiple bandwidth specifications, they must each use a different schedule. For example if the firewall has a "Work Day" schedule, then it must also have an "Off Hours" schedule that contains all the time periods not included in "Work Day" for the second bandwidth specification.

Mask
This drop-down list controls how the limiter will mask addresses in the pipe or queue.

None
When set to *none*, the limiter does not perform any masking. The pipe bandwidth will be applied to all traffic as a whole.

Source / Destination address
When a limiter is set for *Source Address* or *Destination Address*, the pipe bandwidth limit will be applied on a per-IP address basis or a subnet basis, depending on the masking bits, using the direction chosen in the masking.

In general, a limiter should mask the **Source Address** on **Upload** (In) limiters for LAN-type interfaces, and **Destination Address** on **Download** (Out) limiters on LAN-type interfaces. Similar to swapping the directionality of the limiters when applying to LAN and WAN, masking is swapped as well, so the same masked limiter set for **In** on LAN should be used for **Out** on WAN.

Mask Bits
There are separate boxes to control the address masking for IPv4 and IPv6. For IPv4 a value of *32* for **IPv4 mask bits** sets up a per-IPv4 address limit, which is the most common usage. For a per-IPv6-address limit, use *128* as the **IPv6 mask bits** value.

To create per-subnet or similar masks, enter the subnet bits in the appropriate field for either IPv4 or IPv6 mask bits, such as *24* to limit IPv4 in groups of /24 subnets.

Description
An optional bit of text to explain the purpose for this Limiter.

Advanced Options
Additional options that vary when editing a pipe or a queue.

Delay (Pipes)
The **Delay** option is only found on limiter pipes. It introduces an artificial delay (latency), specified in milliseconds, into the transmission of any packets in the limiter pipe. This is typically left blank so that packets are transmitted as fast as possible by the firewall. This can be used to simulate high-latency connections such as satellite uplinks for lab testing.

Weight (Queues)
The **Weight** option is only found on child limiters (queues). This value can range from 1 to 100. Higher values give more precedence to packets in a given queue. Unlike PRIQ and CBQ priorities, a lowly-weighted queue is not in danger of being starved of bandwidth by the firewall.

Packet loss rate
Another method of artificially degrading traffic. The **Packet Loss Rate** can be configured to drop a certain fraction of packets that enter the limiter. The value is expressed as a decimal representation of a percentage, so `0.01` is 1%, or one packet out of a hundred dropped. This field is typically left empty, so every packet is delivered by the firewall.

Queue Size
Sets the size of the queue, specified in queue slots, used for handling queuing delay. Left blank, it defaults to `50` slots, which is the recommended value. Slow speed links may need a lower queue size to operate efficiently. High speed links may need more slots.

Tip: In cases where there are several limiters or limiters with large **Queue Size** values, a **System Tunable** may need to be set which increases the value of `net.inet.ip.dummynet.pipe_slot_limit` above the total number of configured queue slots among all pipes and queues.

Bucket Size
The **Bucket Size**, also specified in slots, sets the size of the hash table used for queue storage. The default value is `64`. It must be a numeric value between `16` and `65536`, inclusive. This value is typically left blank.

See also:

For more information about these values, consult the ipfw(8) man page, in the section titled "Traffic Shaper (dummynet) Configuration". Though current versions of pfSense software utilize dummynet through PF instead of `ipfw`, the configuration options are the same.

25.7.7 Assigning and Using Limiters

Limiters are assigned using firewall rules via the **In/Out Pipe** selectors under **Advanced Options**. Any matching criteria available in firewall rules can assign traffic to a limiter.

The most important thing to remember when assigning a limiter to a rule is that the **In** and **Out** fields are designated **from the perspective of the firewall itself**.

For example, in a firewall configuration with a single LAN and single WAN, inbound traffic on a LAN interface is leaving toward the Internet, i.e. *uploaded* data. Outbound traffic on the LAN interface is going toward the client PC, i.e. *downloaded* data. On the WAN interface the directionality is reversed; Inbound traffic is coming from the Internet to the client (download), and outbound traffic is going from the client to the Internet (upload).

In most cases, a firewall rule will have both an **In** limiter and **Out** limiter, but only the **In** limiter is required by the firewall to limit traffic in a single direction.

Limiters may be applied on normal interface rules, or on floating rules. On floating in the *out* direction, the In/Out selections are flipped conceptually.

25.7.8 Checking Limiter Usage

Information about active limiters may be found under **Diagnostics > Limiter Info**. Here, each limiter and child queue is shown in text format.

The set bandwidth and parameters for each limiter are displayed by the page, along with the current traffic level moving inside the limiter. In the case of masked limiters, the firewall displays the bandwidth of each IP address or masked group.

25.8 Traffic Shaping and VPNs

The following discussions pertain primarily to ALTQ shaping. Limiters will work fine with VPNs as they would with any other interface and rules. Only the ALTQ shaper requires special consideration.

Traffic shaping with VPNs is a tricky topic because VPN traffic is considered separate from, but also a part of, the WAN traffic through which it also flows. If WAN is 10 Mbit/s, then the VPN can also use 10Mbit/s, but there is not actually 20Mbit/s of bandwidth to consider, only 10Mbit/s. As such, methods of shaping that focus more on prioritization than bandwidth are more reliable, such as PRIQ or in some cases, CBQ.

If all traffic inside the VPN must be prioritized by the firewall, then it is enough to consider only the VPN traffic itself directly on WAN, rather than attempting to queue traffic on the VPN separately. In these cases, use a floating rule on WAN to match the VPN traffic itself. The exact type of traffic varies depending on the type of VPN. IPsec and PPTP traffic on WAN can both be prioritized by the shaper wizard, and these rules can be used as an example to match other protocols.

25.8.1 OpenVPN

With OpenVPN, multiple interfaces exist on the operating system, one per VPN. This can make shaping easier in some cases. Features of OpenVPN can also make it easier to shape traffic on WAN and ignore the tunnel itself.

Shaping inside the tunnel

If multiple classes of traffic are carried on the tunnel, then prioritization must be done to the traffic inside the tunnel. In order for the wizard to consider the traffic in this way, the VPN must be assigned as its own interface in the GUI. To accomplish this, assign it as described in *Interface assignment and configuration*, and then use the shaper wizard as if it were a separate WAN interface, and classify the traffic as usual.

Shaping outside the tunnel (passtos)

If the primary concern is shaping VoIP traffic over a VPN, another choice to consider is the `passtos` option in OpenVPN, called **Type-of-Service** in the OpenVPN client or server options. This option copies the TOS bit from the inner packet to the outer packet of the VPN. Thus, if the VoIP traffic has the TOS (DSCP) portion of the packet header set, then the OpenVPN packets will also have the same value.

This option is more useful for signaling intermediate routers about the QoS needs, however. Though the DSCP option on firewall rules can match based on TOS bits, as described in *Diffserv Code Point*, such matching would have to occur in the packet creating a firewall state, and not on specific packets flowing through that state.

Note: Because this option tells OpenVPN to copy data from the inner packet to the outer packet, it does expose information about the type of traffic crossing the VPN. Whether or not the information disclosure, though minor, is worth the risk for the gains offered by proper packet prioritization depends on the needs of the network environment.

25.8.2 IPsec

IPsec is presented to the operating system on a single interface no matter how many tunnels are configured and no matter which WANs are used by the tunnels. This makes shaping IPsec traffic difficult, especially when trying to shape traffic inside one particular IPsec tunnel.

The IPsec interface is also not possible to use on its own as an interface with the wizard. Floating rules can match and queue traffic on the IPsec interface, but in most cases only inbound traffic will be queued as expected. Actual results may vary.

25.9 Traffic Shaping UPnP IGD & PCP Connections

UPnP IGD & PCP rules are generated dynamically by the UPnP IGD & PCP daemon. These dynamic rules exist outside user-defined firewall rules and cannot be edited manually. There is a configuration option for UPnP IGD & PCP where a queue can be defined to which UPnP IGD & PCP will direct traffic matching the rules it creates. This may be set through the pfSense® software GUI at **Services > UPnP IGD & PCP**, and type in a valid **Traffic Shaping Queue**.

See also:

- *ALTQ Traffic Shaper Queue Monitoring*
- *Using the Shaper Wizard to Configure ALTQ Traffic Shaping*
- *Configuring CoDel Limiters for Bufferbloat*
- *Troubleshooting Traffic Shaping*
- *Troubleshooting Traffic Shaping Graphs*

Traffic shaping, or network Quality of Service (QoS), is a means of prioritizing network traffic. Without traffic shaping, packets are processed on a first in/first out basis by the firewall. QoS offers a means of prioritizing different types of traffic, ensuring that high priority services receive the bandwidth they need before lesser priority services.

For simplicity, the traffic shaping system in pfSense® software may also be referred to as the “shaper”, and the act of traffic shaping may be called “shaping”.

25.10 Traffic Shaping Types

There are two types of QoS available in pfSense software: ALTQ and Limiters.

The **ALTQ** framework is handled through PF and is closely tied to network card drivers. ALTQ can handle several types of schedulers and queue layouts. The traffic shaper wizard configures ALTQ and gives firewall administrators the ability to quickly configure QoS for common scenarios, and it allows custom rules for more complex tasks. ALTQ is inefficient, however, so the maximum potential throughput of a firewall is lowered significantly when it is active.

pfSense software also includes a separate shaper concept called **Limiters**. Limiters enforce hard bandwidth limits for a group or on a per-IP address or network basis. Inside those bandwidth limits, limiters can also manage traffic priorities.

25.11 Traffic Shaping Basics

For administrators who are unfamiliar with traffic shaping, it is like a bouncer at an exclusive club. The VIPs (Very Important Packets) always make it in first and without waiting. The regular packets have to wait their turn in line, and "undesirable" packets can be kept out until after the real party is over. All the while, the club is kept at capacity and never overloaded. If more VIPs come along later, regular packets may need to be tossed out to keep the place from getting too crowded.

ALTQ shaping concepts can be counter-intuitive at first because the traffic has to be queued in a place where the operating system can control the flow of packets. Incoming traffic from the Internet going to a host on the LAN (downloading) is shaped *leaving* the LAN interface from the firewall. In the same manner, traffic going from the LAN to the Internet (uploading) is shaped when leaving the WAN.

For ALTQ, there are traffic shaping queues, and traffic shaping rules. The queues allocate bandwidth and priorities. Traffic shaping rules control how traffic is assigned into those queues. Rules for the shaper work the same as firewall rules, and allow the same matching characteristics. If a packet matches a shaper rule, it will be assigned into the queues specified by that rule. In pfSense software, shaper rules are mostly handled on the **Floating** tab using the *Match* action that assigns the traffic into queues, but rules on any interface can assign traffic into queues using the *Pass* action.

Limiter rules are handled differently. Limiters apply on regular pass rules and enforce their limits on the traffic as it enters and leaves an interface. Limiters almost always exist in pairs: One for the "download" direction traffic and one for the "upload" direction traffic.

CHAPTER

TWENTYSIX

CAPTIVE PORTAL

26.1 Captive Portal Zones

Captive Portal zones define separate portals for different sets of interfaces. For example, LAN and Wireless could use one portal, while a conference room would get a separate portal page. Each zone has separate settings for HTML pages, authentication, allowed addresses, and so on. A zone must be created before its settings can be changed.

Note: A zone may have multiple interfaces, but an interface may only be a member of one zone. Attempting to add the same interface to multiple zones will result in an error.

26.1.1 Managing Captive Portal Zones

Captive Portal zones are managed at **Services > Captive Portal**. A list of zones is displayed there, and zones may be added, edited, or deleted from that list.

To create a new Captive Portal zone:

- Navigate to **Services > Captive Portal**
- Click  **Add**
- Enter a **Zone Name**, which may only consist of letters, digits, numbers, and underscores. Spaces and other special characters may not be used
- Enter an optional **Zone Description** to further describe the zone, if desired
- Click **Save & Continue** to move on to the portal settings for the zone

To edit an existing zone, click  at the end of its row.

To delete an existing zone, click  at the end of its row, and then click to confirm the action.

26.2 Common Captive Portal Scenarios

The following are some basic, common scenarios for the use of a Captive Portal. The details of how to perform all the actions described will be covered throughout this chapter.

26.2.1 Portal Configuration Without Authentication

For a simple portal without authentication:

- Create a new Zone
- Check **Enable captive portal**
- Select an **Interface**
- Upload an HTML page with the portal contents as described in *Portal page without authentication*
- Click **Save**

Additional configuration options may be added as detailed in *Zone Configuration Options*.

26.2.2 Portal Configuration Using Local Authentication or Vouchers

To set up a portal with local authentication:

- Create a Zone
- Check **Enable captive portal**
- Select an **Interface**
- Set **Authentication Method** to **Local User Manager / Vouchers**
- Upload an HTML page with the portal contents as described in *Portal page with authentication*.

Additional configuration options may be added as detailed in *Zone Configuration Options*. Then configure the local users in the **User Manager** (*User Management and Authentication*).

To use vouchers, proceed to the **Vouchers** tab and create them there. See *Vouchers* for more information on Vouchers, and use the sample portal page HTML code from *Portal page with Vouchers*.

26.2.3 Portal Configuration Using RADIUS Authentication

To set up a portal using RADIUS authentication:

- Configure the RADIUS server to allow requests from the firewall
- Create a Zone
- Check **Enable captive portal**
- Select an **Interface**
- Set **Authentication Method** to **RADIUS Authentication**
- Fill in the settings for **Primary RADIUS Server** under **Primary Authentication Source**

Read the next section for information on specific configuration options.

26.3 Zone Configuration Options

This section describes each of the configuration options for a Captive Portal zone. Options for a zone are independent of those for other zones. For example, allowed IP address entries in a zone only affect that specific zone.

To reach this page, navigate to **Services > Captive Portal** and edit an existing zone from the list with  , or click  **Add** to create a new zone.

Enable
Check to enable this Captive Portal zone.

Description
Brief text describing the purpose of the zone.

Interface
Determines the interfaces that used by this Captive Portal zone. This **cannot** be a WAN interface. It can be a bridge interface so long as it is the actual bridge (e.g. `bridge0`) and the bridge interface has an IP address assigned.

Maximum concurrent connections
Specifies the maximum number of concurrent connections to the portal web server per IP address. The default value is 4, which is sufficient for most environments. This limit exists to prevent a single host from exhausting all resources on the firewall, whether inadvertent or intentional.

One example where this would otherwise be a problem is a host infected with a worm. The thousands of connections issued will cause the captive portal page to be generated repeatedly if the host is not authenticated already, which would otherwise generate so much load it could leave the firewall unresponsive.

Idle timeout
A timeout, specified in minutes, after which idle users will be disconnected by the portal. Users may log back in immediately.

Hard timeout
A timeout, specified in minutes, after which the portal will forcefully log off users.

Tip: Set either a hard timeout, idle timeout, or both to ensure sessions are removed by the portal when users do not log off manually.

Users may log back in immediately after the hard timeout if their credentials are still valid (for local accounts, not expired, and for RADIUS authentication, user can still successfully authenticate to RADIUS).

Note: If a timeout value is set, the timeout must be less than the DHCP lease time or captive portal sessions can remain active for IP addresses that have switched to different devices. Setting the timeout lower will ensure that the portal sessions end before the lease would be reallocated to a new client.

Traffic Quota
An amount of traffic which, when exceeded by a client, will trigger a disconnect of that client by the portal. This includes both upload and download traffic. Users may log back in immediately if their credentials are still valid

Pass-Through Credits
: These credits give devices a grace period before they must authenticate via the portal. For example, a device could connect 3 times within a day without seeing the portal page, but any more than that, and they must log in. By setting the hard timeout to a value such as 1 hour, the portal would effectively limit a client to three hours of access before forcing it to authenticate. By default, this is disabled, and all clients are presented with the portal login page and must log in.

Note: For this to be effective, set a hard timeout and/or idle timeout.

Pass-through credits allowed per MAC address
: The number of times a specific MAC address may connect through the portal. Once the client uses its credits, it can only log in with valid credentials until the waiting period has expired.

Waiting period to restore pass-through credits
: The number of hours after which the portal will restore the pass-through credits for a client to the original count after it uses the first one. This must be above `0` hours.

Reset waiting period on attempted access
: If enabled, the waiting period is reset by the portal to the original duration if access is attempted when all pass-through credits have already been exhausted. This prevents people who repeatedly attempt to access the portal from gaining open access too quickly.

Logout popup window
: When checked, the portal attempts to show a logout pop up window to the user which allows clients to explicitly disconnect themselves before the idle or hard timeout occurs. Unfortunately, since most browsers block pop up windows, this window may not work for most clients unless.

Pre-authentication redirect URL
: As the name implies, this option redirects users to the specified URL *before* they authenticate. Commonly, this is used to display a custom landing page describing the device location hosted on a server locally or elsewhere. That landing page must contain a link which in turn redirects the users back to the portal page, e.g. `http://x.x.x.x:8002/index.php?zone=somezone&redirurl=http%3A%2F%2Fsomesite.example.com`.

  See also:

  See *Allowed Hostnames* to allow hostnames through the portal without authentication, and *Allowed IP Address* for IP addresses.

  The custom captive portal page must have extra code at the top to properly handle this redirect. In the example code below, the pre-authentication redirect target page must also put its own URL in the `redirurl` parameter of its link back to the portal in order for the login page to appear.

```
<?php
require_once("globals.inc");
$request_uri = urldecode(str_replace("/index.php?zone={$_REQUEST['zone']}&redirurl=", "",
↪ $_SERVER["REQUEST_URI"]));
$portal_redirurl = urldecode("$PORTAL_REDIRURL$");
if(!stristr($portal_redirurl, $request_uri)) {
     Header("Location: $PORTAL_REDIRURL$");
     exit;
}
?>
```

After authentication Redirection URL
: After authenticating or clicking through the portal, it will redirect users to this URL rather than the

one they originally tried to access. If this field is left blank, the portal will redirect the user to the URL they initially attempted to access.

Blocked MAC address redirect URL
URL to which the portal will redirect users with blocked MAC addresses when they attempt access through the portal.

Preserve user database
When set, the database containing logged-in users is preserved by the portal when the firewall reboots.

Concurrent user logins
Controls whether users are allowed to connect multiple times. This is not a total limit for the entire portal, but a per-account limit.

May be set to one of the following:

Disabled
The portal will not allow concurrent logins for a user or voucher.

Multiple
(Default) The portal does not enforce any restrictions on concurrent logins by a user or voucher.

Last Login
The portal will only allow one login per user account or voucher. The most recent login is permitted, and any previous logins are disconnected.

First Login
The portal will only allow one login per user account or voucher. The portal permits the first login and denies any subsequent login attempt.

MAC filtering
When set, the portal disables MAC address filtering. This is necessary in cases where the MAC address cannot reliably be determined, such as when multiple subnets exist behind a separate router using the portal. In that type of situation, all users behind a router will show up to the portal with the MAC address of the intermediate router. If this option is set, the portal will not attempt to ensure that the MAC address of clients stay the same while they are logged into the portal.

Note: This option is not compatible with RADIUS MAC authentication.

Pass-through MAC Auto Entry
In certain use cases, users may only need to authenticate once per device, and then the client should not see the portal login again unless they change devices. Setting up pass-through MAC entries can automatically achieve this goal.

Pass-through MAC automatic additions
If this option is set, the portal automatically adds a MAC pass-through entry after the user has successfully authenticated. Users of that MAC address will never have to authenticate again so long as the entry is present in the configuration. To remove the pass-through MAC entry, log in and remove it manually from the Pass-through MAC tab.

Note: This option is not compatible with RADIUS MAC authentication or the logout window.

Pass-through MAC automatic addition with username
If this option is set, the portal saves the username used during authentication along

with the pass-through MAC entry. To remove the pass-through MAC entry, log in and remove it manually from the **Pass-through MAC** tab.

Per-user bandwidth restrictions
Captive Portal can also optionally rate-limit users to keep them from using too much bandwidth. The **Default download** and **Default upload** fields define the default values for user bandwidth, specified in Kilobits per second. These values can be overridden by RADIUS (*Passing back configuration from RADIUS Servers*) for different limits for specific users. If the fields are blank or set to `0`, then users have unlimited bandwidth.

Use Custom Captive Portal Page
When set, enables upload controls for manually crafted portal pages. See *HTML Page Contents* for details. When unset, simple portal page customization controls are available (*Captive Portal Login Page*).

26.3.1 Captive Portal Login Page

These simple customization controls enable small changes to the portal page without writing custom HTML. For more complicated portal pages, see *HTML Page Contents*.

Display Custom Logo Image
When set, the portal page includes the custom image from **Logo Image** instead of the default logo.

Logo Image
Upload control for setting a custom logo image.

Display Custom Background Image
When set, the portal page includes the custom image from **Background Image** instead of the default background.

Background Image
Upload control for setting a custom background image.

Terms and Conditions
Text displayed by the portal to the user to which the user must agree before they are permitted to log in.

26.3.2 Authentication

This section configures authentication for Captive Portal. If authentication is required for the zone it may be handled by the local user database, RADIUS, or LDAP.

Authentication Method

Use an Authentication backend
This option allows users to authenticate with a username and password or vouchers. The authentication is handled by the local user database (*User Management and Authentication*) or an authentication server (*Authentication Servers*).

Vouchers are pre-generated access codes which grant short-term access to users. Vouchers may be used in addition to, or instead of, user authentication. For more information on using vouchers, see *Vouchers* later in this section.

None, don't authenticate users
The portal only requires users to click through the login page for access. The form must still be submitted, but it does not need to have any user entry fields, only a submit button.

Use RADIUS MAC Authentication
The portal attempts to authenticate users by sending their MAC address as the username and the password entered into **MAC authentication secret** to the RADIUS server.

Note: Users must still attempt an HTTP connection for the portal to see their attempt and perform the initial authentication.

See *RADIUS MAC Authentication Options* for additional options.

This option is not available if MAC filtering is disabled.

Authentication Server
A multi-select control where one or more primary authentication servers, or the local database, can be set for use by the portal. See *Primary Authentication Source* for more information.

Local Database
Captive Portal users in this mode are managed in the pfSense® software GUI. Local users are added in the User Manager (*Manage Local Users*).

Additionally, the **Local Authentication Privileges** option can limit access to only users who possess the proper access privileges.

LDAP Server
When an LDAP server is active in the control, it is used by the portal for authentication as-is. There are no additional options for LDAP server behavior.

RADIUS Server
When a RADIUS server is active in the control, numerous RADIUS server options are displayed by the GUI and Captive Portal users in this zone will be validated against the configured RADIUS server(s).

Secondary Authentication Server
Similar to **Authentication Server**, but sets up an additional separate means of authentication using distinct fields. See *Secondary Authentication Source* for more information.

Primary Authentication Source

The Primary/Secondary authentication servers are used for the main username and password fields on the login form, `auth_user` and `auth_pass`, such as:

```
<tr>
     <td align="right">Username:</td>
     <td><input name="auth_user" type="text" style="border: 1px dashed;"></td>
 </tr>
<tr>
     <td align="right">Password:</td>
     <td><input name="auth_pass" type="password" style="border: 1px dashed;"></td>
</tr>
```

If the first server is down, the portal will attempt authentication using the other servers in the list, in order (top down).

Secondary Authentication Source

The secondary authentication source defines a completely separate authentication setup from the primary. For example, the primary source could be traditional usernames and passwords, while the secondary could be pre-paid card numbers or PINs.

The secondary authentication source uses the form fields `auth_user2` and `auth_pass2` in the captive portal HTML, such as:

```
<tr>
     <td align="right">Username:</td>
     <td><input name="auth_user2" type="text" style="border: 1px dashed;"></td>
</tr>
<tr>
     <td align="right">Password:</td>
     <td><input name="auth_pass2" type="password" style="border: 1px dashed;"></td>
</tr>
```

If the first server is down, the portal will attempt authentication using the other servers in the list, in order (top down).

Local Database Options

Local Authentication Privileges
When **Allow only users/groups with 'Captive portal login' privilege set** is active, the portal will limit access to only users who have Captive Portal privilege. The privilege can be directly on their account or inherited via group membership.

RADIUS Authentication Options

RADIUS is a means of authenticating users against a central server that contains account information. There are many implementations of RADIUS, such as FreeRADIUS, Radiator, and NPS on Windows servers.

RADIUS accounting can be enabled to send usage information for each user to the RADIUS server. Refer to documentation for the RADIUS server for more information.

See also:

To add or edit RADIUS server entries on pfSense software, see *Authentication Servers*.

See also:

For those with a Microsoft Active Directory network infrastructure, RADIUS can be used to authenticate captive portal users from Active Directory using Microsoft NPS. This is described in *Authenticating from Active Directory using RADIUS/NPS*.

Passing back configuration from RADIUS Servers

Some default Captive Portal settings can be overridden by reply attributes from RADIUS servers. The exact attributes can vary by vendor, and may not be compatible with all RADIUS servers.

User bandwidth restrictions
: Defines the bandwidth for the user, drawn from common options such as `WISPr-Bandwidth-Max-Up`/`WISPr-Bandwidth-Max-Down`, or `ChilliSpot-Bandwidth-Max-Up`/`ChilliSpot-Bandwidth-Max-Down`.

Session Timeout
: Drawn from the RADIUS attribute `Session-Timeout`, it will disconnect the user after the time specified by the RADIUS server.

Idle Timeout
: Drawn from the RADIUS attribute `Idle-Timeout`, it will disconnect the user after the time specified by the RADIUS server.

Accounting Interval Interim
: Taken from `Acct-Interim-Interval`, it directs the portal to send interim accounting updates at the specified interval.

URL Redirection
: Allows the after-authentication redirect URL to be defined by the RADIUS server through `WISPr-Redirection-URL`.

RADIUS Options

These options fine-tune how RADIUS authentication behaves.

NAS Identifier
: Configures an alternate NAS Identifier to send with RADIUS requests. The default value is `CaptivePortal-<zone name>`.

Reauthentication
: If enabled, the portal sends `Access-Request` packets to the RADIUS server for each user that is logged in every minute. If an `Access-Reject` is received for a user, that user is disconnected from the captive portal immediately. This allows actively terminating user sessions from the RADIUS server.

> **Warning:** If concurrent login limits are defined in RADIUS this option may not work properly, as the additional request would fail as the reauthentication attempt would be considered a second concurrent login.

Note: If reauthentication is combined with RADIUS accounting, **Interim** accounting updates must be used to track usage during sessions, otherwise the RADIUS server will not know if a user exceeds limits until they log out.

Session-Timeout
: When set, clients will be disconnected after the amount of time set by the RADIUS `Session-Timeout` attribute sent to the portal at login.

Traffic Quota
: When set, the portal uses the `pfSense-Max-Total-Octets` reply attribute sent by the RADIUS

server to set a traffic quota for a user. This determines an amount of traffic which, when exceeded by a client, will trigger a disconnect of that client by the portal. This includes both upload and download traffic.

Per-user Bandwidth Restrictions
When set, the portal uses the `pfSense-Bandwidth-Max-Up` and `pfSense-Bandwidth-Max-Down` reply attribute sent by the RADIUS server to set per-user bandwidth restrictions.

MAC address format
This option changes the MAC address format used in RADIUS. Change this to alter the username format for RADIUS MAC authentication to one of the following styles:

Default
Colon-separated pairs of digits: `00:11:22:33:44:55`

Single Dash
Digits in two groups, separated by a single dash halfway: `001122-334455`

IETF
Hyphen-separated pairs of digits: `00-11-22-33-44-55`

Cisco
Groups of four digits separated by a period: `0011.2233.4455`

Unformatted
All digits together with no formatting or separators: `001122334455`

RADIUS MAC Authentication Options

RADIUS MAC Secret
When the portal attempts RADIUS MAC authentication, it sends the MAC Address as the username and this value as the password.

Login Page Fallback
When set, the portal will redirect a client to the login page if MAC Authentication failed.

Accounting

RADIUS accounting sends session information back to the RADIUS server indicating when a user session starts, ends, and how much data they have transmitted.

Warning: Not all RADIUS servers are capable of using accounting data, or they may not be configured to accept accounting data. Set up the RADIUS server properly before enabling this feature.

Accounting Server
An authentication server entry for a RADIUS server to which the portal will send accounting data (*Authentication Servers*).

Send Accounting Updates
Configures the specific type of accounting messages sent to the server.

No updates
Synonymous with disabling accounting, the portal will not send accounting updates to the server.

Stop/start
The portal sends START and STOP records for a user session only.

Stop/start (FreeRADIUS)
The portal sends START and STOP records for a user session only, in a way that is compatible with FreeRADIUS.

Interim
The portal sends START and STOP records and also periodically sends updates to the server while a user session is active. This method is less likely to lose session data if the firewall restarts without notifying the RADIUS server of a STOP message, but will cause increased database usage on the RADIUS server.

Accounting Style
When **Invert Acct-Input-Octets and Acct-Output-Octets** is enabled, data counts for RADIUS accounting packets will be taken from the client perspective, not the NAS. `Acct-Input-Octets` will represent download, and `Acct-Output-Octets` will represent upload.

Idle Time Accounting
This option changes the time sent in the STOP message for a user disconnected by the portal for idle timeout. When unset (default), the time sent is the last activity time. When set, the idle time is included.

26.3.3 HTTPS login

Login
When set, the portal will listen for and accept HTTPS requests for the portal page. This option requires an **SSL/TLS Certificate**.

HTTPS server name
The FQDN (hostname + domain) used by the portal for HTTPS. This must match the Common Name (CN) on the certificate to prevent users from receiving certificate errors.

SSL Certificate
Select the SSL certificate used by the portal for HTTPS. Certificates are managed in *Certificate Management*.

Disable HTTPS Forwards
When checked, attempts by clients to connect to HTTPS sites on port 443 are not redirected to the portal. This prevents users from receiving invalid certificate errors. Users must attempt a connection to an HTTP site, and will then be forwarded to the portal.

26.3.4 HTML Page Contents

When **Use custom captive portal page** is set on the zone, the portal displays these controls to upload custom HTML pages to alter the look of the page presented to users when they are redirected to the portal.

Customizing these pages is optional. Any page contents left blank will use internal defaults.

Portal pages may contain PHP code, and may also include other resources such as images and CSS files. See *File Manager* for more information on including additional assets in a custom portal page.

> **Warning:** Since custom portal pages can run PHP, audit the code to ensure security so that the page cannot be exploited by connecting users. Also, avoid granting privileges to this page to untrusted administrators.

In each individual section, the pages can be managed by the displayed controls:

- To upload a new page, click **Browse** and select the file to upload. When the portal options are saved, the file will be copied.
- To view an existing page, click  **View Page Contents**
- To download a copy of an existing page, click  **Download**
- To erase the custom page, click  **Restore Default Page**

Portal page contents

This control is for the main portal page presented to users. Depending on the selected options for the portal, use one of the following examples as the basis for a custom page.

Portal page without authentication

This shows the HTML of a portal page that can be used without authentication.

Listing 1: Download: `example-noauth.html`

```
<html>
<head>
        <title>Welcome to our portal</title>
</head>
<body>
        <p>Welcome to our portal</p>
        <p>Click Continue to access the Internet</p>
        <form method="post" action="$PORTAL_ACTION$">
                <input name="redirurl" type="hidden" value="$PORTAL_REDIRURL$">
                <input name="zone" type="hidden" value="$PORTAL_ZONE$">
                <input name="accept" type="submit" value="Continue">
        </form>
</body>
</html>
```

Portal page with authentication

Here is an example portal page requiring authentication.

Listing 2: Download: `example-auth.html`

```
<html>
<head>
        <title>Welcome to our portal</title>
</head>
<body>
        <p>Welcome to our portal</p>
        <p>Enter your username and password and click Login to access the Internet</p>
```

(continues on next page)

(continued from previous page)

```
        <form method="post" action="$PORTAL_ACTION$">
                <input name="auth_user" type="text">
                <input name="auth_pass" type="password">
                <input name="redirurl" type="hidden" value="$PORTAL_REDIRURL$">
                <input name="zone" type="hidden" value="$PORTAL_ZONE$">
                <input name="accept" type="submit" value="Login">
        </form>
</body>
</html>
```

Portal page with Vouchers

Here is an example portal page for use with vouchers.

Listing 3: Download: `example-voucher.html`

```
<html>
<head>
        <title>Welcome to our portal</title>
</head>
<body>
        <p>Welcome to our portal</p>
        <p>Enter your voucher code and click Login to access the Internet</p>
        <form method="post" action="$PORTAL_ACTION$">
                <input name="auth_voucher" type="text">
                <input name="redirurl" type="hidden" value="$PORTAL_REDIRURL$">
                <input name="zone" type="hidden" value="$PORTAL_ZONE$">
                <input name="accept" type="submit" value="Login">
        </form>
</body>
</html>
```

Authentication error page contents

Using this control, optionally upload a custom HTML page to be displayed when authentication errors happen. An authentication error occurs when a user enters a bad username or password, or in the case of RADIUS authentication, potentially an unreachable RADIUS server.

By default, this error page is simply the login page again.

Logout page contents

The logout page is presented to the user after login and it triggers a popup window. The default code uses JavaScript to create the new window in the following way:

Listing 4: Download: `example-logout.html`

```
<html>
<head><title>Redirecting...</title></head>
```

(continues on next page)

(continued from previous page)

```
<body>
<span style="font-family: Tahoma, Verdana, Arial, Helvetica, sans-serif; font-size: 11px;
↪">
<b>Redirecting to <a href="<?=$my_redirurl;?>"><?=$my_redirurl;?></a>...</b>
</span>
<script type="text/javascript">
//<![CDATA[
LogoutWin = window.open('', 'Logout', 'toolbar=0,scrollbars=0,location=0,statusbar=0,
↪menubar=0,resizable=0,width=256,height=64');
if (LogoutWin) {
        LogoutWin.document.write('<html>');
        LogoutWin.document.write('<head><title>Logout</title></head>') ;
        LogoutWin.document.write('<body style="background-color:#435370">');
        LogoutWin.document.write('<div class="text-center" style="color: #ffffff; font-
↪family: Tahoma, Verdana, Arial, Helvetica, sans-serif; font-size: 11px;">') ;
        LogoutWin.document.write('<b>Click the button below to disconnect</b><p />');
        LogoutWin.document.write('<form method="POST" action="<?=$logouturl;?>">');
        LogoutWin.document.write('<input name="logout_id" type="hidden" value="<?=
↪$sessionid;?>" />');
        LogoutWin.document.write('<input name="zone" type="hidden" value="<?=$cpzone;?>"␣
↪/>');
        LogoutWin.document.write('<input name="logout" type="submit" value="Logout" />');
        LogoutWin.document.write('</form>');
        LogoutWin.document.write('</div></body>');
        LogoutWin.document.write('</html>');
        LogoutWin.document.close();
}

document.location.href="<?=$my_redirurl;?>";
//]]>
</script>
</body>
</html>
```

Most browsers have pop-up blockers that will most likely stop that logout window from appearing, so investigate other possible means of creating a JavaScript pop-up using similar code.

26.4 MAC Address Control

The **MACs** tab defines actions for MAC addresses that can be either passed through the portal for this zone without requiring authentication, or blocked from reaching the portal.

To manage these MAC entries:

- Navigate to **Services > Captive Portal**
- Click  on the line for the Zone to edit
- Click the **MACs** tab
- Click  **Add** to add a new entry

- Fill in the form as follows:

  Action
  Defines the action to take on this entry:

  Pass
  Always allow traffic through from this MAC address without authentication.

  Block
  Always deny traffic from this MAC address

  MAC address
  The MAC address of the device to allow. The value must be colon-separated pairs of digits, such as `00:11:22:33:44:55`.

  Description
  Some text describing the entry, if desired.

  Bandwidth up/down
  The amount of bandwidth that this device may use, specified in Kilobits per second. Leave blank to not specify a limit.

- Click **Save**

From this page, an entry may be edited by clicking  on its row, or deleted by clicking  .

26.5 Allowed IP Address

The **Allowed IP Address** tab works similarly to the **MACs** tab, except it checks IP addresses instead of MAC addresses. Traffic matching the specified IP address and the configured direction will always be allowed through the portal with no authentication in this zone.

IP Address
The IP address of the device to always pass through the portal.

Description
Some text describing the entry, if desired.

Direction
The direction to allow traffic matching this IP address.

From
Allow traffic sourced from this IP address through the portal, such as a local client IP address attempting to reach the Internet, or the IP address of a management client that must reach hosts on the portal network.

To
Allow traffic with this IP address as a destination, such as a local web server IP address that must be reached via port forward, or a remote web server IP address which clients must always reach.

Both
Allow traffic both to and from this IP address.

Bandwidth up/down
The amount of bandwidth that this device may use. Leave blank to not specify a limit.

26.6 Allowed Hostnames

Allowed Hostnames work similarly to **Allowed IP Address** entries, except they are configured by hostname instead of IP address. A daemon periodically resolves the hostnames to IP address(es) and allows them through the portal without authentication in this zone.

The most common use of this feature is to make a "walled garden" style portal, where users are permitted to access a restricted set of sites without authenticating to the portal. This is also commonly used with the **Pre-authentication Redirect URL** if that page is hosted externally.

Note: Often sites will use many hostnames, content delivery networks, or ad servers as part of their content. In order to allow a site to load fully, all of these additional sites must be added to the list of allowed hostnames.

Direction
The direction to allow traffic matching this hostname. In most typical use cases for allowing hostnames, the *To* or *Both* directions are the best fit.

To
Allow traffic from local clients to a remote site matching this hostname as a destination without authentication. For example, a remote web server that must always be reachable by local clients, even when they are not logged in.

From
Allow traffic sourced from this hostname through the portal, such as the hostname of a local client attempting to reach the Internet.

Both
Allow traffic both to and from this hostname.

Hostname
The fully qualified domain name (FQDN) of the target host or site. The hostname must exist in DNS so that it can be resolved to an IP address.

Description
Text describing the entry, if desired.

Bandwidth up/down
The amount of bandwidth that traffic to or from this hostname may use. Leave blank to not specify a limit.

26.7 Vouchers

Vouchers are single use codes used to gain Internet access through a Captive Portal. Each roll of vouchers is cryptographically generated and includes a set time limit. Vouchers are common in places where an organization wants Internet access to be authenticated and time-limited, without needing to provide a username and password to users.

This type of configuration is prevalent in travel and hospitality locations such as coffee shops, hotels, and airports. Users enter a voucher code in the portal login form and are the portal grants access for the amount of time allowed by the voucher roll. Voucher rolls can be exported by the GUI as a CSV file, and some companies have even integrated the exported voucher lists into point of sale applications to print voucher codes on customer receipts.

Voucher time does not stop counting down if a user logs out; they allow access only from the start of the session until the duration of the voucher length has elapsed. During that time, the voucher can be re-used by the same or a different computer. If the voucher is used again by another computer, the previous session is stopped.

Vouchers require a public/private RSA key pair to generate and verify. A 32-bit set of keys is generated automatically by the firewall the first time the page loads. A custom key pair may be generated manually by the GUI as well. The maximum key length is 64 Bits. Using shorter keys will make the voucher codes shorter but eventually less secure.

26.7.1 Voucher Options

Voucher options are unique per Captive Portal Zone. To configure vouchers, navigate to the **Vouchers** tab when editing a Captive Portal Zone.

Voucher Rolls
Voucher rolls are managed by this section of the page. The page lists information about each roll along with a link to add new rolls. No options appear here until after the other settings are active. See *Managing Voucher Rolls* for details.

Enable
When checked, this portal zone allows vouchers as a method of authentication, and the page changes to display the other voucher options.

Voucher Keys
The keys which generate and verify vouchers. Before vouchers are active on a zone, the GUI randomly generates new values each time the page loads. After saving voucher settings, the keys are present in the configuration and remain static from that point on.

> **Warning:** Do not change the keys or other bits after creating voucher rolls. If the values change, all current voucher rolls are invalid. Create new voucher rolls using the new settings after making changes.

Voucher Public Key
This key is used by the portal to decrypt vouchers. Use the existing random key, or click **Generate new keys** to make a new public and private key pair. Users may generate keys elsewhere and paste the RSA public key (64 Bit or smaller) in PEM format here.

Voucher Private Key
This key is used by the portal to generate voucher codes and does not need to be available if the vouchers are generated by another system. Use the existing random key or paste in an RSA private key (64 Bit or smaller) in PEM format here.

Character Set
The character set defines which characters are valid for voucher text. The character set is case-sensitive and should contain printable characters (numbers, lower case and upper case letters) that are hard to confuse with others. For example, avoid `0` (Digit zero), `O` (Letter O), `l` (Lowercase L), and `1` (Digit One). It cannot contain a space, double quote, or comma. A smaller character set will result in longer vouchers to ensure sufficient randomness.

Voucher Bits
The following "bit" fields control how the vouchers themselves are generated by the portal. The best practice is to leave these values at their defaults, but they may be adjusted if necessary. The total of all bit fields *must* be less than the RSA key size. For example, the default values are `16`, `10`, and `5`. The sum of these is `31`, which is one less than the default RSA key size of `32`.

# of Roll Bits
Number of bits for the Roll ID. Set this larger to have a lot of rolls active at the same time. Can be from `1-31`, the default value is `16`.

# of Ticket Bits
Number of bits for the Ticket ID. Set this larger if each roll will have many vouchers.

Can be from `1-16`, the default value is `10`.

# of Checksum Bits
Reserves a range in each voucher to store a simple checksum over **Roll bits** and **Ticket bits**. Allowed range is `0-31`, the default value is `5`.

Magic Number
The magic number is present in every voucher, and is verified by the portal during voucher check. The size of the magic number depends on the number of bits remaining after adding together the number of bits for the roll, ticket, and checksum. If there are no bits remaining for use by the magic number, then the portal does not use magic numbers.

Invalid Voucher Message
This message is displayed by the portal to the user if they attempt to enter a voucher that does not exist or is not valid in any way except for being expired.

Expired Voucher Message
This message is displayed to the user by the portal if they enter a voucher that was valid, but has expired.

26.7.2 Enable Vouchers

- Use the pfSense® GUI to navigate to **Services > Captive Portal**
- Click  on the line for the Zone to edit
- Ensure the Zone **Authentication Method** is set to *Use an Authentication backend*, change the value and save if necessary.
- Click the **Vouchers** tab
- Check **Enable**
- Fill in the form based on the options described in *Voucher Options*. In most cases, the options may remain at their default values.
- Click **Save**

With Vouchers enabled, the voucher management controls will be active in the GUI.

26.7.3 Managing Voucher Rolls

Vouchers are created by the portal in batches called **Rolls**. Each roll has specific settings that are unique to that roll. For example, a roll can have an 8-hour time limit and a separate roll can have a 12-hour time limit. Then users may be given voucher codes depending on which level of service they purchased, and they will be limited to the amount of time corresponding to the voucher roll from which their code was picked.

Voucher Roll Options

Roll #
The number of this roll. Each roll must have a unique number. This can be any number from `0` to `65535` with the default number of **Roll Bits**.

Minutes per Ticket
Defines how long the voucher lasts, in minutes. The voucher time starts counting down the moment the voucher is used, and does not stop, so plan the voucher length accordingly. Because this is defined in minutes, ensure the correct length is used, e.g. `1440` minutes is 24 hours.

Count
Defines the number of vouchers in this roll. The value can be from `0` to `1023` with the default number of **Ticket Bits**.

Note: If the count on an existing roll is changed, it will invalidate **all** other vouchers on the roll.

Comment
A description of the roll for reference, such as `2 hour vouchers for coffee purchases`.

Creating Voucher Rolls

To create a voucher roll:

- Use the pfSense® GUI to navigate to **Services > Captive Portal**
- Click  on the line for the Zone to edit
- Click the **Vouchers** tab
- Click  **Add** under the roll list
- Fill in the options as described in *Voucher Roll Options*
- Click **Save**

The new roll is available for immediate use by clients.

Editing Existing Rolls

To edit an existing voucher roll, click  at the end of its row, but be careful when making changes. Changing the **Roll** number or **Count** will invalidate the current vouchers on the roll.

Removing Voucher Rolls

To remove rolls of vouchers, click  at the end of their row. When a roll is removed, *all* the vouchers in that roll become invalid. Do not remove a roll unless it has been completely used, compromised, or otherwise unnecessary.

Exporting/Downloading Voucher Rolls

Click  to download a file containing the vouchers in the specified roll. This action downloads a `.csv` (Comma Separated Value) spreadsheet containing all voucher codes for this roll.

Nearly any spreadsheet editor can open this file, such as LibreOffice Calc, Google Docs, or Excel. Programs such as those can print vouchers, feed them into a POS system, and so on.

Using Vouchers on A Portal Page

The portal page must submit voucher codes via the `auth_voucher` form field. See *Portal page with Vouchers* for an example.

Viewing Active Vouchers

To view the list of currently active vouchers and their timers, navigate to **Status > Captive Portal**, on the **Active Vouchers** tab for a zone, as seen in Figure *Active Vouchers*.

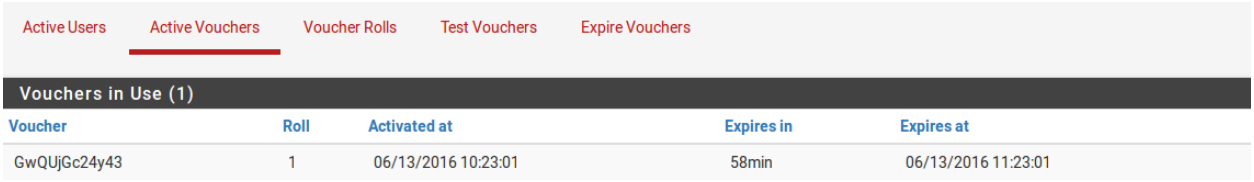


Fig. 1: Active Vouchers

Viewing Voucher Roll Utilization

To view a list of voucher rolls and usage counts, navigate to **Status > Captive Portal**, on the **Voucher Rolls** tab for a zone, as in Figure *Vouchers Roll Usage*

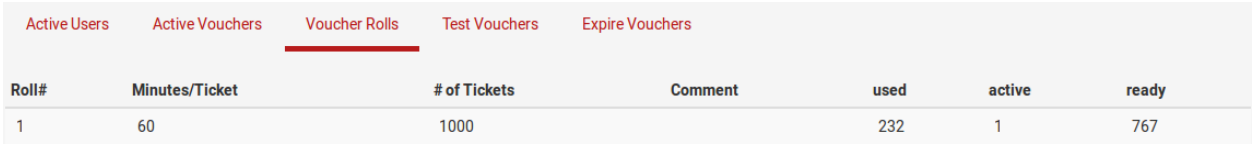


Fig. 2: Vouchers Roll Usage

Testing Vouchers

To test the validity of a voucher code, enter it at **Status > Captive Portal**, on the **Test Vouchers** tab for a zone. Upon submission, the page will display if a code is valid or not, and if it is valid, it will show the voucher time limit, as seen in Figure *Testing Vouchers*. Testing a voucher does not count it as used or expired, it is still free to be used at a later time by a client.

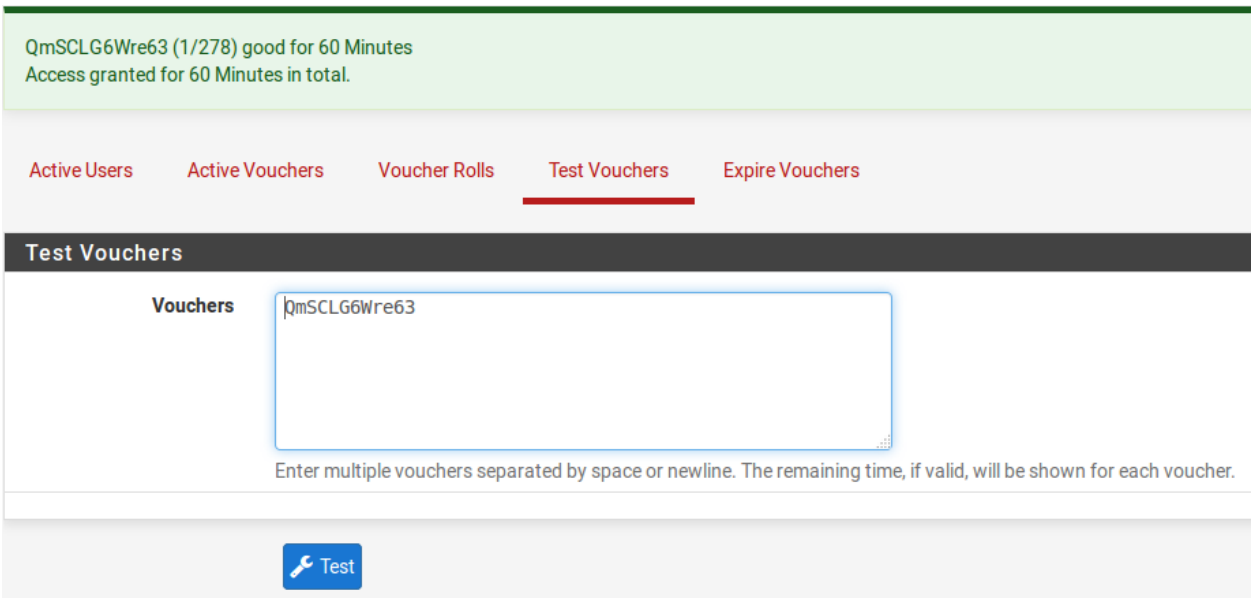


Fig. 3: Testing Vouchers

Expiring Vouchers

To invalidate vouchers, during or before use, enter them at **Status > Captive Portal**, on the **Expire Vouchers** tab for a zone. After submitting, any voucher listed in the form will no longer be allowed by the portal. Active vouchers entered on this page are also immediately expired.

The page can expire any number of vouchers in a batch. Enter vouchers separated by newlines.

26.7.4 Synchronizing Vouchers

When the High Availability option to synchronize Captive Portal data via XMLRPC is enabled on a primary node, pfSense software copies voucher data to the secondary node automatically.

If the primary node is offline, the secondary node may have new data which then must be synchronized back to the primary node. There are options to enable this "reverse" synchronization of voucher data to the primary HA node on the **High Availability** tab in the Captive Portal settings of each zone.

This works similarly to the XMLRPC configuration synchronization options in the high availability settings (*pfSense Software XMLRPC Config Sync Overview*). When active, this function copies the voucher rolls configured target node and also pushes information about active vouchers to the target node as vouchers are consumed by users.

Primary node IP
: The IP address of the primary HA node for voucher synchronization.

Primary node username
: The username for synchronization access. Must have appropriate privileges, for example, admin-level privileges or at least the `System - HA node sync` privilege.

Primary node password
: The password for the **Primary node username** account.

26.8 File Manager

The **File Manager** tab in a Captive Portal zone is used to upload files that can then be utilized inside a captive portal page, such as style sheets, image files, PHP or JavaScript files.

The total size limit for all files in a zone is **1 MB**.

26.8.1 File Name Conventions

When a file is uploaded using the File Manager, the file name will automatically be prefixed with `captiveportal-`. For example, if `logo.png` is uploaded it will become `captiveportal-logo.png`. If a file already has that prefix in its name, the name is not changed.

These files will be made available in the root directory of the captive portal server for this zone. The files may be referenced directly from the portal page HTML code using relative paths.

Example: An image with the name `captiveportal-logo.jpg` was uploaded using the file manager, It can then be included in the portal page as follows:

```
<img src="captiveportal-logo.jpg" />
```

PHP scripts may be uploaded as well, but they may need to be passed extra parameters to work as desired, for example:

```
<a href="/captiveportal-aup.php?zone=$PORTAL_ZONE$&redirurl=$PORTAL_REDIRURL$">
    Acceptable usage policy
</a>
```

26.8.2 Managing Files

To upload files:

- Navigate to **Services > Captive Portal**
- Edit the zone where the files will be uploaded
- Click the **File Manager** tab
- Click 
- Click **Browse**
- Locate and select the file to upload
- Click  **Upload**

The file will be transferred to the firewall and stored in the configuration.

To delete files:

- Navigate to **Services > Captive Portal**
- Edit the zone where the file to delete is located
- Click the **File Manager** tab
- Click  next to the file to remove
- Click **OK** to confirm the delete action

The file will be removed from the portal configuration and will no longer be available for use in portal pages.

See also:

- *Captive Portal Status*
- *Captive Portal Authentication Logs*
- *Troubleshooting Captive Portal*

Captive Portal in pfSense® software forces users on an interface to authenticate before granting access to the Internet. Where possible, the firewall automatically presents a login web page in which the user must enter credentials such as a username/password, a voucher code, or a simple click-through agreement.

This feature is commonly used throughout the hospitality industry (Hotels, restaurants, airports, and more) as well as in corporate and even home environments. It is primarily used for wireless hot spots or for additional authentication before allowing access to internal networks from wireless clients.

Captive Portal is configured under **Services > Captive Portal**.

See also:

Hangouts Archive to view the April 2015 Hangout on Captive Portal.

26.9 Limitations

The Captive Portal implementation in pfSense does have some limitations.

26.9.1 IPv6

Captive Portal is not compatible with IPv6.

26.9.2 Reverse Portal

A reverse portal, requiring authentication for traffic coming into a local network from the Internet, is not possible.

CHAPTER

TWENTYSEVEN

HIGH AVAILABILITY

27.1 High Availability Synchronization Settings

High Availability Synchronization settings for pfSense® software are located in the GUI at **System > High Availability**.

This document covers the settings on that page, but the general topics are covered in more detail throughout this chapter.

27.1.1 State Synchronization Settings (pfsync)

The settings in this section control the behavior of state synchronization and related functions. State synchronization allows firewalls acting as HA nodes to exchange state data so that all nodes in the cluster have knowledge about network connections.

The state synchronization settings should be enabled on **all** members of an HA cluster.

See also:

For details on how state synchronization operates, see *State Synchronization (pfsync) Overview*.

Synchronize States

Controls whether this firewall will perform state synchronization with other HA nodes on a shared segment.

When checked, the firewall will perform state synchronization on the **Synchronize Interface**.

Synchronize Interface

Controls which interface the firewall uses to send and receive state synchronization data with other HA nodes. This interface must have an IP address.

The best practice is to use an interface directly linked between the HA nodes, or at least connected through a switch using an isolated VLAN.

Warning: pfsync does not contain any method of authentication. If this option is set to anything other than an isolated segment it is possible for a user with access to the network on that interface to manipulate the state table. For example, they could insert states into the state table.

Filter Host ID

This option defines a custom PF host identifier carried in state data to uniquely identify which host created a firewall state.

Note: Each node participating in state synchronization must have a **different** filter host ID.

The host IDs from state data are shown on the *CARP status page* which allows administrators to check if an HA node is exchanging state data with other HA nodes.

Using a custom value is ideal but not required. On current versions of pfSense software the default is to use the last 8 characters of the host NDI. On previous versions the default behavior was to generate a randomized value on every filter reload.

The host ID value must be a non-zero hexadecimal string 8 characters or fewer (e.g. `1`, `2`, `ff01`, `abcdef01`).

pfsync Synchronize Peer IP

The IP address to which this firewall will send state synchronization data.

If left blank, the firewall will send state data using multicast to all hosts on the chosen **Synchronize Interface**.

In practice, state synchronization is more reliable when sent directly and not via multicast.

27.1.2 Configuration Synchronization Settings (XMLPRC Sync)

These settings control the behavior of XMLRPC configuration synchronization. XMLRPC configuration synchronization copies settings from compatible sections of the configuration from a primary node to a secondary node.

Warning: XMLRPC configuration synchronization must only be enabled on the primary node! It is not possible to synchronize settings from a secondary node back to the primary node.

Warning: The interfaces on both nodes **must** be assigned **identically**, for example: wan=WAN, lan=LAN, opt1=Sync, opt2=DMZ. Check the `config.xml` contents directly to ensure a match.

If the interfaces do not match up exactly, firewall rules and other configuration items will appear to synchronize to the wrong interface on the secondary node.

See also:

For details on how XMLRPC configuration synchronization operates, see *pfSense Software XMLRPC Config Sync Overview*.

Synchronize Config to IP

The IP address of the firewall to which this node will synchronize its configuration via XMLRPC.

There are a few requirements for this to work properly:

- The target firewall must be running the same version of pfSense software
- The target firewall GUI must be running the same protocol (HTTPS or HTTP)
- The target firewall GUI must be running on the same port (e.g. `443` or `80`)

Remote System Username

The username to use for authenticating against the target firewall.

The sync user must either be `admin` or an account on the target firewall with the *System - HA node sync* privilege.

Note: If XMLRPC is configured to synchronize users, create the sync user on the secondary manually first, as well as on the primary. The redundant copy on the secondary will be removed during the first successful synchronization, but the initial synchronization cannot succeed without that account present.

Remote System Password

The password to use for authenticating against the target firewall.

Synchronize Admin

Controls whether the primary node will synchronize its `admin` account to the target node.

By default, the XMLRPC process does not synchronize the `admin` account, which allows each HA node to have a different password for its `admin` account.

Note: When set, this option automatically updates **Remote System Password** when the password changes on the **Remote System Username** account.

Options to Synchronize

This part of the options is a list of configuration sections which XMLRPC configuration synchronization can copy to the target node. These sections include:

User manager users and groups
: Synchronizes users and groups defined in the user manager.

  If users have associations to certificates (e.g. for OpenVPN), then certificates should also be synchronized.

Authentication servers
: Synchronizes Authentication servers defined in the User Manager settings. For example, LDAP and RADIUS server entries and their settings.

  If these entries require SSL/TLS and are set to use a certificate, then certificates should also be synchronized.

Certificate Authorities, Certificates, and Certificate Revocation Lists
Synchronizes the contents of the Certificate Manager.

This replaces the entire contents of the certificate manager on the target node, which may also cause it to replace the GUI certificate. There are multiple methods to work around this, such as:

- Use the same GUI certificate on both nodes after performing an initial synchronization.
- Import the GUI cert for the secondary into the primary node, allow it to synchronize, and then re-select it on the secondary node.
- Create a new certificate on the primary node and then select it for use on the secondary after it synchronizes.

Tip: Certificates are synchronized when changed, but services depending on those certificates are not automatically restarted. When renewing certificates, services on the secondary which are running must be manually restarted. For example, if the GUI certificate is renewed, then the GUI must manually be restarted on the secondary node.

Firewall rules
Synchronizes the contents of all firewall rule tabs, including assigned interfaces, floating rules, interface groups, VPNs, etc.

If any firewall rules utilize aliases or schedules, those sections should also be set to synchronize.

Firewall schedules
Synchronizes defined firewall schedules.

Firewall aliases
Synchronizes the contents of aliases.

NAT configuration
Synchronizes the contents of NAT rules, including outbound NAT, port forwards, 1:1 NAT, etc.

IPsec configuration
Synchronizes the contents of IPsec tunnels.

If any IPsec tunnels use certificates for authentication, then certificates should also be synchronized.

OpenVPN configuration
Synchronizes the contents of all OpenVPN instances (clients and servers).

When enabled this also synchronizes the contents of the certificate manager as OpenVPN configurations require the use of certificates.

DHCP Server settings
Synchronizes the contents of the IPv4 DHCP server settings.

This synchronization process automatically adjusts the value of failover settings. See *DHCPv4 Server* for details.

DHCP Relay settings
Synchronizes the contents of the IPv4 DHCP relay settings.

DHCPv6 Server settings
Synchronizes the contents of the IPv6 DHCP server and Router Advertisement settings.

This synchronization process automatically adjusts the values of failover settings. See *DHCPv6 Server* for details.

> **Warning:** This is only viable when using the Kea DHCP server backend as the ISC DHCPv6 backend is not compatible with High Availability.

DHCPv6 Relay settings
Synchronizes the contents of the IPv6 DHCP relay settings.

WOL Server settings
Synchronizes the contents of Wake on LAN.

Static Route configuration
Synchronizes the contents of gateways and static routes.

Virtual IPs
Synchronizes the contents of Virtual IP addresses.

Different types of VIPs behave differently with regard to synchronization and some do not synchronize at all. See *Virtual IP Addresses* for details.

Traffic Shaper configuration
Synchronizes the contents of the ALTQ traffic shaper.

If firewall rules reference ALTQ traffic shaper queues, this should be enabled.

Traffic Shaper Limiters configuration
Synchronizes the contents of Limiters.

If firewall rules reference Limiters, this should be enabled.

DNS Forwarder and DNS Resolver configurations
Synchronizes the contents of the DNS Resolver and DNS Forwarder.

Captive Portal
Synchronizes the contents of Captive Portal, which includes additional exchanges of portal user and voucher usage data between HA nodes.

27.2 State Synchronization (pfsync) Overview

pfSense® software uses **pfsync** to synchronize firewall state table data between cluster nodes. Changes to the state table on the primary are sent to the secondary nodes over the Sync interface, and vice versa. When State Synchronization is active and properly configured, all nodes have knowledge of every connection flowing through the cluster. If the primary node fails, the secondary node will take over and most clients will not notice the transition since both nodes knew about the connection beforehand.

See also:

State Synchronization Settings (pfsync)

State synchronization with pfsync uses multicast by default, though an IP address can be defined to force unicast updates. This is ideal for environments with only two firewalls where multicast traffic is unnecessary and may not function properly. Any active interface can be used for sending pfsync updates, however utilizing a dedicated interface is the best practice for security and performance.

> **Warning:** pfsync does not contain any method of authentication. If the interface is set to anything other than an isolated segment it is possible for a user with access to the network on that interface to manipulate the state table. For example, they could insert states into the state table.

In low throughput environments that aren't security paranoid, use of the LAN interface for this purpose may be acceptable. Bandwidth required for this state synchronization will vary significantly from one environment to another, but could be as high as 10% of the throughput traversing the firewall depending on the rate of state insertions and deletions.

Failover can still operate without state synchronization, but it will not be seamless. Without state synchronization, if a node fails and another takes over, user connections are dropped. Users may immediately reconnect through the other node, but they would be disrupted during the transition. Depending on the usage in a particular environment this may go unnoticed, or it could be a significant, but brief, outage.

When state synchronization is in use, **State Synchronization** settings *must be enabled* on **all** nodes participating in state synchronization, including secondary node(s), or state synchronization will not function properly.

27.2.1 pfsync and Firewall Rules

Traffic for pfsync must be explicitly passed on the Sync interface. The rule must pass the *pfsync* protocol from a source of the *Sync network* to *any* destination. A rule passing all traffic of any protocol would also allow the required traffic, but a more specific rule is more secure.

27.2.2 pfsync and Physical Interfaces

States contain information about the interface to which they are bound. Whether or not this impacts pfsync depends on the default State Policy for the node, which can be "Interface Bound States" or "Floating States" (*Firewall State Policy*).

If the default state policy is Floating States and no rules are set to use Interface Bound States, then there is no conflict and state synchronization will work even if the hardware on the nodes is different.

If the default policy is set to Interface Bound States, or any rules are set to use Interface Bound States, then there may be a potential conflict with High Availability nodes which have different hardware.

If the interfaces are not both physically identical and assigned in the same order on both nodes then the states will not properly sync, for example if WAN is `ix0` on one node and `igb0` on the other.

While having identical hardware is always the best practice, mismatched hardware can still function with Interface Bound States by using LAGG interfaces to abstract the assignments. LAGGs can work around this since the states would be bound to the `laggX` interface on each node rather than the underlying physical interface. For example, `lagg0` on primary contains `ix0`, `lagg0` on secondary contains `igb0`, but the states use `lagg0` on both nodes, so sync will function.

27.2.3 pfsync and Upgrades

Normally pfSense software allows HA firewall upgrades without network disruption. Unfortunately, this isn't always the case with upgrades as the pfsync protocol can change to accommodate additional functionality. Always check the upgrade guide linked in all release announcements before upgrading to see if there are any special considerations for CARP users.

27.3 pfSense Software XMLRPC Config Sync Overview

Configuration synchronization makes it easier to maintain two nodes which are nearly identical. This synchronization is optional, but maintaining a cluster without it is significantly more work. Without synchronization, administrators would need to make every change multiple times and ensure the changes were consistent.

pfSense® software uses XMLRPC for configuration synchronization. When XMLRPC Synchronization is enabled, the primary node copies settings from compatible areas to the secondary node and activates them after each configuration change.

See also:

Configuration Synchronization Settings (XMLPRC Sync)

Certain configuration areas cannot be synchronized, such as the Interface configuration, but most other areas can: Firewall rules, aliases, users, certificates, VPNs, DHCP, routes, gateways, and more. See *Options to Synchronize* for a full list. As a general rule, items specific to hardware or a particular installation, such as Interfaces or values under **System > General** or **System > Advanced** do not synchronize. The list of compatible areas can vary depending on the version of pfSense software in use. For a list of areas that will synchronize, see the checkbox items on **System > High Availability** in the XMLRPC section. Most packages will not synchronize, but some contain their own synchronization settings. Consult package documentation for more details.

Configuration synchronization should use the Sync interface, or if there is no dedicated Sync interface, use the same interface configured for firewall state synchronization.

In a two-node cluster the XMLRPC settings must *only* be enabled on the primary node, the secondary node must have these settings *disabled*.

For XMLRPC to function, both nodes must meet the following requirements:

- The GUI must be running on the same port and protocol, for example: HTTPS on port `443`, which is the default setting.
- The interfaces **must** be assigned **identically** on both nodes, for example: wan=WAN, lan=LAN, opt1=Sync, opt2=DMZ. Check the `config.xml` contents directly to ensure a match.

  Warning: If the interfaces do not match up exactly, firewall rules and other configuration items will appear to synchronize to the wrong interface on the secondary node. Additionally, this can also lead to failures in DHCP failover.

- The sync user must either be `admin` or an account with the *System - HA node sync* privilege.

  Note: If XMLRPC will synchronize users, create the sync user on the secondary manually first, as well as on the primary. The redundant copy on the secondary will be removed during the first successful synchronization, but the initial synchronization cannot succeed without it.

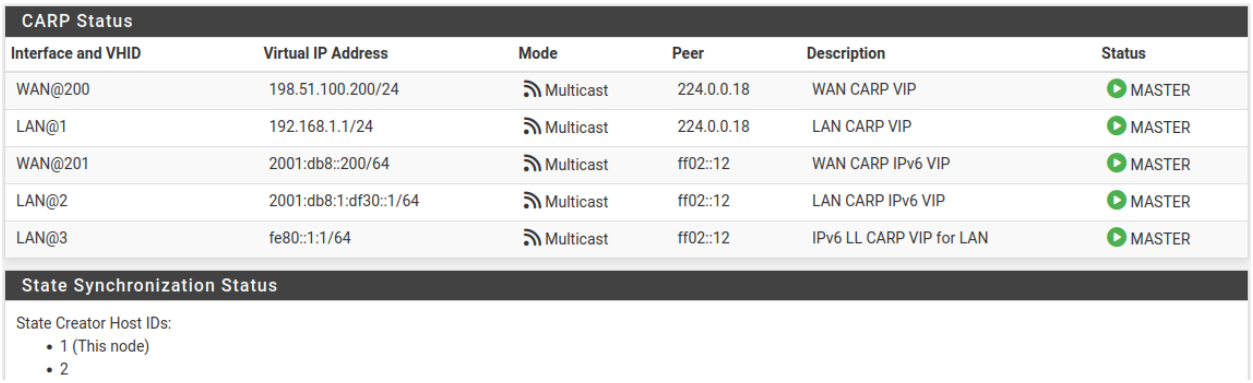
27.4 Verifying Failover Functionality

Since the goal of HA is high availability, thorough testing before placing a cluster into production is a must. The most important part of that testing is making sure that the HA peers will fail over gracefully during outages.

If any actions in this section do not work as expected, see *Troubleshooting High Availability*.

27.4.1 Check CARP status

On both nodes, navigate to **Status > CARP (failover)**. If everything is working correctly, the primary will show **MASTER** for the status of all CARP VIPs and the secondary will show **BACKUP**.



CARP Status

Interface and VHID	Virtual IP Address	Mode	Peer	Description	Status
WAN@200	198.51.100.200/24	Multicast	224.0.0.18	WAN CARP VIP	MASTER
LAN@1	192.168.1.1/24	Multicast	224.0.0.18	LAN CARP VIP	MASTER
WAN@201	2001:db8::200/64	Multicast	ff02::12	WAN CARP IPv6 VIP	MASTER
LAN@2	2001:db8:1:df30::1/64	Multicast	ff02::12	LAN CARP IPv6 VIP	MASTER
LAN@3	fe80::1:1/64	Multicast	ff02::12	IPv6 LL CARP VIP for LAN	MASTER

State Synchronization Status

State Creator Host IDs:
- 1 (This node)
- 2

Fig. 1: HA CARP and State Synchronization Status (Primary Node)

If either node shows **DISABLED**, click the **Enable CARP** button, then refresh the page.

If an interface shows **INIT**, it means the interface containing the CARP VIP does not have a link. Connect the interface to a switch, or at least to the other node. If the interface will not be used for some time, remove the CARP VIP from the interface as this will interfere with normal CARP operation.

27.4.2 Check State Synchronization

The **Status > CARP** page includes **State Synchronization Status** which lists Filter Host ID values for entries in the state table. If the Filter Host ID in the High Availability settings has been changed recently, it may show both old and new values from the primary and secondary nodes. Over time the list should only reflect the current values of the Filter Host ID of each node in the cluster.

If the lists are identical or nearly identical, then state synchronization is working. If the list does not contain an entry for the Filter Host ID of the other node, then states are not being synchronized.

27.4.3 Check Configuration Replication

Navigate to key locations on the secondary node, such as **Firewall > Rules** and **Firewall > NAT** and ensure that rules created only on the primary node are being replicated to the secondary node.

If the example earlier in this chapter was followed, the "temp" firewall rule on the pfsync interface would be replaced by the rule from the primary.

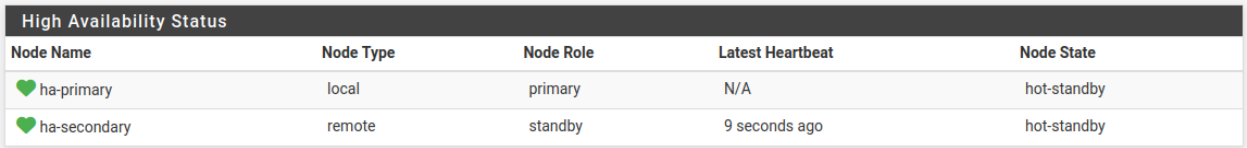
27.4.4 Check DHCP Failover Status

If DHCP failover was configured, its status can be checked at **Status > DHCP Leases** and **Status > DHCPv6 Leases**.

The exact appearance of the status depends on the DHCP service backend

Kea DHCP Failover Status

Failover status is in a section at the bottom of the DHCP and DHCPv6 Leases pages as in figure *Kea DHCP Failover Status - Primary note, both online.*. The failover status works identically for both DHCP and DHCPv6.



High Availability Status

Node Name	Node Type	Node Role	Latest Heartbeat	Node State
ha-primary	local	primary	N/A	hot-standby
ha-secondary	remote	standby	9 seconds ago	hot-standby

Fig. 2: Kea DHCP Failover Status - Primary note, both online.

See also:

See *High Availability Status – Kea DHCP Only* for more details about this section.

ISC DHCP Failover Status

Failover status is in a section at the top of the DHCP Leases page. This section contains the status of all DHCP Failover pools, as in Figure *ISC DHCP Failover Pool Status*.



Pool Status

Failover Group	My State	Since	Peer State	Since
dhcp_lan (LAN)	normal	2023/04/24 19:28:27	normal	2023/04/24 19:28:27

Fig. 3: ISC DHCP Failover Pool Status

See also:

See *Pool Status (HA/Failover) – ISC DHCP Only* for more details about this section.

27.4.5 Test CARP Failover

Now for the real failover test. Before starting, make sure that a local client behind the CARP pair on LAN can connect to the Internet with both nodes online and running. Once that is confirmed to work, it is an excellent time to make a backup.

For the actual test, unplug the primary node from the network or shut it down temporarily. The client will be able to keep loading content from the Internet through the secondary node. Check **Status > CARP (failover)** again on the secondary node, and it will now report that it is **MASTER** for the LAN and WAN CARP VIPs.

Now bring the primary node back online. It will regain its role as **MASTER**, and the secondary node will demote itself to a **BACKUP** state again. At any point during this process, Internet connectivity will still work properly.

Test the HA pair in as many failure scenarios as possible. Additional tests include:

- Unplug the WAN or LAN cable
- Pull the power plug of the primary
- Disable CARP on the primary using both the temporary disable feature and maintenance mode
- Test with each system individually (power off secondary, then power back on and shut down the primary)
- Download a file or try streaming audio/video during the failover
- Run a continuous ICMP echo request (ping) to an Internet host during the failover

27.5 IPsec in High Availability Environments

IPsec is compatible with high availability environments on pfSense® software.

27.5.1 CARP VIP as IPsec Endpoint

CARP type virtual IP addresses are available in the **Interface** drop-down menu on IPsec phase 1 configuration entries. In high availability environments, choose an appropriate CARP VIP address for the WAN where the IPsec tunnel will terminate.

Using a CARP VIP address ensures that the IPsec tunnel will be handled by the currently active High Availability cluster member. Even if the primary node is down, the tunnel will connect to whichever cluster member has taken over the active role.

27.5.2 XMLRPC Configuration Synchronization

The IPsec configuration, static routes (for route-based IPsec), and other similar settings will synchronize via XMLRPC if those functions are enabled on the node.

Warning: When using routed IPsec (VTI) with HA, the interface assignment for the `ipsecX` interface must be performed separately on both nodes. As with all other interfaces in a cluster they must be assigned in identical order.

27.5.3 Initiation Caveats

If the cluster attempts to automatically initiate a tunnel, the cluster member in a backup state may still transmit a message which may confuse the remote peer. pfSense software attempts to minimize the chances of this happening by dynamically setting nodes in a backup state to act as a responder only as well as disabling keep alive functions. When a node becomes active it re-enables these features.

During failover the far end of a tunnel may have to wait until it fully times out before it will rebuild the tunnel to the active cluster member. This process can take several minutes depending on tunnel configuration options (i.e. DPD). This may be faster if the cluster initiates, but depends upon the configuration, environment, and what triggered the failover.

Note: Additional workarounds are present on pfSense Plus software version 22.01 and CE version 2.6.0. On older versions it was easier for the backup node to unintentionally initiate the tunnel before it could be used which delayed failover.

27.6 Layer 2 Redundancy

The diagrams earlier in this chapter did not describe layer 2 (switch) redundancy, to avoid throwing too many concepts at readers simultaneously. This section covers the layer 2 design elements to be considered when planning a redundant network. This document assumes a two node deployment.

If both redundant nodes running pfSense® software are plugged into the same switch on any interface, that switch becomes a single point of failure. To avoid this single point of failure, the best choice is to deploy two switches for each interface (other than the dedicated Sync interface).

Example High Availability Cluster Network Diagram is network-centric, not showing the switch infrastructure. The Figure *Diagram of HA with Redundant Switches* illustrates how that environment looks with a redundant switch infrastructure.

27.6.1 Switch Configuration

When using multiple switches, the switches should be interconnected. As long as there is a single connection between the two switches, and no bridge on either of the firewalls, this is safe with any type of switch. Where using bridging, or where multiple interconnections exist between the switches, care must be taken to avoid layer 2 loops. A managed switch would be required which is capable of using Spanning Tree Protocol (STP) to detect and block ports that would otherwise create switch loops. When using STP, if an active link dies, e.g. switch failure, then a backup link can automatically be brought up in its place.

pfSense software can utilize *lagg(4)* link aggregation and link failover interfaces. This feature allows multiple network interfaces to be plugged into one or more switches for increased fault tolerance.

See also:

See *LAGG (Link Aggregation)* for more information on configuring link aggregation.

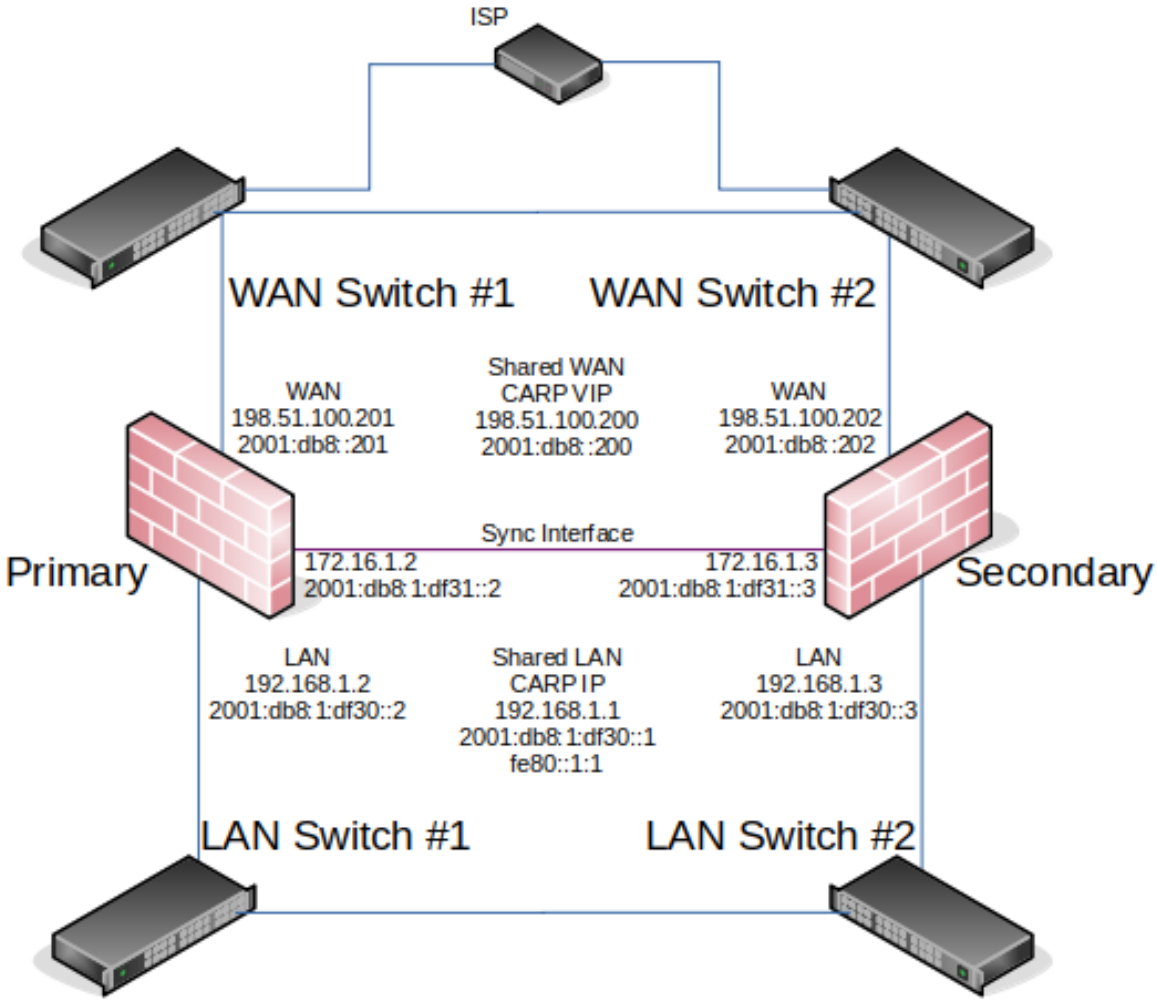


Fig. 4: Diagram of HA with Redundant Switches

27.6.2 Host Redundancy

It is more difficult to obtain host redundancy for critical systems inside the firewall. Each host could have two network cards and a connection to each group of switches using Link Aggregation Control Protocol (LACP) or similar vendor-specific functionality. Servers could also have multiple network connections, and depending on the OS it may be possible to run CARP or a similar protocol on a set of servers so that they would be redundant as well.

Providing host redundancy is more specific to the capabilities of the switches and server operating systems, which is outside the scope of this documentation.

27.6.3 Other Single Points of Failure

When trying to design a fully redundant network, there are many single points of failure that sometimes get missed. Depending on the level of uptime to achieve, there are multiple items to consider beyond a switch failure. Here are a few more examples for redundancy on a wider scale:

- Supply isolated power for each redundant segment
  - Use separate breakers for redundant systems
  - Use multiple UPS banks/generators
  - Use multiple power providers, entering opposite sides of the building where possible
- Even a Multi-WAN configuration is no guarantee of Internet uptime
  - Use multiple Internet connection technologies (Fiber, Cable, DSL, Wireless)
  - If any two carriers use the same pole/tunnel/path, they could both be knocked out at the same time
- Have backup cooling, redundant chillers or a portable/emergency air conditioner
- Consider placing the second set of redundant equipment in another room, another floor, or another building
- Have a duplicate setup in another part of town or another city
- I hear hosting is cheap on Mars, but the latency is killer

27.7 High Availability with Bridging

High availability is not currently compatible with bridging in a native capacity that is considered reliable or worthy of production use. It requires significant manual intervention. The details of the process can be found in *High Availability*.

27.8 Using IP Aliases to Reduce Heartbeat Traffic

If a segment contains numerous CARP VIPs, the segment can have a lot of multicast traffic. The firewall sends one heartbeat per second per CARP VIP. To reduce this traffic, additional VIPs of the same address family (IPv4 or IPv6) may be "stacked" on top of one CARP VIP of the same family on an interface.

- Pick one CARP VIP to be the "main" VIP for an interface for a given address family (IPv4 or IPv6)
- Edit the other CARP VIPs of the same family in the same subnet
  - Change the type to *IP Alias*
  - Select the "main" CARP VIP as the VIP **Interface**
  - Save

  - Repeat for each additional CARP VIP on the same interface
- Apply Changes

This not only reduces the heartbeats on a given segment, but it also causes all the IP alias VIPs to change status along with the "main" CARP VIP, reducing the likelihood that a layer 2 issue will cause individual CARP VIPs to not fail over as expected.

IP Alias VIPs do not normally synchronize via XML-RPC configuration synchronization, however, IP alias VIPs set to use CARP interfaces in this manner will synchronize.

See also:

- *Upgrading High Availability Clusters*
- *High Availability Configuration Example*
- *High Availability Configuration Example with Multi-WAN*
- *High Availability Configuration Example without NAT*
- *Troubleshooting High Availability*
- *Troubleshooting High Availability DHCP Failover*
- *Troubleshooting VPN Connectivity to a High Availability Secondary Node*

pfSense® software is one of very few open source solutions offering enterprise-class high availability capabilities with stateful failover, allowing the elimination of the firewall as a single point of failure.

High Availability on pfSense software is achieved through a combination of features:

- CARP for IP address redundancy
- XMLRPC for configuration synchronization
- pfsync for state table synchronization

With this configuration in place nodes act as an "active/passive" cluster with the primary node working as the active node and the secondary node in a backup "hot standby" style role, taking over as needed if the primary node fails.

Two or more redundant firewalls in this configuration are referred to as a "High Availability Cluster" or "HA Cluster".

Warning: This **should not** be called a "CARP Cluster". CARP is only one of several technologies used to achieve High Availability with pfSense software, and in the future CARP could be swapped for a different redundancy protocol.

One interface on each cluster node is dedicated for synchronization tasks. This is typically referred to as the "Sync" interface, and it is used for configuration synchronization, pfsync state synchronization, and some other function-specific synchronization tasks. Any available interface may be used for this role, but the best practice is to use a dedicated interface directly connected between the two nodes.

Warning: This **should not** be called a "CARP" interface as it is not involved with CARP. CARP heartbeats happen on each interface with a CARP VIP; CARP traffic and failover actions do not utilize the Sync interface.

The only supported High Availability cluster configuration consists of exactly two nodes, which is also the most common configuration in practice. While it is possible to have more than two nodes in a cluster, additional nodes do not provide a significant advantage. Configurations with more than two nodes are not officially supported.

It is important to distinguish between the three functions (IP address redundancy, configuration synchronization, and state table synchronization), because they happen in different places. Configuration synchronization and state synchronization happen on the sync interface, directly communicating between firewall nodes. CARP heartbeats are sent on each interface with a CARP VIP. Failover signaling does not happen on the sync interface, but rather it happens on every CARP-enabled interface.

See also:

Hangouts Archive which contains the June 2015 Hangout also covering High Availability.

27.9 CARP Overview

Common Address Redundancy Protocol (CARP) was created by OpenBSD developers as a free, open redundancy solution for sharing IP addresses among a group of network devices. Similar solutions already existed, primarily the IETF standard for Virtual Router Redundancy Protocol (VRRP). However, Cisco claims VRRP is covered by its patent on their Hot Standby Router Protocol (HSRP), and told the OpenBSD developers that it would enforce its patent. Hence, the OpenBSD developers created a new free, open protocol to accomplish essentially the same result without infringing on Cisco's patent. CARP became available in October 2003 in OpenBSD, and was later added to FreeBSD.

A CARP type *Virtual IP address* (VIP) is shared between nodes of a cluster. CARP refers to the currently active node as the "master". This node receives traffic sent to the CARP VIP address, and the other nodes sharing the VIP maintain "backup" status and monitor for heartbeats to see if they need to take over the active role. Since only one member of the cluster at a time is using the IP address, there is no IP address conflict between CARP VIPs.

For failover to work properly it is important that inbound traffic coming to the cluster, such as routed upstream traffic, VPNs, NAT, local client gateway, DNS requests, etc., be sent to a CARP VIP and for outgoing traffic such as Outbound NAT to be sent from a CARP VIP. If traffic is addressed to a node directly and not a CARP VIP, then that traffic will not be picked up by other nodes.

CARP works similar to VRRP and HSRP and may even conflict in some cases. Heartbeats are sent out on each interface containing a CARP VIP, one heartbeat per VIP per interface. At the default values for skew and base, each VIP sends out heartbeats about once per second. The skew determines which node is active at a given point in time. Whichever node transmits heartbeats the fastest assumes the active role. A higher skew value causes heartbeats to be transmitted with more delay, so a node with a lower skew will be active unless a network or other issue causes the heartbeats to be delayed or lost.

Note: Never access the firewall GUI, SSH, or other management mechanism using a CARP VIP directly. For management purposes, only use the actual IP address on the interface of each separate node and not the VIP. Otherwise, the client cannot determine beforehand which node it is accessing.

27.9.1 IP Address Requirements for CARP

A High Availability cluster using CARP needs **three** IP addresses in each subnet along with a separate unused subnet for the Sync interface. For WANs, this means that each WAN requires a /29 size subnet or larger for an optimal configuration. Each node uses One IP address, plus a shared CARP VIP address for failover. The synchronization interface only requires one IP address per node.

It is technically possible to configure an interface with a CARP VIP as the only IP address in a given subnet, but it is not generally recommended. When used on a WAN, this type of configuration will only allow communication from the primary node to the WAN, which greatly complicates tasks such as updates, package installations, gateway monitoring, or anything that requires external connectivity from the secondary node. It can be a better fit for an internal interface, however internal interfaces do not typically suffer from the same IP address limitations as a WAN, so it is still preferable to configure IP addresses on all nodes.

27.9.2 Switch/Layer 2 Concerns

In the default mode, CARP heartbeats utilize multicast and may require special handling on the switches involved with the cluster. Some switches filter, rate limit, or otherwise interfere with multicast in ways that can cause CARP to fail. Also, some switches employ port security methods which may not work properly with CARP.

At a minimum, to use multicast mode CARP VIPs the switch must:

- Allow Multicast traffic to be sent and received without interference on ports using CARP VIPs.
- Allow traffic to be sent and received using multiple MAC addresses.
- Allow the CARP VIP MAC address to move between ports.

Nearly all problems with CARP failing to properly reflect the expected status are failures of the switch or other layer 2 issues, so be sure the switches are properly configured before continuing.

If multicast mode is not viable, CARP VIPs may be configured in Unicast mode on pfSense Plus software. Unicast mode sends heartbeats to a single defined peer IP address (and vice versa). This mode works across any L3 connectivity through which the heartbeats can be transmitted between the peers without relying on multicast. However, use of unicast mode on traditional infrastructure where multicast is more suitable should be avoided. In unicast mode switches may flood packets for unicast CARP VIPs to all ports, leading to significant security and performance concerns.

See also:

VIP Configuration Options

CHAPTER

TWENTYEIGHT

SYSTEM MONITORING

The data and information that pfSense® software collects and displays is every bit as important as the services it provides. Sometimes it seems that commercial routers go out of their way to hide as much information as possible from users, but pfSense software can provide almost as much information as anyone could ever want (and then some).

This chapter contains a variety of methods for finding information about the firewall status, logs, traffic, hardware, and so on.

28.1 Status

These articles cover various ways to check the status of services or features of the firewall, or the firewall itself.

28.1.1 Dashboard

The main page of the firewall GUI is the **Dashboard**. The Dashboard page provides a wealth of information that can be seen at a glance, contained in configurable widgets. These widgets can be added or removed, and dragged around into different positions in multiple columns.

See also:

There are additional options which control dashboard behavior, such as the number of columns for widgets. These options are located at *System > General Setup* and are also available as per-user options in the user manager.

Managing Widgets

Each widget follows basic conventions for controlling its position, size, settings, and so on.

Adding and Removing Widgets

To add widgets, click the  button in the Dashboard controls area of the breadcrumb bar (*Dashboard Controls in the Breadcrumb Bar*.) This button displays the list of available widgets.

Inside the **Available Widgets** panel, click the name of a widget to add it to the Dashboard (See *Available Widgets List*). The dashboard will reload with the new widget displayed in one of its columns.



Fig. 1: Dashboard Controls in the Breadcrumb Bar

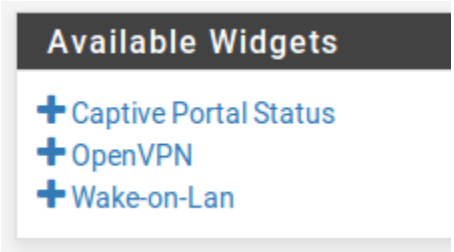


Fig. 2: Available Widgets List

To close and remove a widget from the Dashboard, click the button in its title bar (*Widget Title Bar*), then click in the dashboard controls.



Fig. 3: Widget Title Bar

Rearranging Widgets

Widgets can be rearranged and moved between columns.

To move a widget, click and drag its title bar (Figure *Widget Title Bar*), move the mouse to the desired position, and then release. The widget will "snap" into its new position as it moves, so the new location may be previewed before releasing the mouse button.

After positioning a widget, click in the dashboard controls to store its new location (*Dashboard Controls in the Breadcrumb Bar*).

Minimizing Widgets

To minimize a widget so it hides its content and only shows up as its title bar, click the button in its title bar (*Widget Title Bar*).

To restore the widget to its normal display, click the button.

After changing the widget status, click in the dashboard controls.

Changing Widget Settings

Some widgets have customizable settings which control how they display and update their content.

If a widget is capable of using custom settings, the  button will show up in its title bar as seen in Figure *Widget Title Bar*. Click that button and the settings for the widget will appear. Once the settings have been adjusted, click the **Save** button inside the widget settings panel.

Available Widgets

Each widget contains a specific set of data, type of information, graph, etc. This section lists each of the currently available widgets along with their settings (if any). These are listed in alphabetical order.

Captive Portal Status

This widget shows the current list of online captive portal users, including their IP address, MAC address, and username.

CARP Status

The CARP Status widget displays a list of all CARP type Virtual IP addresses, along with their status as either MASTER or BACKUP.

Disks

The Disks widget contains information on disk layout and usage. This content was formerly part of the System Information widget, but was moved to its own widget and redesigned.

The widget contains a tree view of the disks in the firewall, entries can be expanded to view details about additional ZFS datasets and mount points.

The Disk widget settings allow pinning specific items so the widget always their status. It also allows changing the usage threshold at which items are always shown, which can help identify disk locations which may need attention.

Dynamic DNS

The Dynamic DNS widget displays a list of all configured Dynamic DNS hostnames, their current address, and status.

Firewall Logs

The **Firewall Logs** widget provides an AJAX-updating view of the firewall log. The number of rows shown by the widget is configurable. As with the normal firewall log view, clicking the action icon next to the log entry will show a window displaying which rule caused the log entry. Clicking the source or destination IP address will copy that value to **Diagnostics > DNS** where the address can be resolved.

Gateways

The Gateways widget lists all the system gateways along with their current status. The status information consists of the gateway IP address, Round Trip Time (RTT) also known as delay or latency, the amount of packet loss, and the status (Online, Warning, Down, or Gathering Data). The widget is updated every few seconds via AJAX.

GEOM Mirror Status

This widget will show the status of a `gmirror` RAID array on the system, if one is configured. The widget will show if the array is online/OK (Complete), rebuilding, or degraded.

Installed Packages

The Installed Packages widget lists all the packages installed on the system, along with some basic information about them such as the installed version and whether an update is available.

When a package has an update available,  is displayed next to the version number. Packages may be updated from this widget by clicking the  button at the end of a package's row.

Packages may also be reinstalled by clicking  or removed by clicking  .

Interface Statistics

This widget shows a grid, with each interface on the system shown in its own column. Various interface statistics are shown in each row, including packet, byte, and error counts.

Interfaces

The **Interfaces** widget differs from the **Interface Statistics** widget in that it displays general information about the interface rather than counters. The **Interfaces** widget shows the type and name of each interface, IPv4 address, IPv6 address, the interface link status (up or down), as well as the link speed when available.

IPsec

The IPsec widget has three tabs:

Overview
: A count of active and inactive tunnels.

Tunnel Status
: Lists each configured IPsec tunnel (P1 and P2) and whether that tunnel is up or down. Each entry has controls to connect or disconnect based on its current status.

Mobile
: Shows online remote access IPsec VPN users, such as those using IKEv2 or Xauth.

Netgate Services and Support

Displays the current support status for this firewall instance from Netgate servers. The widget also includes information about support resources and how to contact Netgate TAC.

NTP Status

The **NTP Status** widget shows the current NTP synchronization source and the server time from that source.

OpenVPN

The **OpenVPN** widget displays the status of each configured OpenVPN instance, for both servers and clients. The status of each instance is shown, but the style and type of information shown varies depending on the type of OpenVPN connection. For example, with SSL/TLS servers in client/server mode the widget shows a list of all connected clients. For peer-to-peer mode instances such as shared key clients and servers, the widget displays an up/down status. In each case it displays the IP address of the connecting client with the name and time of the connection.

Picture

The **Picture** widget, as the name implies, displays a picture chosen by the user. This can either be used functionally, for a network diagram or similar, or it can be for style, displaying a company logo or other image.

To add an image:

- Click  on the Picture widget title bar
- Click **Browse** to locate the picture to upload
- Click **Upload** to upload the picture

The size of the picture will adjust to fit the area of the widget, which can vary depending on the size of the browser and platform.

RSS

The RSS (RDFSite Summary, or as it's often called, Really Simple Syndication) widget will display an arbitrary RSS feed. By default, it shows the Netgate blog RSS feed. Some people choose to show internal company RSS feeds or security site RSS feeds, but it can load any RSS feed.

In addition to defining the RSS feeds to display, the number of stories and size of displayed content are also configurable.

S.M.A.R.T. Status

If S.M.A.R.T. is enabled on a drive in the firewall, this widget will show a brief status of the drive integrity as reported by S.M.A.R.T.

Services Status

This widget provides the same view and control of services that appears under **Status > Services**. Each service is listed along with its description, status (Running, Stopped), and start/restart/stop controls.

System Information

This widget is the main widget, displaying a wide array of information about the running system. The information displayed includes:

Name
The configured fully qualified hostname of the firewall.

User
The user viewing the dashboard and their authentication source.

System
The type of system, if the firewall can identify the environment. Can be a specific hardware model, a type of virtual machine, or similar string.

This section also displays the **Netgate Device ID** (NDI) which is used by Netgate to determine the support status for the firewall.

BIOS
Information about the system BIOS, if it can be read by the firewall. May include the BIOS vendor, version, and release date.

Version
The current running version of pfSense® software. The widget displays the version, architecture, and build time at the top. Beneath that, the widget prints the underlying version of FreeBSD.

At the bottom of this section, the widget prints the result of an automatic update check for a more recent version of pfSense software. This automatic update check can be disabled in the update settings.

CPU Type
The version string for the processor, such as `Intel(R) Atom(TM) CPU C2758 @ 2.40GHz`. The widget also prints the CPU count and package/core layout.

If **powerd** is active and the CPU frequency has been lowered, then the current frequency is shown next to the maximum frequency.

If the CPU contains hardware cryptographic features, such as AES-NI or QAT, the widget also prints the status of those items.

Hardware crypto
If hardware cryptographic acceleration is enabled, the widget displays a list of ciphers which the hardware can accelerate.

Uptime
Time since the firewall was last rebooted.

Current date/time
The current date and time of the firewall, including the time zone. This is useful for comparing the log entries, especially when the time zone on the firewall is different from where the user resides.

DNS Server(s)
A list of all configured and automatically located DNS Servers used by the firewall.

Last config change
The date of the last configuration change on the firewall.

State table size
A graphical and numerical representation of active connection states and the maximum possible states as configured on the firewall. Underneath the state counts is a link to view the contents of the state table.

MBUF usage
The number of network memory buffer clusters in use, and the maximum the system has available. These network memory buffers are used for network operations, among other tasks. If the number is close to maximum or at the maximum, increase the number of available mbufs as described in *Hardware Tuning and Troubleshooting*.

Temperature
The current temperature as reported by the hardware, if available.

Load Average
A count of active processes on the firewall which are in a running state during the last 5, 10, and 15 minutes. This is typically `0.00` on an idle or lightly loaded system.

CPU usage
A bar chart and percentage of CPU time used by the firewall.

Note: Viewing the dashboard increases the CPU usage, depending on the platform. On slower platforms this is likely to read significantly higher than it would be otherwise.

Memory usage
The current amount of RAM in use by the system. Note that unused RAM is often allocated for caching and other tasks so it is not wasted or idle, so this number may show higher than expected even when the firewall is operating normally.

Swap usage
The amount of swap space in use by the system. If the system runs out of physical RAM, and there is swap space available, lesser used pages of memory will be paged out to the swap file on the hard drive. This indicator only shows when the system has swap space configured.

Thermal Sensors

The **Thermal Sensors** widget displays the temperature from compatible sensors when present. For many popular Intel and AMD-based chips, the sensors may be activated by choosing the appropriate sensor type under **System > Advanced** on the **Miscellaneous** tab under **Thermal Sensors**.

The widget displays a bar for each sensor, which typically corresponds to each CPU core. The warning and critical thresholds may be configured in the widget settings.

Traffic Graphs

The Traffic Graphs widget contains a live graph for the traffic on each interface. The interfaces displayed are configurable in the widget settings. The default refresh rate of the graphs is once every 10 seconds, but that may also be adjusted in the settings for this widget. The graphs are drawn the same way as those found under **Status > Traffic Graph**.

Wake On LAN

The Wake on LAN widget shows all the WOL entries configured under **Services > Wake on LAN**, and offers a quick means to send a WOL magic packet to each system in order to wake it up. The widget also displays the current status of WOL entries, if possible. To wake up a system, click  next to its entry.

ZFS

This widget is available on pfSense Plus software and displays current status of ZFS pools and their component disks.

28.1.2 Interface Status

The page at **Status > Interfaces** displays the status of each assigned network interface on the firewall.

Status Information

The information available for each interface varies depending on the interface type, but may include:

Header Information
The header of each interface contains the following information:

Friendly Name
The name of the interface as designated by the user (e.g. `DMZ`).

Internal Name
The internal name of the interface (e.g. `wan`, `opt1`).

Assigned Interface
The name of the underlying interface which was assigned to this entry (e.g. `ix2`).

Status
The current status of the interface along with an icon which visually represents the status. The status is typically one of:

Up
The interface is up, has a link, and is operating normally.

Associated
A wireless interface is up and associated with an access point.

Down
A dynamic WAN type is not connected or does not have an IP address.

No Carrier
Typically means that the cable is not plugged in or the device on the other end is malfunctioning in some way.

Disabled
The interface is assigned but manually disabled in the configuration.

Dynamic WAN Controls
Dynamic interfaces have a button to manually change their current state.

DHCP, DHCP6
Interfaces obtaining an IP address from DHCP or DHCPv6 have a **Release** button when there is an active lease, and a **Renew** button when there is not.

The DHCP **Release** action has an optional **Relinquish Lease** checkbox. When set, the release action also sends a special message to the DHCP server which relinquishes its current lease.

PPPoE, PPTP, L2TP, PPP
PPP-based connection types like PPPoE have a **Disconnect** button when connected and a **Connect** button when offline.

If a PPP connection is using dial-on-demand it will reconnect itself when triggered even after a manual disconnect action. Disable dial-on-demand if the interface must remain disconnected.

Note: Clicking 'Renew' or 'Connect' will trigger an attempt to take that action, but the attempt may fail if the interface has a problem or the upstream service is not responding.

PPP Uptime
PPP-based interfaces track how long they have been up based on the time the interface last connected to its upstream provider.

Cellular Status
PPP type interfaces connecting through cellular modems (e.g. 5G, LTE, 4G, 3G, etc.) may show additional statistics from the modem. This varies by hardware and type of connection.

Cell Signal (RSSI)
The signal strength from the cellular provider.

Cell Mode
System mode change indicator.

Cell SIM State
SIM card status (e.g. inserted or removed.)

Cell Service
Service mode change indicator.

Cell Upstream
Measured upload speed.

Cell Downstream
Measured download speed.

Cell Current Up
Maximum upload speed.

Cell Current Down
Maximum download speed.

MAC Address
The hardware MAC Address of the interface.

Tip: Installing the NMAP package activates a feature which allows the page to also display the manufacturer associated with the MAC address, if it is known. Note that this is not effective in some cases, such as for virtual machines which use randomly generated MAC addresses or for wireless clients which utilize privacy features that alter their MAC addresses.

IPv4 Address
The current IPv4 address assigned to the interface.

Note: This does not include Virtual IP addresses.

Subnet mask IPv4
The subnet mask for the current IPv4 address.

Gateway IPv4
The IPv4 gateway defined on this interface, if any.

IPv6 Link Local
The IPv6 link-local address for this interface, including the interface scope.

IPv6 Address
The current IPv6 address assigned to the interface.

Subnet mask IPv6
The length of the prefix for the current IPv6 address.

Gateway IPv6
The IPv6 gateway defined on this interface, if any.

DNS Servers
DNS servers obtained from upstream providers on this interface (e.g. DHCP or PPPoE).

MTU
maximum transmission unit (MTU) of this interface, which is the largest packet it can transmit or receive.

Media
The type of media connected to this interface, including the link speed and type. The exact values depend upon the network interface type and what is connected to that interface.

For example, it may be `1000baseT <full-duplex>` for some types of 1Gbit/s copper Ethernet or `10Gbase-SR <full-duplex>` for some types of 10Gbit/s fiber.

SFP Module Information
If the interface uses an SFP module and the operating system can read the data from the module, the page will also include that information. The data may include the following fields:

- Description (**Plugged**)
- Vendor
- Temperature
- Voltage
- Signal levels

Note: Some interfaces capable of using SFP modules, such as combination copper and SFP interfaces, do not expose this module data to the operating system. As such, the page cannot include

module data from these interfaces.

LAGG Information
If this interface is an assigned LAGG interface, the page displays information about the LAGG itself.

LAGG Protocol
The current protocol for LAGG, for example it could be `failover`, `loadbalance`, `LACP`, etc.

The current LAGG hashing method is also in this field, which typically is `l2,l3,l4` which indicates that when load balancing it takes information from OSI layers 2, 3, and 4 into account when deciding which port to use.

LAGG Ports
The underlying interfaces which are members of this LAGG, along with their current status.

Wireless Information
The page displays additional information specific to wireless interfaces as well, including:

Channel
The wireless channel the interface is using to communicate with peers.

SSID
When acting as an access point, this is the SSID being broadcast to clients.

BSSID
When acting as a wireless client, this is the SSID of the AP to which this interface is connecting.

Rate
When acting as a wireless client, this is the current wireless data transfer rate to the AP.

RSSI
When acting as a wireless client, this is the current signal level to the AP.

In/Out Packets
The number of packets received (in) and transmitted (out) by this interface.

In/Out Packets (Pass)
The number of packets PF has passed on this interface.

In/Out Packets (Block)
The number of packets PF has blocked on this interface.

In/Out Errors
Input and output errors on the interface. This is a total count and can be from a variety of causes. For example, it could be from a hardware issue or packets lost because they could not be processed due to high load.

Hardware issues are typically physical in nature: cabling or port errors. The most common suspect is cables, and they are easy and cheap to replace.

In many cases, occasional errors are not indicative of a problem, however, if the number is large and/or rapidly increasing, there is cause for concern.

Note: Depending on the interface type, more detail may be available from `sysctl`. For example, `ix0` interfaces have information under `sysctl dev.ix.0`, in particular `dev.ix.0.mac_stats` has several fields detailing different types of error conditions.

Collisions
The number of network collisions experienced by this interface.

In most cases this can only happen on half-duplex networks (i.e. hubs, not switches). If this is non-zero it can also indicate that the interface has not linked at the proper duplex.

See also:

See *Interface Configuration* for more about setting the speed and duplex of an interface.

Bridge Interface
If an interface is a member of a bridge, the title of this field contains the name of the bridge and the content is the current status.

Total Interrupts
For physical interfaces this field may show the total number of hardware interrupts generated by this interface. A rapidly increasing number of interrupts can indicate that an interface is highly loaded, but that does not necessarily mean there is a problem if the load is expected.

Example Status

In the first part of Figure *Interface Status*, the firewall has a DHCP WAN connection, and it obtained the IPv4 and IPv6 address, DNS, etc. automatically.

In the lower part of the image, the LAN connection is visible. Since this is a normal interface with a static IP address, only the usual set of items are present.

28.1.3 Service Status

The page at **Status > Services** displays the status of most base system and package service daemons.

The page lists each service with its name, description, and status as seen in Figure *Services Status*. The status is listed as **Running** or **Stopped**.

Normally it is not necessary to control services in this manner, but occasionally there are maintenance or troubleshooting reasons for doing so.

From this view, services can be controlled in various ways:

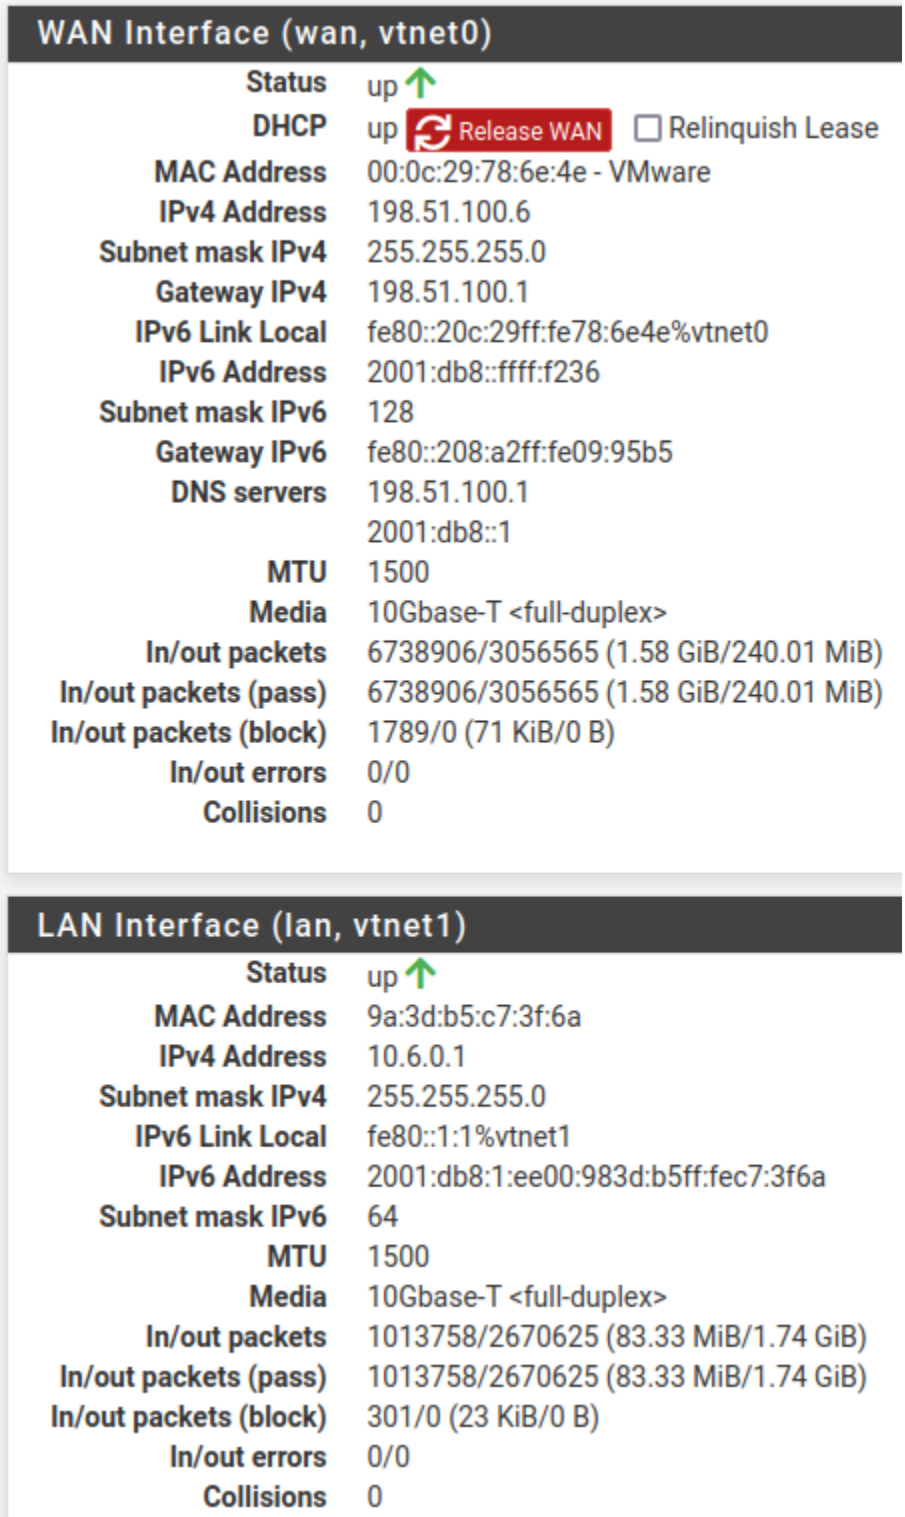
- Click  to restart a running service

---

Note: Some services will stop and start, others reload the configuration. Check the documentation of each service for details.

- Click  to stop a running service
- Click  to start a stopped service

If available, each entry also contains additional shortcuts which navigate to pages related to the service. See *Quickly Navigate the GUI with Shortcuts* for information about shortcut icons.



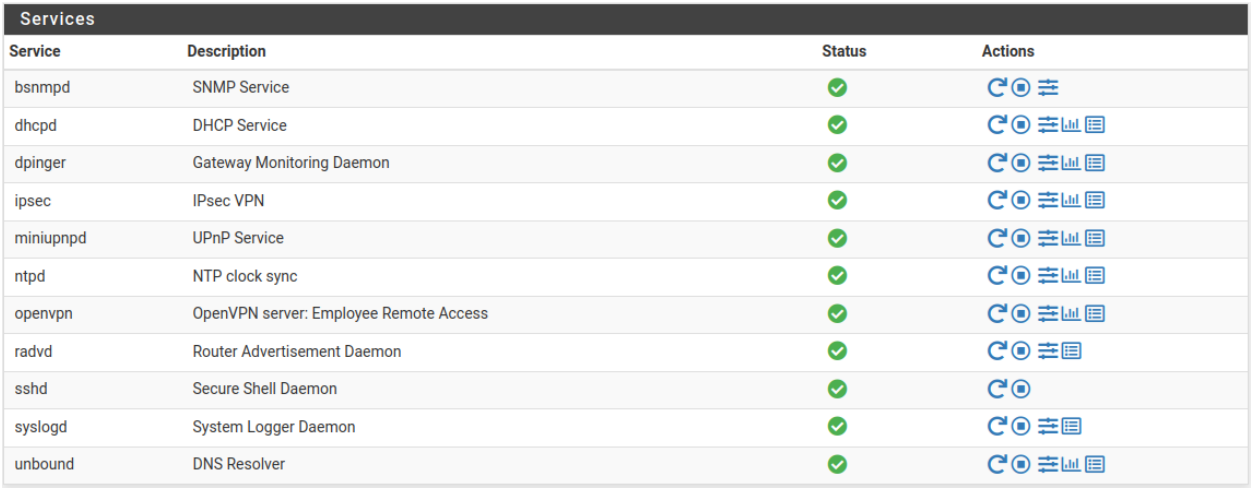
WAN Interface (wan, vtnet0)

Status	up
DHCP	up Release WAN Relinquish Lease
MAC Address	00:0c:29:78:6e:4e - VMware
IPv4 Address	198.51.100.6
Subnet mask IPv4	255.255.255.0
Gateway IPv4	198.51.100.1
IPv6 Link Local	fe80::20c:29ff:fe78:6e4e%vtnet0
IPv6 Address	2001:db8::ffff:f236
Subnet mask IPv6	128
Gateway IPv6	fe80::208:a2ff:fe09:95b5
DNS servers	198.51.100.1
	2001:db8::1
MTU	1500
Media	10Gbase-T <full-duplex>
In/out packets	6738906/3056565 (1.58 GiB/240.01 MiB)
In/out packets (pass)	6738906/3056565 (1.58 GiB/240.01 MiB)
In/out packets (block)	1789/0 (71 KiB/0 B)
In/out errors	0/0
Collisions	0

LAN Interface (lan, vtnet1)

Status	up
MAC Address	9a:3d:b5:c7:3f:6a
IPv4 Address	10.6.0.1
Subnet mask IPv4	255.255.255.0
IPv6 Link Local	fe80::1:1%vtnet1
IPv6 Address	2001:db8:1:ee00:983d:b5ff:fec7:3f6a
Subnet mask IPv6	64
MTU	1500
Media	10Gbase-T <full-duplex>
In/out packets	1013758/2670625 (83.33 MiB/1.74 GiB)
In/out packets (pass)	1013758/2670625 (83.33 MiB/1.74 GiB)
In/out packets (block)	301/0 (23 KiB/0 B)
In/out errors	0/0
Collisions	0

Fig. 4: Interface Status



Services

Service	Description	Status	Actions
bsnmpd	SNMP Service		
dhcpd	DHCP Service		
dpinger	Gateway Monitoring Daemon		
ipsec	IPsec VPN		
miniupnpd	UPnP Service		
ntpd	NTP clock sync		
openvpn	OpenVPN server: Employee Remote Access		
radvd	Router Advertisement Daemon		
sshd	Secure Shell Daemon		
syslogd	System Logger Daemon		
unbound	DNS Resolver		

Fig. 5: Services Status

28.1.4 System Activity (Top)

The **Diagnostics > System Activity** page displays several aspects of system activity as reported by `top` which are updated every few seconds. This is equivalent to running the command `top -aSH` at a shell prompt.

The output contains several key types of information:

- The system load average.
- A total count of processes in various states (running, sleeping, waiting, etc.).
- The current CPU usage.
- A breakdown of memory usage in various areas.

  See also:

  See *Memory Management* for details on each area.
- ZFS ARC usage breakdown, if the system contains any active ZFS pools.

  See also:

  See *ZFS Tuning* for information on reducing this if necessary.
- Swap memory usage.
- A list of the active processes running on the firewall.

Using this view it is easy to see processes that consume the most CPU power during a time of high load. For example, if the highest entry is an interrupt processing queue for one of the network cards, and the system isn't passing enough traffic, it could be one sign that the firewall is trying to push more than the hardware can handle in the current configuration. If the top process is a PHP process, it could be that a browser has requested a GUI page that is processing a large amount of data.

Note: Threads that show **idle** in the **COMMAND** column indicate CPU time that is **not** in use (idle). It is normal for these to show 100% if the firewall has little to no load.

28.1.5 pfInfo

The **Diagnostics > pfInfo** page displays statistics and counters for the firewall packet filter which serve as metrics to judge how it is behaving and processing data.

The **Refresh** checkbox at the top of the page controls whether the page automatically updates every few seconds with new data. To stop the updates, uncheck the box.

The information shown on the page contains items such as:

Host ID
The current 32-bit host ID used by PF. This value is randomized each time the filter reloads, and the value is stored on state table entries to indicate which process created the entry.

Bytes In/Out
Bytes transferred in and out of the firewall.

Packets In/Out
Packets transferred in or out and passed or blocked counters for each direction.

State Table / Source Tracking Table
Statistics about the state table and source tracking table (*Firewall States*).

Current Entries
The number of entries in the table

Searches
How many times the table has been searched and the current rate of searches, which roughly corresponds to the number of packets being passed by the firewall on current open connections.

Inserts
The number of new states added to the table, and the rate at which the states are added.

A high insert rate indicates that there are a lot of new connections being made to or through the firewall.

Removals
The number of old states being removed from the firewall.

Counters
Statistics and counts for various types of special, unusual or badly formatted packets.

Limit Counters
Counters that pertain to packets which have reached or exceeded limits configured on firewall rules, such as max states per IP address.

Table Size Limits
State table max size, source node table size, frag table size, number of allowed tables, and maximum number of table entries.

State Timers
The current configured timeout values for various connection states for TCP, UDP, and other protocols.

Interface Statistics
Per-interface packet counters.

28.1.6 S.M.A.R.T. Hard Disk Status

The firewall can monitor the health of hard drives that implement Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.). This mechanism is intended to allow drives to test and track their own performance and reliability, with the ultimate goal of identifying a failing drive before it suffers data loss or causes an outage.

S.M.A.R.T. availability varies by drive and BIOS, but it is fairly common in modern SSDs and hard drives. S.M.A.R.T. may need to be enabled in the BIOS and on the drive.

Note: S.M.A.R.T. is not a perfect metric of locating a failed drive; Many failed drives still pass S.M.A.R.T. tests. Generally speaking, a problem does exist if S.M.A.R.T. reports one, so it is useful to identify disk failures.

The **Diagnostics > SMART Status** page obtains and displays information from drives, performs or aborts drive tests, and displays drive logs.

In every section of the page, a **Device** must be selected before choosing an option. This **Device** is the disk to be tested by S.M.A.R.T.

Warning: If a drive is not listed in the **Device** list, it either does not have S.M.A.R.T. capabilities, or it is connected to an incompatible controller. In the case of RAID controllers, the controller itself may offer similar functionality or reporting via controller-specific utilities in the shell.

Viewing Drive Information

To view information about a drive:

- Navigate to **Diagnostics > SMART Status**
- Locate the **Information** panel on the page
- Select the **Device** to view
- Select the **Information Type**
- Click  **View**

After reviewing the output, click  **Back** to return to the list of options.

The information types are explained in the next subsections.

Device Information

The **Device Information** option shows information about the drive itself, including the make, model, serial number, and other technical information about the drive capabilities, connection, and operation.

```
Model Family:     Intel 53x and Pro 1500/2500 Series SSDs
Device Model:     INTEL SSDSC2BW120A4
Serial Number:    XXXXXXXXXXXX
LU WWN Device Id: 5 5cd2e4 0003ae43c
Firmware Version: DC32
User Capacity:    120,034,123,776 bytes [120 GB]
```

(continues on next page)

(continued from previous page)

```
Sector Size:      512 bytes logical/physical
Rotation Rate:    Solid State Device
TRIM Command:     Available, deterministic
Device is:        In smartctl database [for details use: -P show]
ATA Version is:   ACS-2 (minor revision not indicated)
SATA Version is:  SATA 3.0, 6.0 Gb/s (current: 6.0 Gb/s)
Local Time is:    Wed Feb  9 13:26:15 2022 EST
SMART support is: Available - device has SMART capability.
SMART support is: Enabled
```

Device Health

The **Health** option gives a brief pass/fail status of the drive.

```
SMART overall-health self-assessment test result: PASSED
```

SMART Capabilities

The **SMART Capabilities** choice gives a report about features and tests implemented in the drive firmware:

```
General SMART Values:
Offline data collection status:  (0x05)      Offline data collection activity
                                     was aborted by an interrupting command from host.
                                     Auto Offline Data Collection: Disabled.
Self-test execution status:      (  33)      The self-test routine was interrupted
                                     by the host with a hard or soft reset.
Total time to complete Offline
data collection:             ( 2930) seconds.
Offline data collection
capabilities:                         (0x7f) SMART execute Offline immediate.
                                     Auto Offline data collection on/off support.
                                     Abort Offline collection upon new
                                     command.
                                     Offline surface scan supported.
                                     Self-test supported.
                                     Conveyance Self-test supported.
                                     Selective Self-test supported.
SMART capabilities:            (0x0003)      Saves SMART data before entering
                                     power-saving mode.
                                     Supports SMART auto save timer.
Error logging capability:        (0x01)      Error logging supported.
                                     General Purpose Logging supported.
Short self-test routine
recommended polling time:     (   1) minutes.
Extended self-test routine
recommended polling time:     (  48) minutes.
Conveyance self-test routine
recommended polling time:     (   2) minutes.
SCT capabilities:           (0x0025) SCT Status supported.
                                     SCT Data Table supported.
```

SMART Attributes

The **SMART Attributes** view is the most useful screen in the majority of cases, but it can also be one of the trickiest to interpret. There are several values displayed but the number and values vary widely by make and model.

The following output is from a 2.5 inch traditional HDD:

```
=== START OF READ SMART DATA SECTION ===
SMART Attributes Data Structure revision number: 16
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG     VALUE WORST THRESH TYPE      UPDATED  WHEN_FAILED␣
↪RAW_VALUE
  1 Raw_Read_Error_Rate     0x000b   099   099   062    Pre-fail  Always       -       ␣
↪65537
  2 Throughput_Performance  0x0005   100   100   040    Pre-fail  Offline      -       0
  3 Spin_Up_Time            0x0007   136   136   033    Pre-fail  Always       -       2
  4 Start_Stop_Count        0x0012   100   100   000    Old_age   Always       -       96
  5 Reallocated_Sector_Ct   0x0033   100   100   005    Pre-fail  Always       -       0
  7 Seek_Error_Rate         0x000b   100   100   067    Pre-fail  Always       -       0
  8 Seek_Time_Performance   0x0005   100   100   040    Pre-fail  Offline      -       0
  9 Power_On_Hours          0x0012   061   061   000    Old_age   Always       -       ␣
↪17502
 10 Spin_Retry_Count        0x0013   100   100   060    Pre-fail  Always       -       0
 12 Power_Cycle_Count       0x0032   100   100   000    Old_age   Always       -       96
191 G-Sense_Error_Rate      0x000a   100   100   000    Old_age   Always       -       0
192 Power-Off_Retract_Count 0x0032   100   100   000    Old_age   Always       -       37
193 Load_Cycle_Count        0x0012   093   093   000    Old_age   Always       -       ␣
↪77869
194 Temperature_Celsius     0x0002   152   152   000    Old_age   Always       -       ␣
↪36 (Min/Max 19/41)
196 Reallocated_Event_Count 0x0032   100   100   000    Old_age   Always       -       0
197 Current_Pending_Sector  0x0022   100   100   000    Old_age   Always       -       0
198 Offline_Uncorrectable   0x0008   100   100   000    Old_age   Offline      -       0
199 UDMA_CRC_Error_Count    0x000a   200   200   000    Old_age   Always       -       0
223 Load_Retry_Count        0x000a   100   100   000    Old_age   Always       -       0
```

There is a thorough article on Wikipedia for S.M.A.R.T. that includes a guide for interpreting the values. Some values are more obvious than others, for example the counts for reallocated sectors should be at or near zero. Others can be harder such as the Raw Read Error Rate, which on most drives should be low, but there are Seagate and similar drives that output gibberish or a random high number in that field that makes it useless on those disks.

A few of the values are informational, such as the Start/Stop Count, Power Cycle Count, and Power On Hours which give a sense of the overall age and usage for the drive. A high value isn't necessarily bad for those, but if the drive is ancient, or has been power cycled a great many times, then have a plan prepared to replace the disk in the near future. The drive's Temperature can give an indication of its environment, and if the temperature is too high, it can lead to stability issues.

The Load Cycle Count is a special value for spinning disks, since it indicates the number of times the heads have been parked. Some laptop drives will automatically park the heads after a short time, but an OS like pfSense® software will want to write periodically, which brings the heads out again. The head parking only makes sense in a mobile device that moves a lot, so the heads have less chance of impacting the platter; In a server/firewall situation, it's completely unnecessary. Drives are only capable of 100,000-300,000 load cycles in their lifetime, which means the count gets run through quickly if the heads are continually parked and unparked. pfSense software attempts to disable the power management features of hard drives at boot time because otherwise the drive could fail prematurely after running this count up high. This cycling happening is typically audible on drives as a soft clicking noise.

To contrast the above, the following output is from an SSD:

```
=== START OF READ SMART DATA SECTION ===
SMART Attributes Data Structure revision number: 10
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG     VALUE WORST THRESH TYPE      UPDATED  WHEN_FAILED␣
↪RAW_VALUE
  5 Reallocated_Sector_Ct   0x0032   100   100   000    Old_age   Always       -       0
  9 Power_On_Hours_and_Msec 0x0032   100   100   000    Old_age   Always       -       ␣
↪40524h+33m+23.020s
 12 Power_Cycle_Count       0x0032   100   100   000    Old_age   Always       -       81
170 Available_Reservd_Space 0x0033   100   100   010    Pre-fail  Always       -       0
171 Program_Fail_Count      0x0032   100   100   000    Old_age   Always       -       0
172 Erase_Fail_Count        0x0032   100   100   000    Old_age   Always       -       0
174 Unexpect_Power_Loss_Ct  0x0032   100   100   000    Old_age   Always       -       18
183 SATA_Downshift_Count    0x0032   100   100   000    Old_age   Always       -       1
184 End-to-End_Error        0x0033   100   100   090    Pre-fail  Always       -       0
187 Uncorrectable_Error_Cnt 0x0032   100   100   000    Old_age   Always       -       0
190 Airflow_Temperature_Cel 0x0032   036   041   000    Old_age   Always       -       ␣
↪36 (Min/Max 20/41)
192 Power-Off_Retract_Count 0x0032   100   100   000    Old_age   Always       -       18
199 UDMA_CRC_Error_Count    0x0032   100   100   000    Old_age   Always       -       0
225 Host_Writes_32MiB       0x0032   100   100   000    Old_age   Always       -       ␣
↪978116
226 Workld_Media_Wear_Indic 0x0032   100   100   000    Old_age   Always       -       ␣
↪65535
227 Workld_Host_Reads_Perc  0x0032   100   100   000    Old_age   Always       -       2
228 Workload_Minutes        0x0032   100   100   000    Old_age   Always       -       ␣
↪65535
232 Available_Reservd_Space 0x0033   100   100   010    Pre-fail  Always       -       0
233 Media_Wearout_Indicator 0x0032   073   073   000    Old_age   Always       -       0
241 Host_Writes_32MiB       0x0032   100   100   000    Old_age   Always       -       ␣
↪978116
242 Host_Reads_32MiB        0x0032   100   100   000    Old_age   Always       -       ␣
↪3326
249 NAND_Writes_1GiB        0x0032   100   100   000    Old_age   Always       -       ␣
↪80031
```

The metrics for an SSD can be significantly different, as seen above. In particular, SSDs can give an estimate of their remaining lifetime, writes of various sizes, errors rates, write failures, and other SSD-specific values in place of the other values that do not apply to an SSD.

All SMART Information

Selecting **All SMART Information** shows all the information above and also includes the drive logs and self-test results.

All SMART and Non-SMART Information

This choice shows all the information above, plus more information that the drive reports. This includes alternate formatting for the attribute list with even greater detail about attribute meanings, a list of available SMART data logs, drive temperature details, additional device statistics, and disk event log content.

Drive Logs

The View Logs section displays the content of various drive logs. These logs contain information and errors, usually related to self-tests and potentially other errors encountered by the disk.

To view drive logs:

- Navigate to **Diagnostics > SMART Status**
- Locate the **View Logs** panel on the page
- Select the **Device** to view
- Select the **Log Type**
- Click  **View**

There are numerous logs available, but some logs are only present on specific types of devices, and some devices may not contain certain logs even if they are the correct type.

Summary Error Log
: The **Error** logs on a drive contain a record of errors encountered during the drive's operation, such as read errors, uncorrectable errors, CRC errors, and so on. Running an **Offline** test will also make the drive print more errors here if they are found during the test.

Extended Error Log
: Similar to the summary error logs, but allows for longer and more detailed error messages.

SMART Self-Test Log
: The **Self-test** logs contain a record of several recent self-tests run on the drive. It shows the type of test, the results of the test, and in the case of tests that were stopped prematurely, it shows the percentage of the test remaining.

  If an error is encountered during a test, the first logical block address (LBA) is printed to help determine where in the disk the problem lies.

Extended Self-Test Log
: Similar to the SMART self-test log but allows for longer and more detailed error messages.

Selective Self-Test Log
: Shows the results of recent selective self-tests and the min/max LBA sets which were included in the test.

Log Directory
: Prints the contents of the device log directory, which includes a list of logs and their current sizes.

Device Temperature Log (ATA Only)
The disk temperature information log from the SMART command transport. Prints both the current temperature and a temperature history with an ASCII graph.

Device Statistics (ATA Only)
Values and descriptions of ATA device statistics logged by the drive.

SATA PHY Events (SATA Only)
Values and descriptions of SATA PHY events logged by the drive.

SAS PHY Events (SAS Only)
Values and descriptions of SAS PHY events logged by the drive.

NVMe Log (NVMe Only)
Prints the contents of the NVMe drive log.

SSD Device Statistics (ATA/SCSI)
Prints either the device statistics or a media percentage used endurance indicator.

Drive Self-tests

To perform a test on a drive:

- Navigate to **Diagnostics > SMART Status**
- Locate the **Perform Self-Tests** panel on the page
- Select the **Device** to test
- Select the **Test Type**
- Click  **Test**

The types of tests are described in the following subsections.

Offline

An **Offline** test is called so because it is done while the disk is idle. This test can make accessing the drive slow while it is happening, but if there is a lot of disk activity, the drive may delay the test until the disk becomes idle again. Because of this variability, the exact time the test takes is hard to predict. An estimate of the time to complete an offline test for a given disk is shown in the **S.M.A.R.T. Capabilities**. An offline test will also cause the drive to update several of the S.M.A.R.T. attributes to indicate the results. After running a test and checking the results, review the **S.M.A.R.T. Attributes** again as well as the **Error** log.

Short

The **Short** test takes around ten minutes and checks the drive's mechanics and reading performance. A more accurate estimate of the length the test will take on a drive can be seen in the **S.M.A.R.T. Capabilities**. To see the results of this test, view the **Self-test** Logs. It can be run at any time and it does not typically impact performance.

Long

The **Long** test is similar to the **Short** test but is more thorough. The time taken by the test depends on the size of the disk, but it is much longer than the short test on its own. A more accurate estimate of the length the test will take on a drive can be seen in the **S.M.A.R.T. Capabilities**. As with the short test, the results end up in the **Self-test** Logs.

Conveyance

Not all drives implement this test. Its primary purpose is to test the drive after it has been physically relocated to determine if any components have been damaged by the move. In most cases it only takes a few minutes to complete. To determine if a drive is capable of a conveyance test, refer to the **S.M.A.R.T. Capabilities** output.

Canceling Active Tests

To cancel an active test on a drive:

- Navigate to **Diagnostics > SMART Status**
- Locate the **Abort** panel on the page
- Select the **Device** currently running a test
- Click  **Abort**

Any active tests on the drive will be stopped.

28.1.7 Filter Reload Status

The page at **Status > Filter Reload** displays the current status of a filter reload. It also contains controls to perform a manual reload, and a button to force an XMLRPC configuration synchronization attempt.

After making a filter change in the GUI and applying changes, the information box on the page contains a link to this page to monitor the filter reload progress.

Reload Status

This section of the page displays the progress of the reload, and it is updated automatically as the process progresses. New messages are placed at the bottom of the output, so scroll down to read the most recent entries.

In most cases updates happen fast enough that `Done.` is the only message shown, indicating that there are no pending changes. With larger configurations, some delays are possible.

Reload Filter

The **Reload Filter** button initiates a manual filter reload. As the process progresses, the page automatically updates with the most recent output.

XMLRPC Synchronization

If this firewall is part of a high availability cluster and has XMLRPC configuration synchronization enabled, the page displays a **Force Config Sync** button.

The **Force Config Sync** button triggers a forced synchronization of the firewall configuration by XMLRPC to the node specified by the *current XMLRPC synchronization configuration*.

As with a typical filter reload, the page updates automatically as the process progresses.

28.1.8 Firewall States

pfSense® software is a *stateful firewall* and uses one state to track each connection to and from the firewall. These states may be viewed in several ways in the GUI and from the console.

Viewing Firewall States in the GUI

The page at **Diagnostics > States** in the GUI contains a listing of the firewall state table contents. Figure *Example States* shows a sample of the output displayed by the GUI.

WAN	tcp	198.51.100.6:37246 -> 162.208.119.39:443	TIME_WAIT:TIME_WAIT	91 / 89	6 KiB / 120 KiB	
LAN	tcp	10.6.0.114:49266 -> 52.88.223.32:443	ESTABLISHED:ESTABLISHED	248 / 244	18 KiB / 20 KiB	
WAN	tcp	198.51.100.6:55087 (10.6.0.114:49266) -> 52.88.223.32:443	ESTABLISHED:ESTABLISHED	248 / 244	18 KiB / 20 KiB	
WAN2	icmp	203.0.113.106:36339 -> 203.0.113.1:36339	0:0	832.827 K / 921	22.26 MiB / 25 KiB	

Fig. 6: Example States

The page displays several columns of information for each state table entry, each with important information:

Interface
: The interface to which the state is bound. This is the interface through which the packet *initially* entered or exited the firewall.

Protocol
: The protocol of the traffic that created the state, such as TCP, UDP, ICMP, or ESP.

Source and Destination
: This column is in two parts, first the source, then an arrow indicating direction, and then the destination.

  The source and destination may also have a port number listed if the protocol in question uses ports.

In cases where NAT is applied (outbound NAT, port forwards, or 1:1 NAT), the address is shown both before and after NAT has been applied. For NAT such as outbound NAT which translates the source, the source section displays the translated source and the original source inside parenthesis. For NAT types that translate the destination, such as port forwards, the destination section shows the translated destination and the original destination in parentheses.

State
The current status of the connection being tracked by this state entry.

The specific values vary depending on the protocol. For example, TCP has many more state types than UDP or other connectionless protocols. See *Interpreting States* for more detail.

The entry in this column contains two parts separated by a colon. The first part is the connection state for the source side, and the second part is the connection state for the destination side.

Packets
The number of packets the firewall has observed which match the state from the source and destination sides.

Bytes
The total size of packets the firewall has observed which match the state from the source and destination sides.

The  icon at the end of each row removes individual states.

See also:

Firewall Maximum States

Filtering States

The **State Filter** panel enables quick searching of the state table contents to find items of interest.

To search for a state:

- Select a specific **Interface** in the **State Filter** panel or leave it on *all* to match all interfaces.
- Enter a **Filter Expression** which is a simple string of text to match exactly in the entry. This field does not perform regular expression matching.
- Click  **Filter** to locate the results.

The search process attempts to match text across all columns, and it only displays entries matching the search parameters.

Tip: Searching for an IP address or subnet will also present a  **Kill States** button which, when clicked, will remove all states originating from or going to the entered IP address or subnet.

For systems with extraordinarily large state tables, filtering can be a requirement. The **Require State Filter** option on *General Configuration Options* prevents the firewall from displaying state table content until the user filters the state table output. This prevents the page from attempting to display too much data.

Interpreting States

The **State** column for each state table entry provides information necessary to determine exactly what is happening with the connection. Each state entry contains two values with a colon between them, marking which value represents the state of the source (left), and which represents the destination (right).

A few of the most common state types are:

SYN_SENT
: For TCP connections this indicates that the side showing this state sent a TCP SYN packet attempting to start a connection handshake.

CLOSED
: For TCP connections the side with this status considers the connection closed, or the firewall has not received any traffic.

ESTABLISHED
: A TCP connection is considered fully established by this side.

TIME_WAIT/FIN_WAIT
: A TCP connection is in the process of closing and finishing up.

NO_TRAFFIC
: The firewall has not received any packets that match the state from this side.

SINGLE
: The firewall has observed a single packet on this state from this side.

MULTIPLE
: The firewall has observed multiple packets on this state from this side.

Common pairings frequently found in the state table include:

ESTABLISHED:ESTABLISHED
: A fully established two-way TCP connection.

SYN_SENT:CLOSED
: The side showing *SYN_SENT* has sent a TCP SYN packet but no response has been received from the far side. Often this is due to the packet not reaching its destination, or being blocked along the way.

SINGLE:NO_TRAFFIC
: Similar to the above, but for UDP and other connectionless protocols. No response has been received from the destination side.

SINGLE:MULTIPLE
: For UDP and other connectionless protocols, commonly observed with DNS where the client sends one packet but receives a large response in multiple packets.

MULTIPLE:MULTIPLE
: For UDP and other connectionless protocols, there are multiple packets in both directions, which is normal for a fully operational UDP connection.

0:0
: Indicates that there is no state level data. This state is typically only found on ICMP states, since ICMP does not have state levels like other protocols.

Firewall States Summary

The **State Table Summary**, accessible from **Diagnostics > States Summary**, provides statistics generated by an in-depth analysis of the state table and the connections therein.

The report includes the IP address, a total state count, and breakdowns by protocol and source/destination ports. Hovering over the ports shows a tooltip display of the full port list instead of the total number of ports. Depending on the firewall environment, high values by any metric may be normal.

The report includes the following categories:

By Source IP Address
: States summarized by the source IP address. This is useful for finding a potential source of attack, or a port scan or similar type probe/attack.

By Destination IP Address
: States summarized by the destination IP address of the connection. Useful for finding the target of an attack or identifying servers.

Total per IP Address
: States summarized by all connections to or from an IP address. Useful for finding active hosts using lots of ports, such as BitTorrent clients.

By IP Address Pair
: Summarizes states between two IP addresses involved in active connections. Useful for finding specific client/server pairs that have unusually high numbers of connections.

Warning: The **States Summary** can take a long time to process and display its output. This is especially true when the state table is large or the hardware is slow.

In cases where the state table is extremely large, the page may not display properly, or the page may fail with a memory error. In these cases, the summary page cannot be used.

Source Tracking States

When using *Sticky Connections* the firewall maintains a source tracking table which records mappings of internal IP addresses to specific external gateways for connections that were passed by a rule utilizing a Load Balancing gateway group (Multiple gateways on the same tier).

Source tracking associations are displayed by the page at **Diagnostics > States** on the **Source Tracking** tab, which is only visible if the sticky connections option is enabled.

By default, these source tracking associations only exist so long as there are active states from the internal IP address. There is a configurable timeout for these source tracking entries to allow them to exist longer if necessary.

See also:

For additional information about sticky connections, see *Sticky Connections*.

The Source Tracking page lists the following information:

Source-to-Destination
: The mapping of a local IP address to a specific load balanced **gateway**.

 # States
: The number of states matching this source IP address to any destination, including traffic that is not load balanced.

 # Connections
: The number of states matching this source IP address which utilize the gateway. For example, connections leaving from this source to an Internet host.

Rate
: The rate of packets matching this source tracking entry.

The **Remove** button at the end of each row will remove these associations individually.

Reset State Table / Source Tracking Table

The page at **Diagnostics > States** on the **Reset States** can reset the contents of the state table and source tracking table.

Certain situations call for resetting the state table to force all existing connections to reestablish. The most notable examples are making changes to NAT rules, firewall block rules, or traffic shaping. When these types of changes are made, resetting the state table is the only way to make sure all connections respect the new ruleset or traffic shaping queues.

> **Warning:** Resetting the state table is disruptive, but clients may immediately reconnect provided the current firewall rules still pass their connections.

Both the state table and the source tracking table may be reset as follows:

- Navigate to **Diagnostics > States**, **Reset States** tab
- Check **State Table** to clear the contents of the state table
- Check **Source Tracking** to clear the contents of the source tracking table
- Click  **Reset**
- Click **OK** to confirm

After confirming the action the firewall will erase the contents of the state table.

> **Warning:** The browser will appear to lose connection with the firewall when resetting the state table. Once the browser realizes the old connection is invalid, it will reconnect. Close and reopen the browser to reconnect faster.

pfTop

pfTop is available from the GUI and the console menu. It offers live views of the firewall ruleset, state table information, and related statistics.

pfTop in the GUI

The GUI page for pfTop is at **Diagnostics > pfTop**. The GUI offers several options to control the output:

View
Controls the type of output displayed by pfTop. Not all views will contain meaningful information for every firewall configuration.

Default
Shows a balanced amount of information, based around the source and destination of the traffic.

Label
Centered around firewall rule descriptions.

Long
Similar to the default view, but tailored for wider displays with longer rows for more columns of information. Shows the gateway after the destination.

Queue
Shows the ALTQ traffic shaping queues and their usage.

Rules
Shows firewall rules and their usage.

Size
Shows states that have passed the most data.

Speed
Shows states that have high-rate traffic.

State
Shows status of states.

Time
Shows long-lived states.

Filter Expression
An expression used to match groups of states to include in the output.

The expression can include several different types of filtering, such as:

- Filter by protocol: `proto <ip|ip6|ah|carp|esp|icmp|ipv6-icmp|pfsync|tcp|udp>`
- Filter by address: `[src|dst|gw] [host|net|port] <host/network/port>`
- Filter by direction: `[in|out]`

Sort By
Some views can be sorted. When sorting is possible, the following sort methods are available. When selected, the view is sorted by the chosen column in descending order:

None
No sorting, the natural order shown by the chosen view.

Age
The age of the states.

Bytes
The amount of data sent matching the state.

Destination Address
The destination IP address of the state.

Destination Port
The destination port number of the state.

Expiry
The expiration time of the state. This is the countdown timer until the state will be removed if no more data matches the state.

Peak
The peak rate of traffic matching a state in packets per second.

Packet
The number of packets transferred matching a state.

Rate
The current rate of traffic matching a state in packets per second.

Size
The total amount of traffic that has matched a state.

Source Port
The source port number of the state.

Source Address
The source IP address of the state.

Maximum # of States
On views capable of sorting, this option limits the number of state entries shown on the page.

pfTop on the Console

To access pfTop from the console or via ssh use option `9` from the menu or run `pftop` from a shell prompt.

While viewing pfTop in this way, there are several methods to alter the view while watching its output.

The most common options are:

- Press `h` to see a help screen that explains the available choices.
- Press `0` through `8` to select different views
- Press `space` for an immediate update
- Press `q` to quit

See the previous section for details on the meaning of the available views and sort orders.

The output is dynamically sized to the terminal width, with wider terminals showing much more information in additional columns.

28.1.9 DHCPv4 Status

The current contents of the DHCPv4 lease database and related information are viewable at **Status > DHCP leases**.

The page contains multiple sections with information about leases. The exact information and layout depends on the selected DHCP backend.

Pool Status (HA/Failover) – ISC DHCP Only

The **Pool Status** section of the page is only present if the firewall is using the ISC DHCP backend and is configured for DHCPv4 failover as a part of a high availability cluster.

This section of the page includes information on DHCP failover pools, including:

Failover Group
The name of the failover group for a given pool, which includes the interface name.

My State
The state of the failover pool from the perspective of this firewall.

Since
The last time the local pool state changed.

Peer State
The state of the failover pool from the perspective of the peer, if known.

Since
The last time the peer pool state changed.

See also:

Troubleshooting High Availability DHCP Failover

DHCPv4 Leases

This section of the page lists client leases and their properties, including:

Status
The first column with no header contains two icons indicating the current status of the lease. Hover over the icon for a tooltip explaining the meaning of the icon.

Lease Type
An icon indicating the type of DHCP lease assigned to this client, which can be one of:

Static
This lease entry is from a static DHCP mapping entry.

Active
A current lease from an active client.

Expired
A lease which has expired because a client did not renew it before its expiration time.

Online
An icon indicating whether the device is currently "online" as determined by the contents of the *ARP table*.

If a system shows online, then it has recently tried to communicate to or through this firewall.

A host marked as offline may be powered on and working, but it has not attempted communication to or through the firewall recently.

IP Address
IP address assigned to the client by DHCP, or static mapping address.

MAC Address
The client MAC address.

Tip: Installing the NMAP package activates a feature which allows the page to also display the manufacturer associated with the MAC address, if it is known. Note that this is not effective in some cases, such as for virtual machines which use randomly generated MAC addresses or for wireless clients which utilize privacy features that alter their MAC addresses.

Hostname
The hostname (if any) that the client sent as part of its DHCP request.

Description
The description for a host with a DHCP static mapping.

Start/End
The beginning and end times of the DHCP lease.

Note: For static mappings the page prints `n/a` as static mapping leases do not have a start or end time, and they do not expire.

Actions
Icons to take action on this lease. See *Actions* for details.

Search

The search box filters the contents of the **Leases** table based on keyword matching.

Enter a search string or UNIX regular expression into the box and click **Search** to filter the list to only matching records.

By default, the search looks at text from all fields in the lease record, but this can be limited to specific fields using the drop-down list.

Actions

Add static mapping

To create a static mapping from a dynamic lease, click  the to the right of the lease. This pre-fills the MAC address of that host into the **Edit static mapping** screen.

Edit static mapping

Entries for existing static mappings have the  icon which takes the user to the page to edit that specific entry.

Wake on LAN Integration

Clicking the icon to the right of the lease sends a Wake on LAN (WOL) packet to that host.

Click to create a WOL entry for the MAC address.

See also:

Wake on LAN

Delete a lease

While viewing the leases, an expired or inactive lease may be manually deleted by clicking at the end of its line. This option is not available for active or static leases, only for offline or expired leases.

Pool Usage Summary

The **Lease Utilization** section summarizes pool usage, giving a count of leases used in each pool configured in the DHCPv4 server.

View inactive leases

By default, the page lists active and static leases. Clicking **Show all configured leases** makes the page display all leases, including inactive and expired leases.

To reduce the view back to normal, click **Show active and static leases only**.

Clear All DHCP Leases

The **Clear all DHCP leases** button stops the DHCP daemon, removes the entire lease database, and then starts the daemon again.

This does not remove static DHCP leases, only dynamic leases.

High Availability Status – Kea DHCP Only

Failover status for Kea DHCP is in a section at the bottom of the DHCP and DHCPv6 Leases pages as in figure *Kea DHCP Failover Status - Primary node, both online*. The failover status works identically for both DHCP and DHCPv6.

High Availability Status

Node Name	Node Type	Node Role	Latest Heartbeat	Node State
ha-primary	local	primary	N/A	hot-standby
ha-secondary	remote	standby	9 seconds ago	hot-standby

Fig. 7: Kea DHCP Failover Status - Primary node, both online

The failover section contains a row for each node with the status of each node using the following fields:

Node Name
An icon indicating the status of a node followed by the configured name for each node.

The possible icons are:

-  : The node is online and has responded recently.
-  : The node is in an interrupted state, indicating it has not responded recently and may be going offline.
-  : The node is in an offline state because it is unresponsive.

Node Type
Indicates whether this entry is local (this node) or remote (the peer).

Node Role
The selected role for this node (primary or standby).

Latest Heartbeat
Elapsed time since this node received a heartbeat from the peer.

Node State
The current state of the node, which may be one of the following:

waiting
The node is waiting for a connection from the peer.

syncing
The node has established a connection with the peer and is synchronizing lease data.

ready
The node has finished synchronizing lease data and is ready to serve clients.

hot-standby
On a primary node this indicates the node is handing out leases for local clients, coordinating lease data with the secondary node, and sending failover heartbeats.

On a standby node this indicates the node is receiving lease data from the primary node, and sending failover heartbeats, but it is not handing out leases.

unavailable
The node is not responding and is considered offline.

partner-down
A standby node has assumed an active role in handing out leases for clients since the primary peer is offline.

28.1.10 DHCPv6 Status

The current contents of the DHCPv6 lease database and related information are viewable at **Status > DHCPv6 leases**.

The page contains multiple sections with information about leases. The exact information and layout depends on the selected DHCP backend.

DHCPv6 Leases

This section of the page lists DHCPv6 client leases and their properties, including:

Status
The first column with no header contains two icons indicating the current status of the lease. Hover over the icon for a tooltip explaining the meaning of the icon.

Lease Type
An icon indicating the type of DHCP lease assigned to this client, which can be one of:

Static
This lease entry is from a static DHCP mapping entry.

Active
A current lease from an active client.

Expired
A lease which has expired because a client did not renew it before its expiration time.

Online
An icon indicating whether the device is currently "online" as determined by the contents of the *ARP table*.

If a system shows online, then it has recently tried to communicate to or through this firewall.

A host marked as offline may be powered on and working, but it has not attempted communication to or through the firewall recently.

IPv6 Address
IPv6 address assigned to the client by DHCP, or static mapping address.

DHCP Unique Identifier (DUID)**
The **DHCP Unique Identifier** (DUID) of this client.

The DUID is meant to be unique *per device*.

IPv6 hosts are identified by a combination of the IAID and DUID.

Note: Due to variances in DHCPv6 clients between operating systems and manufacturers some DUID values are sized differently than others.

Hostname
The hostname of the client (if known).

Start/End
The beginning and end times of the DHCP lease.

Note: For static mappings the page prints `n/a` as static mapping leases do not have a start or end time, and they do not expire.

Actions
Icons to take action on this lease. See *Actions* for details.

Note: This list only contains information on leases handed out by DHCPv6. If a client assigns itself an address some other way (e.g. SLAAC) it cannot be listed on this page.

Search

The search box filters the contents of the **Leases** table based on keyword matching.

Enter a search string or UNIX regular expression into the box and click **Search** to filter the list to only matching records.

By default, the search looks at text from all fields in the lease record, but this can be limited to specific fields using the drop-down list.

Actions

Add static mapping

To create a static mapping from a dynamic lease, click  the to the right of the lease. This pre-fills the DUID of the client into the **Edit static mapping** screen.

See also:

DHCPv6 Static Mappings

Edit static mapping

Entries for existing static mappings have the  icon which takes the user to the page to edit that specific entry.

Wake on LAN Integration

Click  to create a WOL entry for the MAC address.

See also:

Wake on LAN

Delete a lease

While viewing the leases, an expired or inactive lease may be manually deleted by clicking  at the end of its line. This option is not available for active or static leases, only for offline or expired leases.

Delegated Prefixes

When **Prefix Delegation** is enabled, the bottom of the page lists delegated prefixes and their routing.

Entries for each delegated prefix include the following information:

IPv6 Prefix
The IPv6 prefix and length which the firewall has delegated to a DHCPv6 client.

Routed To
The IPv6 address to which the firewall is routing the IPv6 prefix.

The target address will match one of the leases in the list at the top of the screen.

IAID/DUID
The IAID+DUID combination will match one of the leases in the list at the top of the screen and identifies the host to which the firewall is routing this prefix.

Start/End
The start and end time of this delegation. If the client renews its delegation request, the times will update accordingly.

State
The status of the delegation, which can be one of:

Active
A current delegation to an active client.

Expired
A delegation which has expired because a client did not renew it before its expiration time.

Pool Usage Summary

The **Lease Utilization** section summarizes pool usage, giving a count of leases used in each pool configured in the DHCPv6 server.

View inactive leases

By default, the page lists active and static leases. Clicking **Show all configured leases** makes the page display all leases, including inactive and expired leases.

To reduce the view back to normal, click **Show active and static leases only**.

Clear DHCP Leases



The **Clear all DHCPv6 leases** button stops the DHCPv6 daemon, removes the entire lease database, and then starts the daemon again.

This does not remove static DHCPv6 mappings, only dynamic leases.

High Availability Status – Kea DHCP Only

Failover status for Kea DHCP is in a section at the bottom of the DHCP and DHCPv6 Leases pages as in figure *Kea DHCP Failover Status - Primary node, both online*. The failover status works identically for both DHCP and DHCPv6, so refer to *High Availability Status – Kea DHCP Only* for details on how the failover status operates.

28.1.11 DNS Resolver Status

The *DNS Resolver* status page at **Status > DNS Resolver** displays the current contents of the DNS resolver infrastructure cache.

The page contains a variety of statistics for DNS servers contacted by the resolver daemon (Unbound), though the type of content varies based on the current *DNS resolver mode*.

Resolver Mode

In resolver mode the status page contains a list of all authoritative DNS servers which Unbound has recently contacted along with the zone for which they were queried.

The page will typically contain entries for root servers (`.` zone), Top Level Domain (TLD) servers (e.g. `.com`, `.org`) and servers for specific domains. When there are multiple servers for the same zone, the statistics tracked in the cache help Unbound decide which server it will use for subsequent queries.

Forwarding Mode

In forwarding mode the status page contains entries for each of the available forwarding DNS servers. These *DNS Servers* are from the list configured on **System > General Setup** as well as from dynamic WAN sources, when present.

The statistics tracked in the cache help Unbound decide which server it will use for subsequent queries. For example, in a multi-WAN environment a server routed through a failed WAN would be marked as unresponsive and skipped. A server connected to a slow or high latency WAN would be skipped unless better sources were unavailable.

28.1.12 Gateway Status

The Gateway Status page at **Status > Gateways** displays the current status of individual gateways as well as gateway groups.

Gateways Tab

The **Gateways** tab displays the status for each gateway on the firewall, including manually defined gateways as well as dynamic gateways.

See also:

- *Gateways*

The status output includes the following information for each entry:

Name
The name of the gateway.

If this gateway is currently the default gateway for either IPv4 or IPv6, the page will print `(default)` after the name.

Gateway
The IP address of the gateway.

Monitor
The IP address being used by the gateway monitoring system to determine the status of the gateway.

If the gateway has a custom monitor IP address set, this field will be different from the gateway IP address.

If monitoring is disabled for this gateway this column contains the string `(unmonitored)`.

RTT
The round-trip time of the most recent gateway monitoring probes.

RTTsd
The standard deviation of the round-trip time of recent gateway monitoring probes. This indicates how much variance there is between the fastest and slowest recent responses from the monitoring address.

A high value in this column indicates that the latency on the path to the monitor IP address varies significantly, with large differences between the high and low values. This could be due to load or instability on the link, for example.

A low value indicates that the latency on the circuit is consistent, which can mean it is in good condition, has a light load, or is otherwise operating optimally.

Loss
The amount of packet loss the firewall has experienced recently while probing the monitor IP address. This may indicate that the circuit has a problem somewhere along the path to the monitor IP address.

Some hosts and routers drop or throttle ICMP messages during times of high load. Thus, experiencing packet loss from monitoring probes does not always indicate a problem with the circuit.

Tip: If the circuit appears to be working properly despite showing loss, it's possible that the monitoring probes have been dropped by a router somewhere in between the firewall and the monitor IP address host. In this case the best course of action is to choose a different monitor IP address.

Status
The status field shows the current state of the gateway. The status may be one of: Online, Offline, or Warning.

When in an offline or warning state the field also contains a text description of the problem. For example, it may indicate that the gateway is offline due to packet loss.

Description
The text description of the gateway, either from the manually configured settings or a default string for dynamic gateways.

Action
Selectively kill firewall states using this gateway in various ways.

-  : Kills all firewall states created by policy routing rules using this specific gateway by name.

  This does not include states which used this gateway as a part of a gateway group.

-  : Kills all firewall states created by policy routing rules using this specific gateway by IP address.

  The states must have matched a rule using policy routing with this gateway alone, in a gateway group, or via `reply-to` on a WAN-type interface.

-  : This icon is present only on default gateway entries. It will kill states using the default gateway (`0.0.0.0` for IPv4, `::` for IPv6). This affects states which did not match policy routing rules (e.g. default routing behavior, automatic gateway switching, etc.).

  Warning: This can be very disruptive as it may terminate sessions to firewall services and/or on multiple WANs and internal interfaces.

Gateway Groups Tab

The **Gateway Groups** tab shows the status of gateway group members and the groups as a whole.

See also:

- *Gateway Groups*

The status output includes the following information for each entry:

Group Name
The name of the gateway group.

Gateways
A table containing a row for each gateway group member arranged in tiers on different columns.

If all the gateways on a tier are down, the firewall will use the gateways on the next available tier. For example, if all the gateways in tier 1 are offline, the firewall will look for gateways in tier 2, then tier 3, and so on.

Description
The text description of the gateway group.

Action
Selectively kill firewall states using this gateway group.

Clicking the  icon kills all firewall states created by policy routing rules using this specific gateway group by name.

This does not affect other uses of the gateway group, such as services bound to a gateway group as their interface.

28.1.13 CARP Status

The CARP status page is a part of the pfSense® software GUI at **Status > CARP (failover)**. This page shows the current status of all configured CARP *Virtual IP addresses*. The page also provides troubleshooting and maintenance controls.

CARP Maintenance Controls

The top section of the page contains buttons to manage the CARP behavior of this node.

> **Warning:** After changing the enable/disable status or maintenance mode, it may take a few moments for a node to completely take over the **MASTER** status on all VIPs.
>
> After clicking one of the buttons the page may refresh before this process is complete. To ensure the status is accurate, wait a few moments and manually reload the page by clicking the page title in the breadcrumb bar.

Enable/Disable CARP

The first button toggles the enable/disable status of CARP temporarily, and will have one of two labels depending on the current status:

Temporarily Disable CARP
When CARP is active this button will temporarily disable CARP and remove the CARP VIP configuration from the operating system.

If this is the primary node, the secondary node will take over the **MASTER** role when the process completes.

This setting is not retained across reboots. If CARP is temporarily disabled and the firewall reboots, CARP will be active after the reboot.

Enable CARP
When CARP is disabled this button will enable CARP and reconfigure the CARP VIPs on the interfaces.

If this is the primary node it will take over the **MASTER** role when the process completes.

Maintenance Mode

The next button toggles CARP maintenance mode. In maintenance mode the VIP configuration remains on the interfaces and a node participating in CARP demotes itself naturally by increasing the advertising frequency skew of its VIPs to the maximum value, `254`. This allows other CARP nodes to take over the **MASTER** role naturally.

For example, the secondary node typically has a skew of `100`. If the primary node enters maintenance mode, the secondary node now has a lower skew (`100` is less than `254`) and the secondary node will assume the **MASTER** role as it will be advertising faster than the demoted primary node.

Maintenance mode persists across reboots so it can ensure that a node does not take back over prematurely before it is ready. This makes it useful for performing upgrades or other maintenance on the primary node.

The button has one of two labels depending on the current status:

Enter Persistent CARP Maintenance Mode
Sets the skew of all VIPs to `254` and sets the maintenance mode flag in the firewall configuration. If this flag is present in the configuration at boot time, the node will remain in maintenance mode.

Leave Persistent CARP Maintenance Mode
Sets the skew of all VIPs to the value specified in the VIP configuration and clears the maintenance mode flag in the firewall configuration.

Warning: If all nodes in a cluster are in maintenance mode, the result is unpredictable as they will all be using the same skew value. Only put one node in a cluster into maintenance mode at a time.

Reset Demotion Status

The system keeps track of a demotion value which can change based on the status of interfaces with CARP VIPs. For example, if an interface with a CARP VIP is down, the system increases the demotion value by `240`, and it adds that value internally to the VIP skews. This allows a node to automatically demote itself when it detects a problem. When an interface recovers it decreases the demotion value by the same amount.

When the demotion status is non-zero the status page displays a warning box at the top explaining that the demotion status may be incorrect with a button to reset the value. This **Reset CARP Demotion Status** button resets the demotion value back to the default of `0`.

Warning: Before resetting this value check all interfaces to ensure there is not an ongoing problem that needs resolved first. Resetting the demotion status while there is a problem could result in the demotion status becoming incorrect again when that problem is corrected.

Fixing the underlying problem will naturally correct the demotion value.

In rare cases a node may have a problem properly setting or clearing its own demotion status after processing interface events, and that is the only time the button should be used to return to a working status.

CARP Status

The **CARP Status** table includes entries for each CARP VIP configured on the firewall and also shows IP Alias VIPs which use a CARP VIP as a parent.

Each entry contains the following information:

Interface and VHID
The interface and VHID for a given CARP VIP entry.

For example, a CARP VIP on WAN with a VHID of `11` will be listed as `WAN@11`.

Virtual IP Addresses
The IP addresses associated with the CARP VIP. This includes the CARP VIP itself as well as any IP alias type VIPs which utilize this CARP VIP as a parent.

Status
The **Status** column shows one of the following status strings:

MASTER
Indicates this node is accepting all traffic for this VIP

BACKUP
Indicates this node is monitoring CARP advertisements and not accepting traffic for the VIP.

INIT or blank
Generally indicates a problem with the VIP. Either the VIP is not configured at the OS level, the interface upon which it is configured is down, or the interface has a problem.

When operating normally the primary node should show each VIP in **MASTER** status. On the secondary node each VIP should show **BACKUP** for the status.

If both nodes show **MASTER** there is usually a problem at layer 2 (the switch) preventing the nodes from seeing advertisements from the other node.

See also:

See *Troubleshooting High Availability* for help troubleshooting CARP.

State Synchronization Status / pfsync Nodes

The bottom section of the page contains a list of state creator host IDs.

On current versions of pfSense software the default ID for a host is the last 8 characters of its NDI, but there is an option to set a specific custom ID (See *Filter Host ID*). On previous versions the default behavior was to generate a randomized value on every filter reload.

When a cluster is configured for state synchronization each node should see states created by IDs from other nodes in this list, indicating that they are properly synchronizing state table data.

There can be some slight differences in the list depending on timing (e.g. when changing the host ID to a custom value) but the list should be nearly identical on all nodes.

Widget

There is a **CARP Status** widget available for the *Dashboard* which shows similar information in a condensed format without the maintenance controls.

28.1.14 Captive Portal Status

The **Captive Portal Status** page is available through the pfSense® software GUI at **Status > Captive Portal**.

The status page defaults to the **Active Users** tab which displays information about active user sessions.

See also:

The other tabs on this page are related to managing voucher rolls and not user sessions. See *Vouchers* for details.

Active Users Tab

The **Active Users** tab lists online users for captive portal zone configurations with and without authentication.

If the captive portal configuration contains more than one zone, the first option presented by the GUI is a drop-down list of zones. To view the sessions for a zone, select it from the **Display Zone** drop-down.

The list of online users contains the following information for each entry:

IP address
The IP address of the captive portal user.

MAC Address
The MAC Address of the captive portal user.

This field may be blank if the zone does not utilize MAC filtering.

Username
Can be one of several things, depending on the zone configuration:

- The name of an authenticated captive portal user if they logged in with a username and password.
- The voucher code if the user authenticated using vouchers.
- The MAC address of a host if the zone uses **RADIUS** with **MAC Authentication**.
- The string `unauthenticated` if the zone does not require authentication.

Session Start Time
The time when the user authenticated their current session.

Actions

A  icon that, if clicked and confirmed, will terminate a captive portal session. The user must authenticate again to get back online.

Active Users Status Examples

These figures are examples of several different user styles that may appear in the list.

For **Local User Manager** or **RADIUS** authentication with username/password, status entries look like Figure *Online Captive Portal Users: User Authentication*:



Users Logged In (1)

IP address	MAC address	Username	Session start	Actions
10.7.0.10	00:0c:29:4c:b3:9b	jimp	06/13/2016 12:29:51	

Fig. 8: Online Captive Portal Users: User Authentication

If a zone allows voucher-based authentication, users that signed on using a voucher look like Figure *Online Captive Portal Users: Vouchers*:



Users Logged In (1)

IP address	MAC address	Username	Session start	Actions
10.7.0.10	00:0c:29:4c:b3:9b	dCKPsPLLiL83	06/13/2016 12:31:23	

Fig. 9: Online Captive Portal Users: Vouchers

If a zone is set to **No Authentication**, status entries look like Figure *Online Captive Portal Users: No Authentication*:

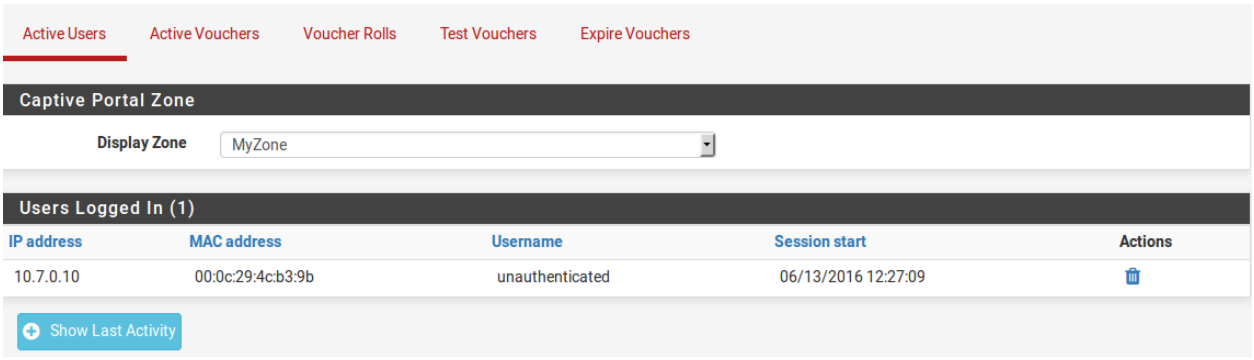


Fig. 10: Online Captive Portal Users: No Authentication

28.1.15 Viewing Active Network Sockets

The **Diagnostics > Sockets** page prints a list of active TCP/IP sockets for both IPv4 and IPv6 *used by the firewall itself*.

Note: The output of this command only shows sockets used by the firewall OS for daemons or other programs on the firewall. It does **not** show connections for traffic passing through the firewall.

This list is useful for determining which IP addresses and ports are in use by various firewall processes and/or packages. The firewall interprets the contents of the page from the output of the FreeBSD command `sockstat`.

By default, the page only displays *listening* sockets. Click **Show all socket connections** to also display sockets in use by the firewall for connections to external hosts.

Each row in the output contains the following information:

User
: The operating system user who owns the socket (e.g. `root`)

Command
: The command which holds the socket. This might be a daemon or a program making an outbound connection.

PID
: The process ID of the command holding the socket.

FD
: The file descriptor number of the socket.

Proto
: The transport protocol and address family combined (e.g. TCP4, UDP6, UDP46).

Local
: The local IP address and port number associated with this socket.

Foreign
: The remote IP address and port number associated with this socket.

28.1.16 ARP Table

IPv4 Hosts use ARP (Address Resolution Protocol) to locate IPv4 neighbors by MAC address on a directly connected network.

The ARP table in pfSense® software displays a list of IPv4 hosts on the network which have attempted to talk to or through the firewall within the past few minutes. If a host is up but has not talked to or through the firewall it will not appear in the ARP table.

See also:

For IPv6 hosts, see *NDP Table*.

To view the contents of the ARP table in pfSense software, navigate to **Diagnostics > ARP Table**.

The page contains the following items for each ARP table entry:

Interface
The interface where the firewall observed the host. If the interface is assigned, this field contains the given name of the interface in pfSense software. Otherwise, the page displays the operating system interface name.

IP Address
The IPv4 address of the host.

MAC Address
The MAC address of the host.

A MAC address listed as **(Incomplete)** indicates that the firewall has attempted to discover the host via ARP, but it has not yet received a valid response.

Tip: Installing the NMAP package activates a feature which allows the page to also display the manufacturer associated with the MAC address, if it is known. Note that this is not effective in some cases, such as for virtual machines which use randomly generated MAC addresses or for wireless clients which utilize privacy features that alter their MAC addresses.

Hostname
The fully qualified domain name, or at least the hostname portion, of the host. This can be discovered via DHCP lease database content or by a reverse lookup of the IP address via DNS.

Status
The status of the entry, typically one of two types:

Permanent
A static entry either located on the firewall itself (e.g. interface address, VIP) or a static ARP entry.

Expires in <time>
A dynamic ARP entry which will expire in `<time>` unless the host communicates to or through the firewall again. The default maximum age is `1200` seconds (20 minutes).

Link Type
The type of network link through which this host can be reached (e.g. **Ethernet**).

Actions
Contains the  icon that, if clicked and confirmed, will remove this ARP table entry. This can nudge the firewall to discover a new MAC address for a host if it changes.

The  **Clear ARP Table** button purges the entire contents of the ARP table. Clearing the ARP table is not typically necessary but can help the firewall in situations where multiple hosts have changed MAC addresses and the firewall is still attempting to communicate with the old addresses.

28.1.17 NDP Table

IPv6 Hosts use NDP (Neighbor Discovery Protocol) to locate IPv6 neighbors by MAC address on a directly connected network.

The NDP table in pfSense® software displays a list of IPv6 hosts on the network which have attempted to talk to or through the firewall within the past few minutes. If a host is up but has not talked to or through the firewall it will not appear in the NDP table.

See also:

For IPv4 hosts, see *ARP Table*.

To view the contents of the NDP table in pfSense software, navigate to **Diagnostics > NDP Table**.

The page contains the following items for each NDP table entry:

IPv6 Address
The IPv6 address of the host.

MAC Address
The MAC address of the host.

A MAC address listed as **(Incomplete)** indicates that the firewall has attempted to discover the host via NDP, but it has not yet received a valid response.

Tip: Installing the NMAP package activates a feature which allows the page to also display the manufacturer associated with the MAC address, if it is known. Note that this is not effective in some cases, such as for virtual machines which use randomly generated MAC addresses or for wireless clients which utilize privacy features that alter their MAC addresses.

Hostname
The fully qualified domain name, or at least the hostname portion, of the host. This can be discovered using reverse lookup of the IPv6 address via DNS.

Interface
The interface where the firewall observed the host. If the interface is assigned, this field contains the given name of the interface in pfSense software. Otherwise, the page displays the operating system interface name.

Expiration
The expiration status of the entry, typically one of two types:

Permanent
A static entry either located on the firewall itself (e.g. interface address, VIP) or a static NDP entry.

<time>
A dynamic NDP entry which will expire in `<time>` unless the host communicates to or through the firewall again.

Actions

Contains the  icon that, if clicked and confirmed, will remove this NDP table entry. This can nudge the firewall to discover a new MAC address for a host if it changes.

The  **Clear NDP Table** button purges the entire contents of the NDP table. Clearing the NDP table is not typically necessary but can help the firewall in situations where multiple hosts have changed MAC addresses and the firewall is still attempting to communicate with the old addresses.

28.1.18 Hardware Temperature Monitoring

FreeBSD, and thus pfSense® software, can monitor hardware temperature on compatible devices.

Widgets and Graphs

The *System Information* dashboard widget contains a single **Temperature** entry on systems with one or more available temperature sensors.

The *Thermal Sensors* widget displays the temperature from every available thermal sensor on the firewall. It also has configurable behavior for warning thresholds and output contents.

When the firewall detects compatible temperature monitoring hardware, it also graphs the temperatures over time. See *Monitoring Graphs* for details.

Intel and AMD Temperature Monitoring

There is an option to select either an Intel or AMD temperature monitor module under **System > Advanced** on the **Miscellaneous** tab. These modules work with Intel Core series and later Intel chips, and similarly recent AMD chips, respectively. Compatibility is not universal, but it is common, especially on Intel chips, even Atom-based chips. The temperature can be observed using the Thermal Sensors dashboard widget or by `sysctl`:

```
# sysctl -a | grep "dev.cpu.*.temperature"
dev.cpu.0.temperature: 46.0C
dev.cpu.1.temperature: 47.0C
dev.cpu.2.temperature: 47.0C
dev.cpu.3.temperature: 47.0C
```

These are typically on-die sensors so they only represent the CPU core temperatures, not other zones in the system.

ACPI Thermal Monitoring

FreeBSD also handles some settings through ACPI. To see if the hardware is capable of temperature monitoring, try the following command from a shell prompt or **Diagnostics > Command**:

```
# sysctl hw.acpi.thermal
```

If the hardware is compatible, the command output will be similar to the following:

```
hw.acpi.thermal.min_runtime: 0
hw.acpi.thermal.polling_rate: 10
hw.acpi.thermal.user_override: 0
hw.acpi.thermal.tz0.temperature: 22.5C
hw.acpi.thermal.tz0.active: -1
hw.acpi.thermal.tz0.passive_cooling: 1
hw.acpi.thermal.tz0.thermal_flags: 0
hw.acpi.thermal.tz0._PSV: 85.0C
hw.acpi.thermal.tz0._HOT: -1
hw.acpi.thermal.tz0._CRT: 100.0C
hw.acpi.thermal.tz0._ACx: 85.0C -1 -1 -1 -1 -1 -1 -1 -1 -1
```

In this example, there is only one thermal zone, and its temperature is 22.5C (72.5F).

From the acpi_thermal(4) man page:

```
hw.acpi.thermal.min_runtime
     Number of seconds to continue active cooling once started.  A new
     active cooling level will not be selected until this interval
     expires.

hw.acpi.thermal.polling_rate
     Number of seconds between polling the current temperature.

hw.acpi.thermal.user_override
     If set to 1, allow user override of various setpoints (below).
     The original values for these settings are obtained from the BIOS
     and system overheating and possible damage could occur if
     changed.  Default is 0 (no override).

hw.acpi.thermal.tz%d.active
     Current active cooling system state. If this is non-negative,
     the appropriate _AC%d object is running.  Set this value to the
     desired active cooling level to force the corresponding fan
     object to the appropriate level.

hw.acpi.thermal.tz%d.passive_cooling
     If set to 1, passive cooling is enabled. It does cooling without
     fans using cpufreq(4) as the mechanism for controlling CPU speed.
     Default is enabled for tz0 where it is available.

hw.acpi.thermal.tz%d.thermal_flags
     Current thermal zone status.  These are bit-masked values.

hw.acpi.thermal.tz%d.temperature
     Current temperature for this zone.

hw.acpi.thermal.tz%d._PSV
     Temperature to start passive cooling by throttling down CPU, etc.
     This value can be overridden by the user.

hw.acpi.thermal.tz%d._HOT
     Temperature to start critical suspend to disk (S4).  This value
```

(continues on next page)

(continued from previous page)

```
     can be overridden by the user.

hw.acpi.thermal.tz%d._CRT
     Temperature to start critical shutdown (S5).  This value can be
     overridden by the user.

hw.acpi.thermal.tz%d._ACx
     Temperatures at which to switch to the corresponding active cooling
     level.  The lower the _ACx value, the higher the cooling power.

All temperatures are printed in Celsius.  Values can be set in Celsius
(by providing a trailing "C") or Kelvin (by leaving off any trailing letter).
When setting a value by sysctl(8), do not specify a trailing decimal (i.e.,
90C instead of 90.0C).
```

The defaults for these values are taken from the BIOS and some systems will not allow changes. The only way to know is to try. Before attempting to alter any of these values, set this OID:

```
# sysctl -w hw.acpi.thermal.user_override=1
```

These values may be set by adding the appropriate lines for the OIDs to **System > Advanced** on the **Tunables** tab. Try setting them at run time:

```
# sysctl -w hw.acpi.thermal.tz0._CRT=120C
```

28.1.19 IPsec Status

The IPsec status page at **Status > IPsec** displays the current state of all IPsec tunnels configured on the firewall.

This page is divided into four tabs.

See also:

- *IPsec Logs*
- *IPsec Troubleshooting*

Overview Tab

This tab lists all enabled IPsec tunnels. Each entry contains the tunnel description, links to its settings, outer and inner IP addresses, various properties of the tunnel, counters, and current status.

Connected tunnels are listed first, followed by disconnected tunnels. There are buttons on each row to connect or disconnect entries manually.

By default, only the IKE portion of a tunnel (phase 1) is listed to keep the display compact. Click  **Show child SA entries** to display the child SA (phase 2) entries.

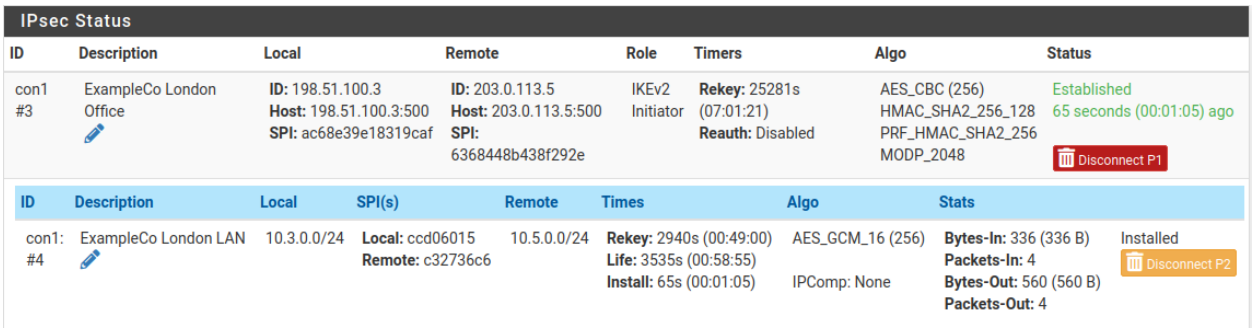


Fig. 11: Connected Tunnel with Child SA List expanded

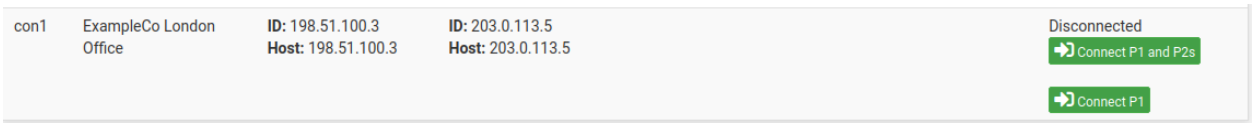


Fig. 12: Disconnected Tunnel

Leases Tab

Lists current usage statistics for mobile IPsec client leases from configured pools. Current and recently connected clients are also listed along with the IP address they were assigned by the firewall.

SAD Tab

Shows the contents of the IPsec Security Association Database (SAD) which contains data about current IKE SA entries and corresponds with active phase 1 entries.

The page contains one entry in the list for each direction between *public peer addresses* of an active IPsec tunnel. For example, one entry for `x.x.x.x` to `y.y.y.y` and a corresponding entry for `y.y.y.y` to `x.x.x.x`.

SPD Tab

Shows the contents of the IPsec Security Policy Database (SPD). These policies define the networks which are interesting to IPsec and corresponds with phase 2 entries.

The page contains one entry for each direction between *private networks* of all IPsec tunnels whether they are connected.

28.1.20 OpenVPN Server and Client Status

The OpenVPN status page at **Status > OpenVPN** shows the status of each OpenVPN server and client. The status includes service controls for each separate server and client instance on the status page.

SSL/TLS Client/Server Mode

For OpenVPN servers in SSL/TLS client/server mode (tunnel network larger than `/30`), the status provides a list of connected remote clients along with their usernames or certificate common names and connection data as seen in Figure *OpenVPN status for an SSL/TLS server with one connected client*.

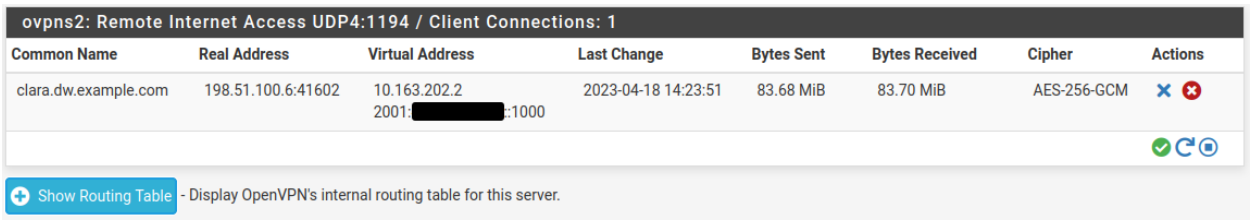


Fig. 13: OpenVPN status for an SSL/TLS server with one connected client

The status of an instance has a header bar which includes the OS interface name (e.g. `ovpns1`), the custom text description of the VPN, its protocol and port, plus a total count of connected clients.

The status output includes the following columns in separate blocks for each server instance in this mode:

Common Name
: The certificate common name and/or username of the client. For VPNs utilizing user authentication, both are printed in this column. The values may differ depending on whether the server has SSL/TLS enabled and whether the **Username as Common Name** option is enabled.

Real Address
: The external/public IP address of the client, as it would appear on the WAN.

Virtual Address
: The tunnel network IPv4 and/or IPv6 addresses assigned to the client for use inside the VPN.

Connected Since
: A timestamp indicating when this client connected to the server or the last status change of the connection.

Bytes Sent
: The amount of data the OpenVPN server has sent to this client.

Bytes Received
: The amount of data the OpenVPN server has received from this client.

Ciphers
: The encryption algorithm in use for this client, which may vary due to cipher negotiation.

Actions
: This column includes icons which control the client.

  The  icon will appear at the end of each client row if that client authenticated via RADIUS and has firewall rules received from RADIUS. Clicking the icon will open a modal dialog displaying the contents of that user's personal firewall ruleset.

  The  icon at the end of each client row clears the client session, which disconnects the client while allowing them to reconnect.

  The  icon at the end of each client row sends a command which halts the remote client. If the client honors the request, its process terminates, and it will not automatically reconnect without

manual intervention. This can be useful for stopping an unattended client from conflicting with a different active session for a user.

The  **Show Routing Table** button under each server's list of clients displays a table of networks and IP addresses connected through each client connected to that server.

Peer-to-Peer Mode

For OpenVPN instances in peer-to-peer mode (shared key or SSL/TLS with a `/30` tunnel network), the output is slightly different. OpenVPN does not report the same amount of information for instances running in peer-to-peer mode, so it cannot offer the same functionality as SSL/TLS client/server mode.

As each instance in this mode is limited to one client per server, the entries are shown in a single table each for clients and servers, with one instance listed per row.

The **Name** column prints the OS interface name for the VPN (e.g. `ovpns1`) and its configured text description, along with the protocol and port number.

For a server instance the **Status** column indicates whether the instance is running and waiting on connections or if the remote client has connected.

For client instances the **Status** column indicates whether a connection is pending or active.

The status column may display more detailed information if it's available during certain stages of configuration and connection.

Peer to Peer Server Instance Statistics

Name	Status	Last Change	Virtual Address	Remote Host	Bytes Sent	Bytes Received	Service
ovpns5 S2S [redacted] UDP4:1198	Connected (Success)	Mon Apr 17 8:41:48 2023	10.18.105.1	198.51.100.19	3.37 MiB	3.37 MiB	
ovpns6 S2S [redacted] UDP4:1199	Waiting for peer connection	Tue Apr 25 15:42:43 2023			0 B	0 B	

Client Instance Statistics

Name	Status	Last Change	Local Address	Virtual Address	Remote Host	Bytes Sent	Bytes Received	Service
ovpnc6 Client UDP4	Reconnecting (Ping Restart)	Tue Apr 25 15:43:37 2023	(pending)		(pending)	0 B	0 B	

Fig. 14: OpenVPN status showing peer-to-peer instances including a server that is up, a server waiting for a connection, and a client attempting to reconnect

28.1.21 Route Table Contents

The current contents of the firewall route table are displayed by the GUI page at **Diagnostics > Routes**. The CLI can also be display the route table using the command `netstat -rWn`.

Route Table GUI

The GUI route table contents looks like Figure *Route Table Display*.

IPv4 Routes

Destination	Gateway	Flags	Use	Mtu	Netif	Expire
default	198.51.100.1	UGS	1797	1500	igb1	
10.2.0.0/24	link#2	U	0	1500	igb0	
10.2.0.1	link#2	UHS	0	16384	lo0	
127.0.0.1	link#11	UH	204	16384	lo0	
198.51.100.0/24	link#3	U	907	1500	igb1	
198.51.100.1	00:08:a2:09:95:b6	UHS	2519	1500	igb1	
198.51.100.2	link#3	UHS	0	16384	lo0	

Fig. 15: Route Table Display

The route table contents are described in detail *later in this document*.

Routing Table Display Options

The list of routes displayed by the GUI allows pagination and filtering to aid with viewing large routing tables, such as those found with a full BGP feed. The top section of the page contains the following options which control the behavior of the page:

Resolve Names
: This option controls whether the firewall attempts to resolve items using DNS. The default is unchecked, which disables DNS resolution. When checked, the firewall attempts a DNS lookup to show hostnames rather than IP addresses for route table entries.

> **Warning:** Enabling this feature causes a delay and performance penalty as the page attempts to resolve all the entries. As the size of the table increases, the delay will also increase and performance will degrade further.

Rows to display
: This option controls the number of rows output from each route table. By default, the page displays `100` rows. Choose a new value to show more or less rows.

Filter
: This text entry box defines a string or pattern which the page uses to search the route table for matching entries. The field performs regular expression matching for advanced filtering.

Click  **Update** to re-display the routing table with the current settings.

Route Table CLI

Viewing the route table in the CLI is similar to the GUI. The *same information* is present, and the labels are similar.

The `netstat -rWn` command can be run from a console or SSH shell:

```
$ netstat -rWn
Routing tables

Internet:
Destination        Gateway            Flags       Use    Mtu      Netif Expire
default            198.51.100.1       UGS         294   1500     vtnet0
127.0.0.1          link#4             UH        20976  16384        lo0
192.168.1.0/24     link#2             U             1   1500     vtnet1
192.168.1.1        link#2             UHS           0  16384        lo0
198.51.100.0/24    link#1             U           116   1500     vtnet0
198.51.100.1       ca:1d:62:6c:c6:9c  UHS         191   1500     vtnet0
198.51.100.103     link#1             UHS           0  16384        lo0
```

Omit the `-n` flag and the command will attempt to use DNS to resolve IP addresses to hostnames where possible.

IPv4 and IPv6 Route Table Content

The route table information output by either the GUI or the CLI contains the following fields:

Destination
: The destination network or host for this route.

  The default route for each address family is listed as `default`. Otherwise, hosts are listed as an IP address and networks are listed with an IP address and CIDR mask or prefix.

Gateway
: The next hop through which the firewall will route traffic going to the **Destination**.

  If this column shows a link, such as `link#1`, then that network is directly reachable by that interface and no special routing is necessary. If a host is visible with a MAC address, then it is a locally reachable host with an entry in the ARP table, and packets are sent there directly.

Flags
: Properties of this route. See *Route Table Flags* for the meanings of each flag.

Uses
: The total number of packets the firewall has sent via this route.

  This is helpful for determining if the firewall is actively using a route as the value will continually increment as packets utilize the route.

MTU
: The MTU for packets using this route.

Interface
: The interface through which the firewall will route traffic for **Destination**.

Expire
: An expiration time for temporary routes, such as those added from ICMP redirects.

Route Table Flags

There are quite a few flags, all of which are covered in the FreeBSD man page for *netstat(1)*. The portion of the content from that document covering flags is reproduced in *Route Table Flags*.

Table 1: Route Table Flags

Letter	Flag	Meaning
1	RTF_PROTO1	Protocol specific routing flag #1
2	RTF_PROTO2	Protocol specific routing flag #2
3	RTF_PROTO3	Protocol specific routing flag #3
B	RTF_BLACKHOLE	Discard packets during updates
b	RTF_BROADCAST	Represents a broadcast address
D	RTF_DYNAMIC	Created dynamically by redirect
G	RTF_GATEWAY	Destination requires forwarding by intermediary
H	RTF_HOST	Host entry (net otherwise)
L	RTF_LLINFO	Valid protocol to link address translation
M	RTF_MODIFIED	Modified dynamically (by redirect)
R	RTF_REJECT	Host or net unreachable
S	RTF_STATIC	Manually added
U	RTF_UP	Route usable
X	RTF_XRESOLVE	External daemon translates proto to link address

For example, a route flagged as `UGS` is a usable route, packets are sent via the gateway listed, and it is a static route.

28.1.22 Wireless Status

The Wireless Status page at **Status > Wireless Status** displays a list of access points within range of the wireless radio along with various statistics about the access points detected.

If a wireless card is acting as an access point the page prints a list of associated clients along with related information.

Note: This page is only available when a wireless interface is detected in the firewall and that interface is enabled.

Access Point/Ad-hoc Peer Capabilities

Name	Description
WME	Wireless Multimedia Extensions (QoS)
WPA	Wi-Fi Protected Access
WPS	Wi-Fi Protected Setup
RSN	802.11i - Robust Secure Network
HTCAP	802.11n - High Throughput (HT)
ATH	Atheros protocol extensions
VEN	Unknown vendor-specific extensions

Access Point/Ad-hoc Peer Flags

Flag	Description
E	ESS Mode (Access Point)
I	IBSS Mode (Infrastructure/Client)
c	CF Pollable
C	CF Poll Request
P	Privacy
S	Short Preamble
B	PBCC
A	Channel Agility
s	Short Slot Time
R	RSN
D	DSSSOFDM

Peer Status Flags

Flag	Description
A	Authorized for Data
Q	QoS (WME)
E	Extended Rate (ERP), 802.11g
P	Power Save Mode
H	High Throughput (HT)
H+	HT Compat mode (Setup with vendor OUIs)
W	Wi-Fi Protected Setup (WPS)
N	Transitional Security Network (TSN) association
T	Aggregated MAC Protocol Data Unit (AMPDU) Transmit
R	AMPDU Receive
M	MIMO Power Save
M+	MIMO Power Save with RTS
I	Reduced Interframe Space (RIFS)
S	Short GI in HT40
S+	Short GI in HT40 and HT20
s	Short GI in HT20
t	Aggregated MAC Service Data Unit (AMSDU) Transmit
r	AMSDU Receive

28.1.23 ALTQ Traffic Shaper Queue Monitoring

The page at **Status > Queues** displays ALTQ traffic shaper queue usage. Monitor this page to ensure that traffic shaping is working as intended.

Figure *Basic WAN Queues* shows each queue listed by name, its current usage, and other related statistics.

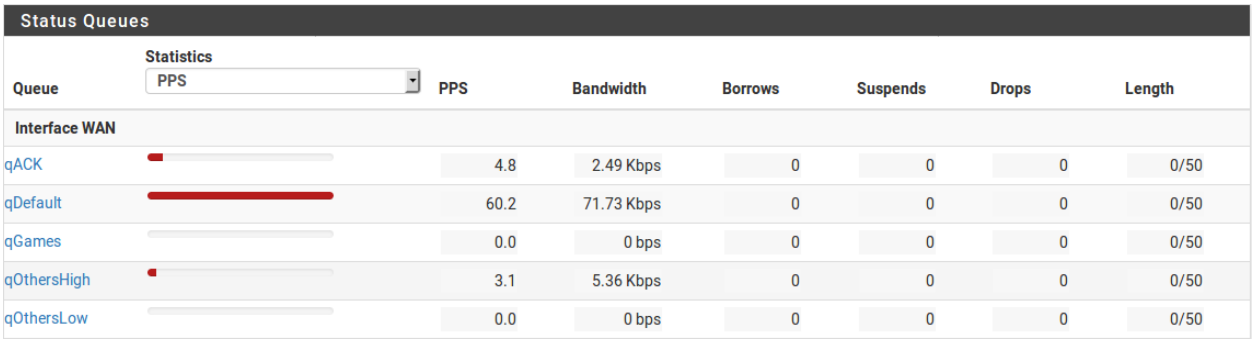


Fig. 16: Basic WAN Queues

Queue Status Columns

Queue
The name of the traffic shaper queue.

Statistics
A graphical bar which shows how "full" this queue is.

PPS
The rate of queued data in packets per second (PPS)

Bandwidth
The rate of queued data in bits per second (e.g. Mbps, Kbps, bps).

Borrows
Borrows happen when a neighboring queue is not full and capacity is borrowed from there.

Suspends
The suspends counter indicates when a delay action happens. The suspends counter is only used with the CBQ scheduler and should be zero when other schedulers are in use.

Drops
Drops happen when traffic in a queue is dropped in favor of higher priority traffic. Drops are normal and this does not mean that a full connection is dropped, only a packet. Usually, one side of the connection will see that a packet was missed and then resend, often slowing down in the process to avoid future drops.

Length
The number of packets in the queue waiting to be transmitted, over the total size of the queue.

Queue Status Options

There are two options on the page which control the output:

Refresh Rate
Controls how frequently the page updates itself automatically. This defaults to `1` second, but can set between `0.5` seconds and `5` seconds.

Statistics
Controls which data point is used to when generating bar graphs for queue content. This can be *PPS* (Default) or *Bandwidth*.

28.1.24 NTP Daemon Status

The NTP status page at **Status > NTP** shows the status of the NTP daemon. This status includes information on upstream NTP servers as well as other items such as GPS information if such a device is connected and enabled.

An example of the status is shown in Figure *NTP Daemon Status With GPS Output*.

Network Time Protocol Status

Status	Server	Ref ID	Stratum	Type	When	Poll (s)	Reach	Delay (ms)	Offset (ms)	Jitter (ms)
False Ticker	127.127.20.0	.GPS.	0	l	9	16	377	0.000	-127.53	1.879
Pool Placeholder	0.pfsense.pool.ntp.org	.POOL.	16	p	-	64	0	0.000	+0.000	0.004
Pool Placeholder	1.pfsense.pool.ntp.org	.POOL.	16	p	-	64	0	0.000	+0.000	0.004
Pool Placeholder	2.pfsense.pool.ntp.org	.POOL.	16	p	-	64	0	0.000	+0.000	0.004
Active Peer	66.220.9.122	.CDMA.	1	u	38	256	377	50.304	+0.264	1.763
Candidate	162.159.200.123	10.27.8.152	3	u	136	256	377	16.325	-0.033	1.472
Candidate	162.159.200.1	10.27.8.152	3	u	68	256	377	16.051	-0.082	0.213

GPS Information

Clock Latitude	Clock Longitude
38.[redacted] (38°[redacted]N)	-86.[redacted] (86°[redacted]W)
Google Maps Link	

Fig. 17: NTP Daemon Status With GPS Output

The status screen contains one line for every peer, and each entry contains the following items:

Status
The current status of the NTP server entry, which may be one of:

Pool Placeholder
Placeholder entry for pools defined in the configuration. When using a pool, these are the entries in the NTP configuration and the NTP daemon dynamically discovers the servers based on the DNS response for the pool address.

Unreach/Pending
The server is either unreachable or it has not yet responded with sufficient information to determine a better status.

Active Peer
The NTP server with which the daemon is currently synchronizing its clock.

Candidate
A peer which has a high enough quality status to become an active peer.

PPS Peer
A peer utilizing PPS signaling.

Selected
A peer which is included by the combine algorithm.

Excess Peer
An excess peer which is not included because there are too many other viable peers.

False Ticker
A peer which has been rejected by the intersection algorithm.

Outlier
A peer which has been discarded by the cluster algorithm.

See also:

For more information on the algorithms involved in determining status, see the NTP documentation page on Clock Select Algorithm.

Server
The hostname or IP address of the NTP server. For pool placeholders, this is the FQDN of the pool.

Some special sources such as a local NMEA reference clock like a GPS use reserved addresses. A GPS, for example, appear as an address in `127.127.20.x`.

Ref ID
A reference ID indicating how this NTP server is determining its own time.

This is typically a source IP Address for higher stratum servers (2 or higher) which is used by clients to prevent a loop where servers might reference each other which could lead to inaccuracies.

Stratum 1 servers are considered reference clocks. For these servers this reference ID is a four character string indicating the type of reference used to determine the NTP server's own time. There are a handful of pre-defined codes in RFC 5905, but servers can use custom codes as well.

Some commonly observed reference ID codes are:

POOL
This is a pool placeholder entry and not a source.

GPS
This NTP server determines its time based on a GPS device.

PPS
This NTP server determines its time based on a PPS device.

CDMA
This NTP server determines its time based on a CDMA network reference.

Stratum 0 servers are invalid as a time source but offer a four character "Kiss code" indicating a reason why they are not a valid time source at that moment.

See also:

For a full list of pre-defined Kiss and Reference ID codes, see the IANA NTP parameter lists.

Stratum
The stratum of this NTP server, in the following ranges:

Stratum 0
Invalid or unavailable time source.

Stratum 1
Reference clocks which use a non-NTP time source to keep time, such as a GPS, atomic clock, etc.

These servers are highly accurate but should not be used directly by clients. They should be used primarily by servers which will offer NTP service to many of their own local clients.

Stratum 2-14
NTP servers which obtain their own time from other NTP servers on a lower stratum.

These are typically the best sources for clients to use. Though they are potentially not as accurate as stratum 1 servers, in nearly all cases they are accurate enough for the majority of purposes.

Stratum 15
The server is not suitable for synchronizing clients (See the NTP documentation page on Clock Select Algorithm).

Stratum 16
Unsynchronized source.

Type
A code indicating the type of NTP server:

u
Unicast or manycast client

b
Broadcast or multicast client

p
Pool source

l
Local (reference clock)

s
Symmetric (peer)

A
Manycast server

B
Broadcast server

M
Multicast server

When
The amount of time since the NTP daemon last received a packet from this server. The value is - if NTP has not received any packets.

Poll(s)
The interval between polls to this server.

Reach
Reach shift register, displayed in octal.

Delay (ms)
Round trip delay to this NTP server.

Offset (ms)
Clock offset of this NTP server relative to this firewall.

Jitter (ms)
Offset RMS error estimate.

If a serial GPS is connected and configured, the page also prints the coordinates reported by the GPS device along with a link to the coordinates on Google Maps.

Note: The quality of GPS data can vary widely depending on the signal level, the GPS device, and how it is connected. Traditional serial ports are higher quality and better suited to GPS clock usage. USB serial GPS units may be acceptable, but due to how USB functions, the timing of signals cannot be guaranteed the way it can be with a traditional hard-wired serial port.

28.2 Graphs

These articles cover graphs for monitoring pfSense software itself as well as for traffic on interfaces and using additional packages for more detailed monitoring of user throughput/usage.

28.2.1 Monitoring Graphs

pfSense® software has many built-in graphs that monitor different aspects of the system, and they work out-of-the-box with no intervention.

The firewall collects and maintains data about how the system performs, and then stores this data in Round-Robin Database (RRD) files. Graphs created from this data are available under **Status > Monitoring**.

These graphs measure things such as CPU usage, memory usage, state table usage, throughput (in bytes as well as packets), link quality, traffic shaping queue usage, and more.

The graph on that page can be configured to show items from several categories, and a category and graph may be chosen for both the left axis and right axis for easy comparison.

Working with Graphs

The firewall displays a graph showing its CPU usage by default. To view other graphs or to add a second category on another axis, the graph settings must be changed as described in the next section, *Graph Settings*.

Inside the graph, the labels in the top left corner note the sources for the data in the left axis and right axis.

The graph contains a legend at the top right with each of the data sources plotted on the graph. Clicking a data source in the legend will hide it from view.

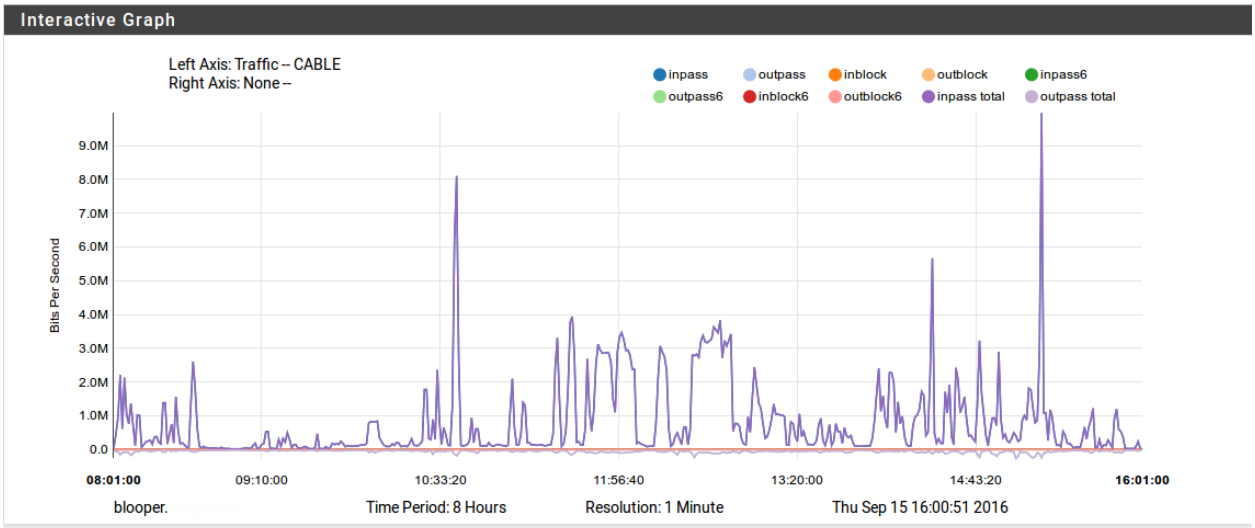
Tip: If a data source has a large spike, click its name in the legend to remove it from the graph. With the larger data source removed, more detail from the other remaining sources will be visible.

The firewall hostname, graph time period, and graph resolution are printed along the bottom of the graph, along with the time the graph was generated.

The firewall prints a table below the graph itself with a summarization of the data. This table contains minimums, averages, maximums, current values, in some cases 95th percentile values. In cases where units are given, hovering the mouse pointer over the unit will display a more detailed description of the unit.

Note: Totals are not displayed because the way data is stored in RRD files, accurate totals are not possible. To see total usage for traffic on network interfaces, install the **Status Traffic Totals** package.

Figure *WAN Traffic Graph* shows an example of an 8-hour graph of traffic on a firewall interface named *CABLE* with inverse enabled. The interface has a maximum utilization of 9.96Mbit/s during a one-minute period.



Data Summary

	Minimum	Average	Maximum	Last	95th Percentile
inpass	16.52 kb/s	817.76 kb/s	9.96 Mb/s	0.00 b/s	3.12 Mb/s
outpass	12.44 kb/s	56.21 kb/s	262.97 kb/s	0.00 b/s	122.66 kb/s
inblock	0.00 b/s	7.27 b/s	676.90 b/s	0.00 b/s	
outblock	0.00 b/s	0.04 b/s	11.54 b/s	0.00 b/s	
inpass6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s
outpass6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s
inblock6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	
outblock6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	
inpass total	16.52 kb/s	817.76 kb/s	9.96 Mb/s	0.00 b/s	3.12 Mb/s
outpass total	12.44 kb/s	56.21 kb/s	262.97 kb/s	0.00 b/s	122.66 kb/s

Fig. 18: WAN Traffic Graph

Graph Settings

To change the graph, click  in the breadcrumb bar to display the graph settings panel.

Tip: The graph settings panel is hidden by default, but this behavior can be changed. Navigate to **System > General Setup** and check **Monitoring Settings** to always display the settings panel by default.

The options on the settings panel are:

Left Axis / Right Axis
The options here control the data displayed on each axis. By default, only the **Left Axis** is populated with a value, but both may be utilized to compare areas. First pick a **Category** (or *None*), then pick a **Graph** inside that category. The list of available categories and graphs will vary depending on the firewall configuration.

Category
The general area of the desired graph: System, Traffic, Packets, Quality, Captive Portal, NTP, Queues, QueueDrops, DHCP, Cellular, Wireless, and VPN Users. These are covered in more detail later in this section.

Graph
The specific graph to display from the chosen category.

Options
This section of the settings panel controls how the graph itself looks, including the time span and style.

Time Period
The length of time to show on the graph. The default ranges cover from 1 hour up to 4 years, or a *Custom* period may be chosen. Selecting *Custom* displays the **Custom Period** controls. All the periods are displayed even if there is no data in a graph database going back that far. The graph will be empty for times when graphing was not active.

Resolution
The smallest slice of time for which data is available on this graph. Over time, data is consolidated over longer periods, so resolution is lost. For example, on a 1-hour graph it is possible to see data from one minute intervals, but on a graph including older data, it is not possible to show data that accurately since it has been averaged out. Depending on the time period of the graph it may contain *1 Minute*, *5 Minute*, *1 Hour*, or *1 Day* averages for data. Resolutions which are not possible for the given time period cannot be selected.

Inverse
Used on graphs such as the traffic graph, to separate incoming and outgoing data. For example, with **Inverse** set to *On*, outbound data is represented as a negative value to more easily differentiate it from inbound data.

Custom Period
When **Time Period** is set to *Custom*, the GUI displays this section to configure the custom time period for the graph.

Start Date
The start date for the graph. Clicking in the field will show a calendar date picking control. Only today, or days in the past, may be selected.

Start Hour
The hour of the day to start the graph using 24-hour style (`0-23`).

End Date
The end date for the graph.

End Hour
The end hour for the graph, exclusive. The chosen hour is not included in the graph. For example, on a graph starting at hour `10` to hour `12`, the graph covers 10:00am to 12:00pm.

Settings

Click  **Show Advanced** to display additional advanced controls not typically required for average use.

Export as CSV
Click this button to download the data from the graph as a `.csv` (Comma Separated Values) spreadsheet file, which can then be imported into another program for analysis.

Save as Defaults
Click this button to store the current graph settings as the default configuration, so this specific graph will be displayed by default on future visits to this page.

Disable/Enable Graphing
This toggle will disable or enable the collection of graph data. Graphing is enabled by default. Normally this would only be disabled for diagnostic purposes or if all required graphing is handled externally.

Reset Graphing Data
Clicking this button will erase all graph database files and create new, empty files.

Click  **Update Graphs** to change the graph to the selected view.

Graph Category List

There are a several different categories of graph data that the firewall can plot. Each category is covered here, but not all categories will be visible on every firewall. Some graphs must be enabled separately or will only be present if a specific feature or piece of hardware is enabled.

System Graphs

The graphs under the *System* category show a general overview of the system utilization, including CPU usage, memory usage, and firewall states.

Mbuf Clusters

The **Mbuf Clusters** graph plots the network memory buffer cluster usage of the firewall. Firewalls with many interfaces, or many CPU cores and NICs that use one interface queue per core, can consume numerous network memory buffers. In most cases, this usage will be fairly flat, but depending on various circumstances, such as unusually high load, the values may increase. If the usage approaches the configured maximum, increase the number of buffers.

See also:

Refer to *Hardware Tuning and Troubleshooting* for information on how to increase the amount of mbufs available to the OS.

The **Mbuf Clusters** graph contains the following data sources:

Current
: The current number of consumed mbuf clusters

Cache
: The number of cached mbuf clusters

Total
: The total of Current and Cache

Max
: The maximum allowed number of mbuf clusters

Memory Graph

The **Memory** graph shows the system RAM usage broken down into multiple areas. These areas are described in detail at *Memory Management*.

Active
: Active (in use) memory pages referenced by userland (non-kernel).

Inactive
: Memory pages which were in use but have not been referenced recently.

Free
: Memory available for immediate use.

Cache
: Memory used by the operating system for caching. On systems using ZFS, this is the ZFS ARC cache (23.05+). On UFS systems, it is the UFS directory hash.

  See also:

  ZFS Tuning.

Wire
: Memory allocated by the kernel, including the kernel itself, which cannot be paged/swapped and cannot be freed until explicitly released.

  Note: In the OS, the ZFS ARC cache and UFS buffers sizes are included in wired memory. In the graphs on pfSense Plus software version 23.05 and later, however, these values are removed from the Wired total and graphed separately. ZFS ARC usage is graphed under Cache and UFS buffers are graphed under Buffers.

UserWire
: Similar to Wired, but memory wired by user processes, not the kernel.

Laundry
Memory pages which are considered "dirty" and are due to be "cleaned".

Buffers
Memory used for UFS buffers.

Processor Graph

The processor graph shows CPU usage for the firewall using the following data sources:

User Utilization
The amount of processor time consumed by user processes.

Nice Utilization
The amount of processor time consumed by processes with a high priority.

System Utilization
The amount of processor time consumed by the operating system and kernel.

Interrupts
The amount of processor time consumed by interrupt handling, which is processing hardware input and output, including network interfaces.

Processes
The number of running processes.

States Graph

The states graph shows the number of system states but also breaks down the value in several ways.

State Changes
The number of state changes per second, or "churn". A high value from this source would indicate a rapid number of new or expiring connections.

Filter States
The total number of state entries in the states table.

NAT States
The total number of state entries involving NAT (e.g. outbound NAT, port forwards, 1:1 NAT, etc.).

Source Addresses
The number of active unique source IP addresses.

Destination Addresses
The number of active unique destination IP addresses.

Traffic Graphs

Traffic graphs show the amount of bandwidth used on each available interface in *bits per second* notation. The **Graph** list contains entries for each assigned interface, as well as IPsec and individual OpenVPN clients and servers.

The traffic graph is broken down into several data sources. Aside from the total, each has an IPv4 and IPv6 equivalent. The IPv6 data sources have `6` appended to the name.

inpass
The rate of traffic entering this interface that was *passed* into the firewall.

outpass
The rate of traffic leaving from this interface that was *passed* out of the firewall.

inblock
The rate of traffic attempting to reach this interface that was *blocked* from entering the firewall.

outblock
The rate of traffic attempting to leave this interface that was *blocked* from leaving the firewall.

inpass total
The total rate of traffic (IPv4 and IPv6) that was passed inbound.

outpass total
The total rate of traffic (IPv4 and IPv6) that was passed outbound.

Note: The terms "inbound" and "outbound" on these graphs are from the perspective of the firewall itself. On an external interface such as a WAN, "inbound" traffic is traffic arriving at the firewall from the Internet and "outbound" traffic is traffic leaving the firewall going to a destination on the Internet. For an internal interface, such as LAN, "inbound" traffic is traffic arriving at the firewall from a host on the LAN, likely destined for a location on the Internet and "outbound" traffic is traffic leaving the firewall going to a host on the LAN.

Packet Graphs

The packet graphs work much like the traffic graphs and have the same names for the data sources, except instead of reporting based on bandwidth used, it reports the number of *packets per second* (PPS) passed. The **Graph** list contains entries for each assigned interface, as well as IPsec and individual OpenVPN clients and servers.

Packets Per Second (PPS) is a better metric for judging hardware performance than Traffic throughput as it more accurately reflects how well the hardware handles packets of any size. A circuit may be sold on a certain level of bandwidth, but hardware is more likely to be bottlenecked by an inability to handle a large volume of small packets. In situations where the hardware is the limiting factor, the **Packets** graph may show a high plateau or spikes while the traffic graph shows usage under the rated speed of the line.

Quality Graphs

The **Quality** category contains **Graph** entries that track the quality of WAN or WAN-like interfaces such as interfaces with a gateway specified or those using DHCP or PPPoE. The firewall contains one **Graph** entry per gateway, including gateways that were configured previously, but no longer exist. Graph data files for old gateways are not automatically removed so that historical data is available for future reference.

The following data sources are used to track gateway reliability:

Packet Loss
The percentage of attempted pings to the monitor IP address that were lost. Loss on the graph indicates connectivity issues or times of excessive bandwidth use where pings were dropped.

Delay Average
The average delay (Round-trip time, RTT) on pings sent to the monitor IP address. A high RTT means that traffic is taking a long time to make the round trip from the firewall to the monitor IP address and back. A high RTT could be from a problem on the circuit or from high utilization.

Delay Standard Deviation
The standard deviation on the RTT values. The standard deviation gives an impression of the variability of the RTT during a given calculation period. A low standard deviation indicates that the

connection is relatively stable. A high standard deviation means that the RTT is fluctuating up and down over a large range of values, which could mean that the connection is unstable or very busy.

Captive Portal

The **Captive Portal** category contains **Graph** entries for each Captive Portal zone, past and present. Graph data files for old zones are not automatically removed.

Concurrent
The *Concurrent* graph choice shows how many users are logged in at a given point in time. As users log out or their sessions expire, this count will go down. Many concurrent users will not necessarily cause a strain on the portal, but it can be useful for judging overall capacity and bandwidth needs.

Logged In
The *Logged In* graph shows the number of login events that occur during each polling interval. This is useful for judging how busy the captive portal daemon is at a given point in time. Many users logging in around the same time will put more stress on the portal daemon compared to logins that are spread out over the course of a day.

NTP

The **NTP** graph displays statistics about the NTP service and clock quality. This graph is disabled by default because it is not relevant for most use cases. The graph can be enabled at **Services > NTP**. On that page, check **Enable RRD Graphs of NTP statistics**.

See also:

For more information about these values, see the NTP Configuration Manual, NTP Query Manual, and the NTPv4 Specification.

Offset
Combined clock difference between from server relative to this host.

System Jitter (sjit)
Combined system jitter, which is an estimate of the error in determining the offset.

Clock Jitter (cjit)
Jitter computed by the clock discipline module.

Clock Wander (wander)
Clock frequency stability expressed in parts per million (PPM)

Frequency Offset (freq)
Offset relative to hardware clock (In PPM)

Root Dispersion (disp)
Total difference between the local clock and the primary reference clock across the network.

Queue/Queuedrops Graphs

The queue graphs are a composite of each traffic shaper queue. Each individual queue is shown, represented by a unique color.

The **Queues** category shows individual queue usage in bytes.

The **QueueDrops** category shows a count of packet drops from each queue.

DHCP

The **DHCP** category contains a graph for each interface with a DHCP server enabled. The data sources shown for DHCP are:

Leases
: The number of leases in use out of the configured DHCP range for the interface.

Static Leases
: The number of static mapping leases configured for the interface.

DHCP Range
: The total size of the DHCP pool available for use on the interface.

If the **Leases** count approaches the **Range** value, then a larger pool may be required for the interface. Static mappings exist outside the range, so they do not factor into the amount of leases consumed in the pool.

Cellular

On select 3G/4G devices, the firewall is able to collect signal strength data for the **Cellular** graph. The signal strength is the only value plotted on the graph.

Wireless

The **Wireless** category is present on systems containing an 802.11 wireless network device that is enabled and in-use as a client (Infrastructure, BSS mode). The following data sources are collected and displayed when acting as a wireless client:

SNR
: The signal-to-noise ratio for the AP the client is connected to.

Channel
: The wireless channel number used to reach the AP.

Rate
: The wireless data rate to the AP.

VPN Users

The **VPN Users** category shows the number of OpenVPN users logged in concurrently for each individual OpenVPN server.

28.2.2 Traffic Graphs

Real time traffic graphs drawn with JavaScript using NVD3 are available which update continually. These graphs can be viewed at **Status > Traffic Graph**, and an example of the graph can be found in Figure *Example LAN Graph*.

These traffic graphs show interface traffic as it happens, and give a clear view of what is happening "now" rather than relying on averaged data from the RRD graphs which are better for long-term views.

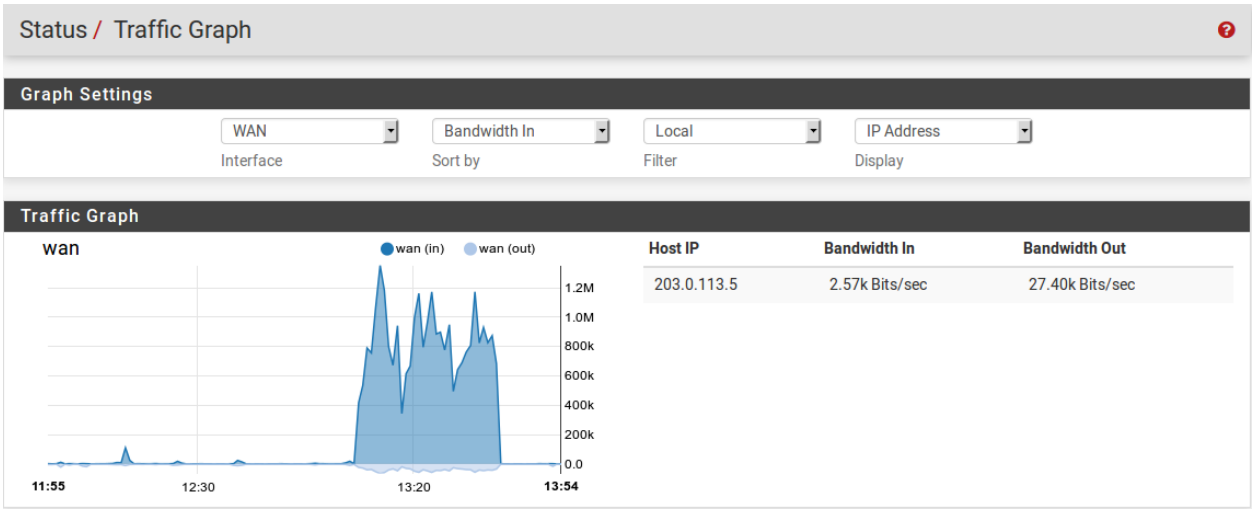


Fig. 19: Example LAN Graph

Only one interface is visible at a time, and this interface can be changed using the **Interface** drop-down list. Once an interface is chosen, the page will automatically refresh and start displaying the new graph.

Similar style traffic graphs can also be viewed on the Dashboard by adding the **Traffic Graphs** widget. Using the widget, multiple traffic graphs can be displayed simultaneously.

See also:

For more about the Dashboard, see *Dashboard*.

A table containing momentary glimpses of data being transferring from specific IP addresses is also displayed next to the traffic graph. These are limited to only displaying briefly, so ongoing transfers are more likely to show up than quick connections. Also, only connection from within that interface's primary subnet will be shown.

The display of the graph and table can be controlled using the following options:

Interface
: The firewall interface to use as the traffic source for the graph and the table.

Sort By
: Selects the sort order of the graph, either *Bandwidth In* or *Bandwidth Out*.

Filter
: Selects which type of hosts to display in the table

Local
Shows only IP addresses within the interface network

Remote
Shows only IP addresses that are not within the interface network

All
Shows all IP addresses, inside and outside the interface network

Display
Controls the display of the **Host IP** column using one of the following choices:

IP Address
The IP address of the host.

Host Name
The short hostname that corresponds to the IP address, as listed in DHCP static mappings, DNS Resolver host overrides, or DNS Forwarder host overrides.

Description
The description that corresponds to the IP address, as listed in DHCP static mappings, DNS Resolver host overrides, or DNS Forwarder host overrides.

FQDN
The fully qualified domain name that corresponds to the IP address, as listed in DHCP static mappings, DNS Resolver host overrides, or DNS Forwarder host overrides.

28.2.3 Monitoring Bandwidth Usage

With pfSense® software, there are several methods for monitoring bandwidth usage, with different levels of granularity.

pftop

If a connection is currently active, connect to the firewall console (physical access or ssh) and watch the traffic flow with pftop (Option `9`).

The output can be changed to show several views (press `0`-`8` or `v` to cycle) and may be sorted in various ways. Press `?` for a list of available command keys while running pftop.

iftop

Run `iftop` from the shell (console or SSH) as follows:

```
iftop -nNpPi em0
```

Change `em0` to an appropriate interface to monitor.

In the above example, `-nNpP` tells iftop to not resolve hostnames (`n`) or port numbers (`N`), and to run in promiscuous mode (`p`) and also display ports in the output (`P`).

Press `t` to cycle through various views.

trafshow

Another option for viewing real time throughput is trafshow, which can be installed from the CLI with `pkg install trafshow` followed by `rehash`.

Once installed, run it at an SSH command prompt:

```
trafshow
```

Then select the interface.

Built-in Graphs

If overall per-interface usage is all that is required, there are built-in RRD graphs in pfSense software, which can be found under **Status > Monitoring**.

BandwidthD

If more detail is required, such as by client IP on the LAN interface, there is a package for bandwidthd that can be installed under **System > Packages**. Once installed, it appears under **Services > BandwidthD**.

Darkstat

Darkstat is also available in **System > Packages**. Once installed, it appears under **Services > darkstat**. It also offers bandwidth graphs for an interface, as well as traffic to/from specific IP addresses.

ntopng

If even more detail is required, the ntopng package, which can also be found under **System > Packages**, can help. It can break down detail by IP, protocol, and so on. Once installed, it appears under **Diagnostics > ntopng**. It will even track where connections were made by local PCs, and how much bandwidth was used on individual connections.

Note: Due to the resource requirements of ntopng, it is not suited for with low CPU or RAM.

Monitoring on Multiple Interfaces

The **bandwidthd** package cannot listen on multiple interfaces.

The **darkstat** and **ntopng** packages can listen on multiple interfaces.

NetFlow

NetFlow is another option for bandwidth usage analysis. NetFlow is a standard means of traffic accounting implemented on many routers and firewalls. NetFlow collector running on a host inside the network is required to collect the data.

pfSense software can export NetFlow data to a collector using the `pflow(4)` feature in PF (Plus only, see *Firewall Packet Flow Data*) or the *softflowd* package.

Traffic Totals

Traffic Totals is another bandwidth monitoring tool available to install as a package. See *Status Traffic Totals* for more information.

See also:

- *Exporting NetFlow with softflowd*

28.3 Logs

Logs in pfSense software contain recent events and messages from daemons. These messages can be stored locally on a limited basis, or forwarded to a central logging server for long-term storage, better reporting, alerting, and so on.

28.3.1 System Logs

pfSense® software logs a lot of data by default, but does so in a manner that attempts to avoid overflowing the storage on the firewall. The GUI has pages which display and manage logs under **Status > System Logs** and the log files themselves are under `/var/log/` on the file system.

Some services, such as DHCP and IPsec, generate enough logs that they have their own log files and tabs in the GUI. This reduces clutter in the main system log and makes it easier to troubleshoot these individual services. To view other logs in the GUI, click the tab for the subsystem to view. Certain areas, such as **System**, and **VPN**, have sub-tabs with additional related options.

Log Format

pfSense® Plus software version 21.02 and pfSense CE software version 2.5.0 use plain text log files. The firewall periodically rotates these log files to keep their size in check. The rotation behavior is controlled by the log settings (*Log Rotation Settings*).

For each separate log, the firewall has is one main log file plus a number of rotated log files. The firewall compresses rotated log files by default on most installations, but this compression is disabled for systems running ZFS as ZFS already compresses this data. The GUI understands each compression option and displays and searches contents of rotated log files in addition to the main log file. This adds processing time but vastly increases the amount of log data available to the GUI.

pfSense® software versions older than 21.02/2.5.0 used a binary circular log format known as `clog` to maintain a constant log size without the need for rotation. As `syslogd` wrote new entries to a `clog` file, it removed older entries automatically. As such, the older data was lost.

Though there were multiple benefits to binary circular logs, such as restricting log file sizes, the downsides were too significant on modern systems. Among other reasons, binary circular logs were not very flexible, could not be used directly by shell utilities, were susceptible to corruption, and could not reliably store larger amounts of log data.

Furthermore, the original justification for size restrictions were primarily based on hardware choices from over a decade ago. Hardware, even embedded system hardware, is much more capable now.

Tip: If log retention is an issue for an organization, the logs can be copied to another server with syslog where they may be permanently retained or rotated with less frequency. See *Remote Logging with Syslog* later in this chapter for information about syslog.

On normal installations where logs are kept on disk, they are retained across reboots. When `/var` is in a RAM disk, the system attempts to back up the logs at shutdown and restore them when booting. If the system does not shut down cleanly, the logs will reset.

Viewing System Logs

The GUI interface to view system logs is located at **Status > System Logs**, on the **System** tab. This includes log entries generated by the host itself in addition to those created by services and packages which do not have their logs redirected to other tabs/log files.

As shown by the example entries in *Example System Log Entries*, there are log entries from several different areas in the main system log. Many other subsystems will log here, but most will not overload the logs at any one time. Typically, services which generate a significant volume of log entries have their own tab and log file.

2021-12-27 01:01:17.556514-05:00	php	26665	rc.dyndns.update: phpDynDNS: Not updating clara [redacted] A record because the IP address has not change d.
2021-12-27 01:01:17.556709-05:00	php	26665	rc.dyndns.update: phpDynDNS: Not updating clara [redacted] AAAA record because the IPv6 address has not c hanged.
2021-12-27 03:16:00.684185-05:00	ACME	47943	Checking if renewal is needed for: clara-gui
2021-12-27 03:16:00.685407-05:00	ACME	47943	Renewal number of days not yet reached.
2021-12-27 03:16:00.685519-05:00	ACME	47943	Checking if renewal is needed for: pfsense [redacted]
2021-12-27 03:16:00.685668-05:00	ACME	47943	Renewal number of days not yet reached.
2021-12-27 11:51:50.073658-05:00	php-fpm	2237	/index.php: Successful login for user 'admin' from: 198.51.100.142 (Local Database)
2021-12-27 11:54:57.678705-05:00	php-fpm	2237	/pkg_edit.php: Configuration Change:
2021-12-27 11:54:57.831290-05:00	check_reload_status	2266	Syncing firewall
2021-12-27 11:54:57.833034-05:00	php-fpm	2237	/pkg_edit.php: Beginning configuration backup to https://acb.netgate.com/save
2021-12-27 11:54:57.939999-05:00	php-fpm	2237	/pkg_edit.php: miniupnpd: Restarting service on interface: lan
2021-12-27 11:55:03.541425-05:00	sshd	56924	Accepted publickey for root from 198.51.100.142 port 60800 ssh2: RSA SHA256 [redacted]
2021-12-27 11:55:07.201307-05:00	php	36520	/usr/local/sbin/acbupload.php: End of configuration backup to https://acb.netgate.com/save (success).

Fig. 20: Example System Log Entries

Filtering Log Entries

The GUI can search and filter every log to find entries matching specific patterns. This is useful for tracking down log entries from a single service or log messages containing a specific username, IP address, and so on.

To search for log entries:

- Navigate to **Status > System Logs**
- Click the tab for the log to search

- Click  in the breadcrumb bar to open the **Advanced Log Filter** panel
- Enter the search criteria, for example, enter text or a regular expression in the **Message** field
- Click  **Apply Filter**

The filtering fields vary by log tab, but may include:

Message
: The body of the log message itself. A word or phrase may be entered to match exactly, or use regular expressions to match complex patterns.

Time
: The timestamp of the log message. Uses month names abbreviated to three letters.

Process
: The *name* of the process or daemon generating the log messages, such as `sshd` or `check_reload_status`.

PID
: The process ID number of a running command or daemon. In cases where there are multiple copies of a daemon running, such as `openvpn`, use this field to isolate messages from a single instance.

Quantity
: The number of matches to return in filter results. Setting this value higher than the number of log entries in the log file will have no effect, but setting it higher than the current display value will temporarily show more log messages.

The **Firewall** log tab has a different set of filtering fields:

Source IP Address
: The source IP address listed in the log entry.

Destination IP Address
: The destination IP address listed in the log entry.

Pass
: Check this option to only match log entries that passed traffic.

Block
: Check this option to only match log entries that blocked traffic.

Interface
: The friendly description name of the interface to match (e.g. `WAN`, `LAN`, `OPT2`, `DMZ`.)

Source Port
: The source port of the log entry to match, if the protocol uses ports.

Destination Port
: The destination port of the log entry to match, if the protocol uses ports.

Protocol
: The protocol to match, such as TCP, UDP, or ICMP.

Protocol Flags
: For TCP, this field matches the TCP flags on the log entry, such as `SA` (SYN+ACK) or `FA` (FIN+ACK)

Tip: The filter pane is hidden by default, but the **Log Filter** setting under **System > General Setup** can alter this behavior. When set, the GUI expands the filter pane on the log pages at all times.

See also:

- *Log Settings*
- *Remote Logging with Syslog*
- *Working with Log Files*
- *Adjusting the Size of Log Files*

28.3.2 Log Settings

Log settings on pfSense® software may be adjusted in two different ways:

- Globally at **Status > System Logs** on the **Settings** tab
- On each log tab where settings can override the global defaults

  To change these settings click  in the breadcrumb bar while viewing a log.

Each of these methods will be explained in detail in this section.

The global options area contains more options than the per-log settings. Only differences will be covered in detail for the per-log settings.

Global Log Settings

The global log options under **Status > System Logs** on the **Settings** tab include:

In the GUI, the **Settings** tab under **Status > System Logs** controls how the logging system behaves.

Log Message Format
: The format of messages logged by the system log daemon (`syslogd`) for local and remote logs. Both formats are handled the same way locally, but remote syslog servers may prefer one format or the other. Check the documentation of the syslog server for details.

  BSD (RFC 3164, default)
  : The default log format used by previous versions of pfSense software and natively used by FreeBSD.

  syslog (RFC 5424, with RFC 3339 microsecond-precision timestamps)
  : A modern syslog message format with more precise timestamps. Also includes the hostname.

Forward/Reverse Display
: By default the logs are displayed in their natural order with the oldest entries at the top and the newest entries at the bottom. Some administrators prefer to see the newest entries at the top, which can be accomplished by checking this box to flip the order.

GUI Log Entries
: The number of log entries to display in the log tabs of the GUI by default. This does not limit the number of entries in the file, only what is shown on the page at the time. The default value is `50`. The actual log files may contain much more than the number of lines to display, depending on the **Log File Size**.

Raw Logs
: When checked, this setting disables log parsing, displaying the raw contents of the logs instead. The raw logs contain more detail, but they are much more difficult to read. For many logs it also stops the GUI from showing separate columns for the process and PID, leaving all of that information contained in the **Message** column.

See also:

For more information on raw firewall logs, see *Raw Filter Log Format*.

Show Rule Descriptions
Controls if, and where, the Firewall Log tab will show descriptions for rules which generated the entries. Displaying the rule descriptions causes extra processing overhead that can slow down the log display, especially in cases where the view is set to show many entries.

Don't load descriptions
When selected this choice will not display any rule descriptions. The description may still be viewed by clicking the action column icon in the firewall log view (e.g.  or ).

Display as column
The default for new installations. Adds the rule description in a separate column. This works best if the descriptions are short, or the display is wide.

Display as second row
Adds a second row to each firewall log entry containing the rule description. This choice is better for long rule descriptions or narrow displays.

Tip: If the firewall logs display slowly with rule descriptions enabled, select *Don't load descriptions* for faster performance.

Local Logging
When checked, local logs are not retained. They are not written to disk nor are they kept in memory. While this saves on disk writes, it necessitates the use of remote logging so that information is not lost. This is not a best practice, as having local logs is vital for the vast majority of use cases.

Reset Log Files
This button clears the data from all log files and reinitialize them as new, empty logs. This can be used to clear out irrelevant/old information from logs if necessary.

> **Warning:** Resetting the log files will not save the other options on the page. If options on this page have been changed, click **Save** before attempting to reset the log files.

Logging Preferences

These options control whether certain items create log entries. Administrators may wish to disable some of these options to reduce logging or perform custom logging using other methods.

Default Firewall "block" Rules
Checked by default. When enabled, the "default deny" firewall rules, which block traffic not matched by other rules, will log entries to the firewall log.

Typically, these log entries are beneficial, but in certain rare use cases they may produce undesirable log entries that are made redundant by custom block rules with logging enabled.

Default Firewall "pass" Rules
Unchecked by default. When set, the firewall will create log entries for packets matching the default outbound pass rules on interfaces.

This option will generate a large amount of log data for connections outbound from the firewall. The best practice is to only enable this for brief periods of time while performing troubleshooting or diagnostics.

Default "Bogon Networks" Block Rules
Checked by default. When checked, if an interface has **Block Bogon Networks** active, the firewall will log packets matching those rules. Uncheck to disable this logging.

Default "Private Networks" Block Rules
Checked by default. When checked, if an interface has **Block Private Networks** active, the firewall will log packets matching those rules. Uncheck to disable this logging.

Default "IPv4 link-local" Block Rules
Checked by default. When checked, the firewall will log packets blocked by the default internal rules which block IPv6 link-local packets which should never traverse a firewall or router. Uncheck to disable this logging.

See also:

Allow APIPA

Hosts blocked by IDS
Checked by default. When checked, the firewall will log packets blocked by IDS software, such as Snort or Suricata.

While these logs are a best practice for security, they can be very high volume.

Web Server
Checked by default. When checked, the firewall will log messages from the Web GUI and Captive Portal web server processes (`nginx`).

On occasion, especially with Captive Portal active, these messages can be frequent but irrelevant and clutter the log contents.

Configuration Changes
Checked by default. When checked, the firewall creates a log entry any time a configuration change occurs. The log message includes the description of the configuration change when possible.

Click **Save** to store the new settings. The remaining options on this screen are discussed in *Remote Logging with Syslog*.

Log Rotation Settings

Starting with pfSense Plus software version 21.02 and pfSense CE software version 2.5.0, the system logs are kept in a plain text format and periodically rotated. The options in this section control how the firewall handles log rotation.

Note: The options in this section of the page are global only, and cannot be changed for individual logs.

Log Rotation Size (Bytes)
This field controls the size at which the firewall rotates logs. The default size is `500` KiB per log file. There are nearly 20 log files, so plan space accordingly.

This **does not** account for space used by rotated log files.

Note: Increasing this value allows every log file to grow to the specified size, so disk usage can increase significantly. The firewall checks log file sizes once per minute to determine if rotation is necessary, so a rapidly growing log file may exceed this value.

Log Compression
The type of compression the firewall uses when rotating log files. Compressing rotated log files saves disk space, and the compressed logs remain available for display and searching in the GUI. Though processing large compressed files can be time-consuming, most use cases will not notice significant slowness.

The types of compression available are `bzip2` (default), `gzip`, `xz`, `zstd`, and `none` (disables compression). All the options which use compression are reasonably fast and offer good compression rates. Some may compress better than others, others are slightly faster, but ultimately the decision is up to the environment and the administrators.

> **Warning:** The type of compression used by all log files must be identical. When changing this value, the firewall must remove all previously rotated compressed log files.

On certain systems, disabling compression (set to `none`) is the best course of action. Examples include:

- Firewalls using large log file sizes, which may take too long to compress
- Slower firewalls which may take too long to compress or search the log files even at default sizes
- Firewalls using ZFS which by default will already compress disk contents

Log Retention Count
The number of rotated log files to keep before the oldest copy is removed. Keeping more log files will consume more disk space, but compressed logs files do not consume nearly as much space as decompressed logs.

Per-Log Settings

To change per-log settings, visit the log tab to change and then click  in the breadcrumb bar to expand the settings panel.

On this panel, several options are displayed. Most of the options will show the global default value or have a **General Logging Options Settings** choice which will use the global value and not the per-log value.

The per-log settings panel for each tab only displays options relevant to that log. For example, the options to log default block or pass rules are displayed only when viewing the **Firewall** log tab.

Each per-log settings panel has at least the following options: **Forward/Reverse Display**, **GUI Log Entries**, and **Formatted/Raw Display**. For each of these, a value which will only apply to this log may be set.

See also:

For more information on how these options work, see *Global Log Settings* above.

Click **Save** to store the new log settings.

See also:

- *Remote Logging with Syslog*
- *Accessing Firewall Services over IPsec*
- *Working with Log Files*
- *Adjusting the Size of Log Files*

28.3.3 Remote Logging with Syslog

The **Remote Logging** options under **Status > System Logs** on the **Settings** tab enable `syslog` to copy log entries to a remote server.

The logs kept by pfSense® software on the firewall itself are of a finite size. Copying these entries to a syslog server can aid troubleshooting and allow for long-term monitoring. Having a remote copy can also help diagnose events that occur in other scenarios. For example, before a firewall restarts, when older entries are cycled out of the log, after administrators clear local logs, and cases when local storage has failed while the network remains active.

Warning: Corporate or local legislative policies may dictate the length of time an organization must retain log data from firewalls and similar devices. If an organization requires long-term log retention for their own or government purposes, a remote syslog server is required to receive and retain these logs.

Warning: Logs sent using this method are delivered in the clear (not encrypted) unless the logs are sent through a VPN or using a mechanism such as *Stunnel package*. As an alternative, consider using the **syslog-ng** package which is capable of sending encrypted syslog messages.

The following options are available for remote logging:

Source Address
Controls where the `syslog` daemon binds for sending out messages. In most cases, the default (*Any*) is the best option, so the firewall will use the address nearest the target. If the destination server is across a tunnel mode IPsec VPN, however, choosing an interface or Virtual IP address inside the local Phase 2 network will allow the log messages to flow properly over a tunnel.

IP Protocol
When choosing an interface for the **Source Address**, this option gives the `syslog` daemon a preference for either using IPv4 or IPv6, depending on which is available. If there is no matching address for the selected type, the other type is used instead.

Remote Log Servers
Enter up to three remote servers using the boxes contained in this section. Each remote server can use either an IP address or hostname, and an optional UDP port number. If the port is not specified, the default *syslogd* port, `514`, is assumed.

A syslog server is typically a server that is directly reachable from the firewall on a local interface. Logging can also be sent to a server across a VPN.

Warning: Do not send log data directly across any WAN connection or unencrypted site-to-site link, as it is plain text and could contain sensitive information.

Note: The syslog daemon can only send messages over UDP. To send syslog messages over TCP, consider using the **syslog-ng** package.

Remote Syslog Contents
The options in this section control which log messages will be sent to the remote log server.

Everything
When set, all log messages from all areas are sent to the server.

System Events
Main system log messages that do not fall into other categories.

Firewall Events
Firewall log messages in raw format. The format of the raw log is covered in *Raw Filter Log Format*.

DNS Events
Messages from the DNS Resolver (`unbound`), DNS Forwarder (`dnsmasq`), and from the `filterdns` daemon which periodically resolves hostnames in aliases.

DHCP Events
Messages from the IPv4 and IPv6 DHCP daemons, relay agents, and clients.

PPP Events
Messages from PPP WAN clients (PPPoE, L2TP, PPTP)

General Authentication Events
Log messages about authentication events, such as for the GUI or certain types of VPNs.

Captive Portal Events
Messages from the Captive Portal system, typically authentication messages and errors.

VPN Events
Messages from VPN daemons such as IPsec and OpenVPN, as well as the L2TP server and PPPoE server.

Gateway Monitor Events
Messages from the gateway monitoring daemon, `dpinger`

Routing Daemon Events
Routing-related messages such as UPnP IGD & PCP, IPv6 routing advertisements, and routing daemons from packages like OSPF, BGP, and RIP.

Network Time Protocol Events
Messages from the NTP daemon and client.

Wireless Events
Messages from the Wireless AP daemon, `hostapd`.

To start logging remotely:

- Navigate to **Status > System Logs** on the **Settings** tab
- Check **Send log messages to remote syslog server**
- Configure the options as described above
- Click Save to store the changes.

If a syslog server is not already available, it is fairly easy to set one up. Almost any UNIX or UNIX-like system can be used as a syslog server. FreeBSD is described in the following section, but others may be similar.

Setup Syslog on the Logging Host

FreeBSD

First, configure the syslog server to accept remote connections which means running it with the `-a <subnet>` or similar flag.

On FreeBSD, edit `/etc/rc.conf` and add this line:

```
syslogd_flags=" -a 192.168.1.1 "
```

Where `192.168.1.1` is the IP address of the pfSense firewall.

More complex allow rules for syslog are also possible, like so:

```
syslogd_flags=" -a 10.0.10.0/24:*"
```

Using that parameter, syslog will accept from any IP address in the 10.0.10.0 subnet (mask 255.255.255.0) and the messages may come from any UDP port.

Now, edit `/etc/syslog.conf` and add a block at the bottom:

```
!*
+*

+pfsense
*.*                /var/log/pfsense.log
```

Where `pfSense` is the hostname of the pfSense firewall. An entry may also need to be added in **/etc/hosts** for that system, depending on the DNS setup. Logs may be split separate files. Use the **/etc/syslog.conf** file on the pfSense firewall for more details on which logging facilities are used for specific items.

```
192.168.1.1            pfsense     pfsense.example.com
```

The log file may also need to be created manually with proper permissions:

```
# touch /var/log/pfsense.log
# chmod 640 /var/log/pfsense.log
```

Now restart syslog:

```
# /etc/rc.d/syslogd restart
```

Windows

Setting this up on Windows entirely depends on which syslog server is being used. Consult the documentation for more information on configuration.

There is a free multipurpose utility that can act as a syslog server, which can be found here: http://tftpd32.jounin.net/

Kiwi Syslog Server is free for up to 5 devices. https://www.solarwinds.com/free-tools/kiwi-free-syslog-server

Linux

Configuration of the system logger on Linux depends on the distribution. Consult the distribution's documentation on how to change the behavior of **syslogd**. It should be similar in many cases to the alterations in the FreeBSD section.

OpenBSD

The configuration for OpenBSD is similar to FreeBSD, with the following notes:

1. The option to accept remote syslog events is `-u`.
2. This option may be enabled using *rcctl(8)*:

```
# rcctl set syslogd flags -u
```

1. To restart the syslogd service:

```
# rcctl restart syslogd
```

Other Logging Servers

Other log systems or styles such as Splunk, ELSA (Enterprise Log Search and Archive), Graylog, ELK (Elasticsearch, Logstash, and Kibana), or OpenSearch (open source fork of ELK components) may also be used but the methods for implementing them are beyond the scope of this document. If such a system is syslog-compatible, then the pfSense software side should be fairly simple to set up as it would be for any other syslog system.

28.3.4 Adjusting the Size of Log Files

pfSense® software manages log files automatically and attempts to limit their size. The default size is 500 KiB per log file, and there are around 20 log files.

When increasing log sizes, keep disk space in mind. There is a disk space indicator for the filesystem containing the logs under the **Log Rotation Size (Bytes)** text description on *Log Settings*.

Tip: These log files are held in `/var/log` which may optionally be a RAM disk.

Recent versions of pfSense software (pfSense Plus software version 21.02, pfSense CE software version 2.5.0) and later use plain text log files and log rotation keeps the size in check. Space for rotated log files is **in addition to** the default log size limit.

For example, the firewall will keep multiple rotated copies of the log by default, but rotation is triggered by the size of the main log file. If the firewall keeps 7 rotated log files in addition to the main log, and has disabled compression for rotated log files, then the actual consumed space for logs could be up to 8 times the rotation size. The GUI displays a "worst case" disk usage amount under the **Log Rotation Size (Bytes)** text description on *Log Settings*.

See also:

Log rotation settings are covered in *Log Rotation Settings*.

The GUI only shows 50 lines per log by default, but the files contain many more entries. See *Log Settings* for more information on that setting.

Short Version

To change the log file sizes:

- Navigate to **Status > System Logs**, **Settings** tab
- Enter a new value in **Log Rotation Size (Bytes)**, being careful not to overfill the disk containing the logs.
- Click **Save**

Details

The rotation settings in *Log Rotation Settings* cover this topic thoroughly.

The default size for a log file is `512000` (500KiB) which can generally hold between 2000-3000 log entries but varies by entry size. The time span covered by logs depends entirely on how much data is logged. A quiet log file could contain months or even years of information, a busy log file may only contain minutes.

Underneath the text for **Log Rotation Size (Bytes)** the current and available disk space is displayed based on the current log file sizes and their location. For example:

```
Disk space currently used by log files: 4.3M
Worst case disk usage for base system logs based on current global settings: 58.11 MiB
Remaining disk space for log files: 9.3G
```

There are approximately 20 log files affected by the size control. The value entered in **Log Rotation Size (Bytes)** is for a single log file, so the actual usage will be approximately **20 times** that value, not including rotated log files. Increasing this value allows every log file to grow to the specified size, so disk usage can increase significantly. The firewall checks log file sizes once per minute to determine if rotation is necessary, so a rapidly growing log file may exceed this value.

As shown above, with the default value of `512000`, the firewall could consume up to around 10MB of total log space for the main log files, plus extra for the rotated copies. If the size of the logs is increased to `1024000` (1MB), then nearly 20MB could be used for logs. Be certain before changing the **Log Rotation Size (Bytes)** value that the disk has enough space to hold all the log files.

> **Warning:** Do not increase the log file size to huge values in an attempt to retain log data on the firewall long term. Not only does that consume large amounts of disk space and make it more difficult for the firewall to display and manage logs, but it is unreliable compared to properly storing log files on a remote syslog server dedicated to that purpose.

28.3.5 Working with Log Files

The format of log files is described in *Log Format*, read that section before proceeding.

pfSense® Plus software version 21.02, pfSense CE software version 2.5.0, and later versions utilize plain text log files which can be used by a variety of traditional shell utilities. There are also utilities compatible with the various types of compressed rotated log files.

pfSense® software versions older than 21.02/2.5.0 use a binary circular log format known as `clog` to maintain a constant log size without the need for rotation. As `syslogd` writes new entries to a `clog` file, it removes older entries automatically. As such, the older data is lost. These binary log files cannot be processed directly by shell utilities and must first be unwrapped with the `clog` utility.

Viewing Log Contents (21.02/2.5.0 and later)

To view the contents of a log, use common shell utilities, such as `cat`, `grep`, and so on:

```
# cat /var/log/filter.log
# grep -i "error" /var/log/system.log
```

To follow the contents of a log file in real time, use `tail -f` or `tail -F`. The latter form follows the log to a new file after rotation.

```
# tail -F /var/log/filter.log
```

In addition to the main log file, the rotated log files can be viewed and searched by passing them through utilities specific to the format with which they are compressed. For example, the default compression type is `bzip2`, so use `bzcat`, or `bzgrep`:

```
# bzcat /var/log/filter.log.0.bz2
# bzgrep -i "error" /var/log/system.log.0.bz2
```

Additional utilities can be utilized by piping the output.

The following list contains the different compression options and a sample of utilities which can parse their contents:

bzip2 (`*.log.<number>.bz2`)
: `bzcat`, `bzgrep`, `bzless`.

gzip (`*.log.<number>.gz`)
: `zcat`, `zgrep`, `zless`.

xz (`*.log.<number>.xz`)
: `xzcat`, `xzgrep`, `xzless`.

zstd (`*.log.<number>.zst`)
: `zstdcat`, `zstdgrep`, `zstdless`.

none (`*.log.<number>`)
: `cat`, `grep`, `less`, plus anything else capable of parsing text files.

Viewing Log Contents (< 21.02/2.5.0, clog)

On versions of pfSense software before 21.02/2.5.0, the contents of binary circular log files can only be read using the `clog` command:

```
# clog /var/log/filter.log
```

The output of that command may then be piped to tools like `grep`:

```
# clog /var/log/system.log | grep -i "error"
```

To follow the log files in a manner like `tail -f`, use `clog -f`:

```
# clog -f /var/log/filter.log
```

The command prints the entire contents of the log file to the console, and then prints new entries as they are written.

28.3.6 Viewing the Firewall Log

The firewall creates log entries for each rule configured to log and for various other internal rules such as default deny rules. There are several ways to view these log entries, each with varying levels of detail. There is no particular best method as it depends on preferences and skill level of the firewall administrators, though using the GUI is the easiest method.

Tip: The logging behavior of the default deny rules and other internal rules can be controlled using the **Settings** tab under **Status > System Logs**. See *Logging Preferences* for details.

Like other logs, the firewall log only retains a certain number of entries. If the needs of an organization require a permanent record of firewall logs for a longer period of time, see *Remote Logging with Syslog* for information on copying these log entries to a syslog server as they happen.

See also:

- *Troubleshooting Blocked Log Entries for Legitimate Connection Packets*
- *Working with Log Files*
- *Log Settings*
- *Logging Preferences*
- *Filtering Log Entries*

Viewing in the GUI

To view firewall logs in the GUI, navigate to **Status > System Logs**, **Firewall** tab.

By default, this page parses and renders firewall log entries in an easy-to-read format.

Action	Time	Interface	Rule	Source	Destination	Protocol
✕	Dec 15 05:02:22	WAN2	Default deny rule IPv4 (1000000103)	203.0.113.1:30069	203.0.113.129:68	UDP
✕	Dec 15 06:02:21	WAN2	Default deny rule IPv4 (1000000103)	203.0.113.1:30069	203.0.113.129:68	UDP
✕	Dec 15 07:02:22	WAN2	Default deny rule IPv4 (1000000103)	203.0.113.1:30069	203.0.113.129:68	UDP
✕	Dec 15 08:02:23	WAN2	Default deny rule IPv4 (1000000103)	203.0.113.1:30069	203.0.113.129:68	UDP
✕	Dec 15 09:02:23	WAN2	Default deny rule IPv4 (1000000103)	203.0.113.1:30069	203.0.113.129:68	UDP
✕	Dec 15 10:02:23	WAN2	Default deny rule IPv4 (1000000103)	203.0.113.1:30069	203.0.113.129:68	UDP

Fig. 21: Example Log Entries Viewed From The GUI

See also:

Several aspects of firewall log display behavior are controlled by options on the **Settings** tab. See *Log Settings* for information on how to view and change these log settings.

The parsed GUI logs, shown in Figure *Example Log Entries Viewed From The GUI*, are in multiple columns:

Action
: Icon representing the firewall rule action which resulted in the log entry (e.g. pass or block).

  Hovering the mouse over the action icon displays several details about the log entry in a popup, including which rule triggered the log entry. These details can be useful when troubleshooting rule issues.

  The details displayed by this view include:

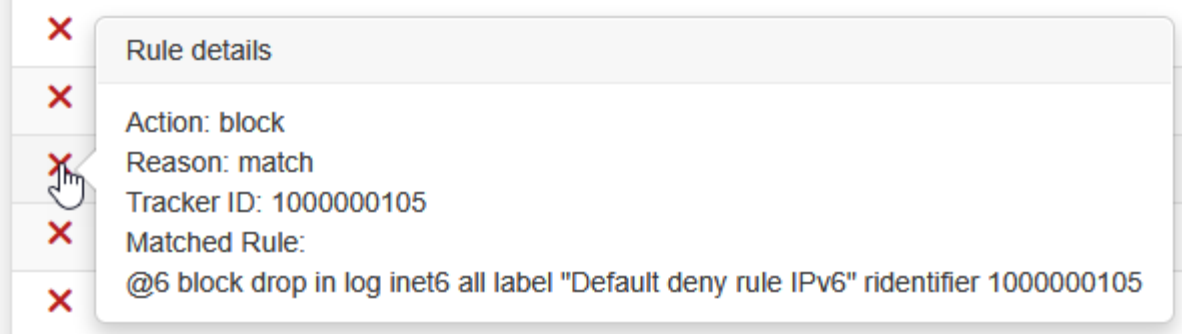


Fig. 22: Firewall Log Entry Detail

Action
Full text of the firewall rule action which triggered this log entry (e.g. Pass, Block, etc.)

Reason
The reason the log entry was generated, e.g. `match`. There are several possible reason strings which are typically self-explanatory based on the text, such as `match` being from a packet matching a rule. Reasons can also include things like state mismatches, checksum mismatches, time mismatches, hitting state limitations, or even the presence of IP options.

Tracker ID
The firewall *Rule Tracking ID* which uniquely identifies this rule.

Matched Rule
The rule which matched this packet, if the firewall was able to locate the rule.

Associated Rules
A list of rules with the same tracking ID as the rule which generated the log entry. This can help narrow down which rule may have triggered a log entry when the firewall could not locate an exact match.

Time
The time that the packet arrived.

Interface
Where the packet entered the firewall.

The GUI prints a ► character next to the interface if a rule matched a packet in the *outbound* direction. The vast majority of rules match in the inbound direction, so the direction is omitted in that case.

Rule
The firewall rule description and *rule tracking ID* which generated the log entry, if available. This column only appears when rule descriptions are configured to appear in a separate column. They may also be shown in a separate row, or disabled entirely.

See also:

- *Log Settings*
- *Rule Tracking ID*

Source
The source IP address and port.

Clicking the  icon next to the source and destination IP addresses makes the firewall perform a

DNS lookup on the IP address. If the address has a valid hostname, the page will display it underneath the IP address in all instances of that address on the page.

Clicking the  icon next to the source IP address or the  icon next to the destination IP address will add firewall rules based on those addresses using EasyRule.

See also:

Using EasyRule to Manage Firewall Rules

Destination
The destination IP address and port.

Protocol
The protocol of the packet, e.g. ICMP, TCP, UDP, etc.

Log entries for TCP packets have extra information appended to the protocol field displaying TCP flags present in the packet. These flags indicate various connection states or packet attributes.

See also:

TCP Flags

The GUI can also filter log output to find specific entries, so long as they exist in log files on the device. Click  to display the filtering options. See *Filtering Log Entries* for more information.

Viewing from the Console Menu

Option `10` from the console menu views and follows the `filter.log` in real time. An easy example is a log entry like that seen above in Figure *Example Log Entries Viewed From The GUI*:

```
Aug  3 08:59:02 master filterlog: 5,16777216,,1000000103,igb1,match,block,in,4,0x10,,128,
↪0,0,
none,17,udp,328,198.51.100.1,198.51.100.2,67,68,308
```

This single line shows that the log entry was triggered by rule id `1000000103`, which resulted in a block action on the `igb1` interface. The source and destination IP addresses are shown near the end of the log entry, followed by the source and destination port. Packets from other protocols may show significantly more data.

See also:

See *Raw Filter Log Format* for details on the format of the filter log file.

Viewing from the Shell

When using the shell, either from SSH or from the console, there are numerous options available to view the filter logs.

When directly viewing the contents of the log file, the log entries can be quite complex and verbose.

For information on viewing logs from the shell, see *Working with Log Files*.

Viewing parsed log output in the shell

There is a simple log parser written in PHP which can be used from the shell to produce reduced output instead of the full raw log. To view the parsed contents of the current log, run:

```
# cat /var/log/filter.log | filterparser.php
```

The script prints the log entries one per line, with simplified output:

```
Aug  3 08:59:02 block igb1 UDP 198.51.100.1:67 198.51.100.2:68
```

Finding the rule which caused a log entry

When viewing one of the raw log formats, the log includes the *Rule Tracking ID* for an entry. This tracking ID can be used to find the rule which caused the match.

The following example locates the rule with a tracking ID of `1000000103`:

```
# pfctl -vvsr | grep 1000000103
@4 block drop in log inet all label "Default deny rule IPv4" ridentifier 1000000103
```

As shown in the above output, this was the default deny rule for IPv4.

28.3.7 Raw Filter Log Format

The raw content of the filter log consists of single lines containing comma-separated values.

pfSense software generates this log format for its internal filter log, and the log output it transmits over `syslog` to remote hosts.

Plain text layout

This section describes the content of the raw log in general terms.

See also:

For a more technical representation with greater detail, see *BNF / Grammar*.

Note: The hostname is not included in log messages the firewall forwards to remote log hosts with the default **BSD** log message format, but it is included in the optional **syslog** (RFC5424) format. See *Log Settings* for details.

```
<Timestamp> <Hostname> filterlog: <CSV data>
```

CSV Data has many common fields and some that vary by protocol:

Common fields:

- Rule Number
- Sub rule number
- Anchor

- Tracker - unique *Rule Tracking ID* per rule, stored with the rule in `config.xml` for user added rules, or check `/tmp/rules.debug`
- Real interface (e.g. `em0`)
- Reason for the log entry (e.g. `match`)
- Action taken that resulted in the log entry (e.g. `block`, `pass`)
- Direction of the traffic (`in`/`out`)
- IP version (`4` for IPv4, `6` for IPv6)

IPv4:

- TOS
- ECN
- TTL
- ID
- Offset
- Flags
- Protocol ID
- Protocol text (TCP, UDP, etc.)

IPv6:

- Class
- Flow Label
- Hop Limit
- Protocol
- Protocol ID

IPv4 or IPv6:

- Length
- Source IP
- Destination IP

For TCP and UDP (Proto ID 6 or 17) on IPv4 or IPv6

- Source Port
- Destination Port
- Data Length

TCP Only:

- TCP Flags
- Sequence Number
- ACK
- Window
- URG

- Options

ICMP:

- ICMP Type, used to choose between the following possibilities

ICMP Echo Request/Reply

- ICMP ID
- ICMP Sequence

ICMP Protocol Unreachable

- ICMP Destination IP
- ICMP Protocol ID

ICMP Port Unreachable

- ICMP Destination IP
- ICMP Protocol ID
- ICMP Port Number

ICMP unreachable (other), time exceeded, parameter problem, redirect, mask reply:

- ICMP Description

ICMP Need Frag

- ICMP Destination IP
- ICMP MTU

ICMP tstamp

- ICMP ID
- ICMP Sequence

ICMP tstamp reply

- ICMP ID
- ICMP Sequence
- ICMP otime
- ICMP rtime
- ICMP ttime

ICMP default:

- ICMP Description

CARP (Protocol ID 112):

- Type
- TTL
- VHID
- Version
- Advskew
- Advbase

BNF / Grammar

For more technical purposes, this is a BNF format representation of the log output. It is not a 100% complete BNF as the exact contents of many of the fields are beyond the scope of this document as they aren't generally relevant to typical logging, but they are included in the log entries for completeness. Consult a reference on IP packet headers for more information.

```
<log-entry> ::= <time-stamp> <host-name> "filterlog:" <log-data>

<log-data> ::= <rule-number>,<sub-rule-number>,<anchor>,<tracker>,<real-interface>,
↪<reason>,<action>,<direction>,<ip-version>[,<ip-specific-data>]

<rule-number> ::= <integer> -- Rule number in the PF Ruleset
<sub-rule-number> ::= <integer> -- Sub rule number in the PF Ruleset (not typically␣
↪significant for general use)
<anchor> ::= <text> -- Anchor name in which the rule exists
<tracker> ::= <integer> -- Unique rule tracking ID per rule, stored with the rule in␣
↪config.xml for user added rules, or check /tmp/rules.debug
<real-interface> ::= <text> -- Real interface for the log entry (e.g. em0)
<reason> ::= <text> | "unkn(%u)" -- Reason for the log entry (typically "match")
<action> ::= "pass" | "block" | "unkn(%u)" -- Action taken that resulted in the log entry
<direction> ::= "in" | "out" | "unkn(%u)" -- Direction of the logged traffic
<ip-version> ::= "4" | "6" -- IPv4 or IPv6
<ip-specific-data> ::= (<ipv4-specific-data>|<ipv6-specific-data>),<ip-data>[,<protocol-
↪specific-data>]

<ipv4-specific-data> ::= <tos>,<ecn>,<ttl>,<id>,<offset>,<flags>,<protocol-id>,<protocol-
↪text>

<tos> ::= <empty> | <hex> -- Type of Service identification
<ecn> ::= <empty> | -- Explicit Congestion Notification
<ttl> ::= <integer> -- Time To Live (TTL) of the packet
<id> ::= <integer> -- ID of the packet
<offset> ::= <integer> -- Fragment offset
<flags> ::= "none" | <text> -- IP Flags (NOT TCP flags -- those are later)
<protocol-id> ::= <integer> -- IP protocol ID (e.g. 6 for TCP, 17 for UDP)
<protocol-text> ::= "tcp" | "udp" | "icmp" | <text> -- IP protocol text (examples given)
<ipv6-specific-data> ::= <class>,<flow-label>,<hop-limit>,<protocol-text>,<protocol-id>

<class> ::= <hex> -- ToS traffic class
<flow-label> ::= <data> -- Flow label
<hop-limit> ::= <integer> -- Hop Limit (similar to IPv4 TTL)
<protocol-text> ::= "tcp" | "udp" | "icmp" | <text> -- IP protocol text (examples given)
<protocol-id> ::= <integer> -- IP protocol ID (e.g. 6 for TCP, 17 for UDP)
<ip-data> ::= <length>,<source-address>,<destination-address>

<length> ::= <integer> -- Length of the packet in bytes
<source-address> ::= <ip-address> -- The source IP address of the logged traffic
<destination-address> ::= <ip-address> -- The destination IP address of the logged␣
↪traffic
<protocol-specific-data> ::= <tcp-data> | <udp-data> | <icmp-data> | <carp-data>

<tcp-data> ::= <source-port>,<destination-port>,<data-length>,<tcp-flags>,<sequence-
```

(continues on next page)

(continued from previous page)

```
↪number>,<ack-number>,<tcp-window>,<urg>,<tcp-options>

<source-port> ::= <integer> -- Source port number
<destination-port> ::= <integer> -- Destination port number
<data-length> ::= <integer> -- Data/payload length
<tcp-flags> ::= [S][A][.][F][R][P][U][E][W] -- TCP Flags
<sequence-number> ::= <integer> -- TCP Sequence ID
<ack-number> ::= <integer> -- ACK number
<tcp-window> ::= <integer> -- Windows size
<urg> ::= <data> -- Urgent pointer data
<tcp-options> ::= <data> -- TCP Options
<udp-data> ::= <source-port>,<destination-port>,<data-length>

<icmp-data> ::= <icmp-type>,(<echo-data> | <unreachproto-data> | <unreachport-data> |
↪<other-unreachable-data> | <needfrag-data> | <tstamp-data> | <tstampreply-data> |
↪<icmp-default-data>)

<icmp-type> ::= <echo-type> | "unreachproto" | "unreachport" | <other-unreachable> |
↪"needfrag" | "tstamp" | "tstampreply" | <text>
<echo-type> ::= "request" | "reply"
<other-unreachable> ::= "unreach" | "timexceed" | "paramprob" | "redirect" | "maskreply"
<echo-data> ::= <icmp-id>,<icmp-sequence>

<icmp-id> ::= <integer> -- ID of the echo request/reply
<icmp-sequence> ::= <integer> -- Sequence number of the echo request/reply
<unreachproto-data> ::= <icmp-destination-ip-address>,<unreachable-protocol-id>

<icmp-destination-ip-address> ::= <ip-address> -- Original destination address of the␣
↪connection that caused this notification
<unreachable-protocol-id> ::= <integer> -- Protocol ID number that was unreachable
<unreachport-data> ::= <icmp-destination-ip-address>,<unreachable-protocol-id>,
↪<unreachable-port-number>

<unreachable-port-number> ::= <integer> -- Port number that was unreachable
<other-unreachable-data> ::= <icmp-description>

<icmp-description> ::= <text> -- Description from the ICMP packet
<needfrag-data> ::= <icmp-destination-ip-address>,<icmp-mtu>

<icmp-mtu> ::= <integer> -- MTU to use for subsequent data to this destination
<tstamp-data> ::= <icmp-id>,<icmp-sequence>

<tstampreply-data> ::= <icmp-id>,<icmp-sequence>,<icmp-otime>,<icmp-rtime>,<icmp-ttime>

<icmp-otime> ::= <unix-timestamp> -- Originate Timestamp
<icmp-rtime> ::= <unix-timestamp> -- Receive Timestamp
<icmp-ttime> ::= <unix-timestamp> -- Transmit Timestamp
<icmp-default-data> ::= <icmp-description>

<carp-data> ::= <carp-type>,<carp-ttl>,<vhid>,<version>,<advbase>,<advskew>

<carp-type> ::= <text> -- Type of CARP/VRRP
```

(continues on next page)

(continued from previous page)

```
<carp-ttl> ::= <integer> -- Time to Live
<vhid> ::= <integer> -- Virtual Host ID
<version> ::= <integer> -- CARP Version
<advbase> ::= <integer> -- Advertisement base timer interval (seconds)
<advskew> ::= <integer> -- Advertisement skew (1/256 of a second)
```

28.3.8 Gateway Logs

The gateway logs can be found through the pfSense® software GUI under **Status > System Logs** on the **System/Gateways** sub-tab.

This log contains entries from the gateway monitoring daemon, *dpinger*, which can generate a significant amount of logging with many gateways to monitor.

The entries found here will record events such as when a gateway is down, or in an alarm state, or has returned to an online state.

Interpreting Gateway Logs

The gateway logs contain several types of messages that some users may find initially confusing. The most common log messages are:

Monitoring startup
: These messages are logged when gateway monitoring starts or restarts, one line per gateway.

Note: These logs are normal and **do not** indicate a problem.

```
dpinger[34996]: send_interval 500ms  loss_interval 2000ms  time_period 60000ms
  report_interval 0ms  data_len 1  alert_interval 1 000ms  latency_alarm 500ms
  loss_alarm 20%  dest_addr 198.51.100.1 bind_addr 198.51.100.10 identifier "WAN_GW"
```

Gateway Alarm
: This type of message indicates that a gateway entered an "Alarm" state. For example, due to connectivity loss between the firewall and the monitor address.

```
dpinger[11000]: WAN_GW 198.51.100.1: Alarm latency 4807us stddev 1790us loss 21%
```

Gateway Recovery
: This type of log message indicates a gateway is recovering from an Alarm state and is now available for use.

```
dpinger[11000]: WAN_GW 198.51.100.1: Clear latency 4464us stddev 1028us loss 5%
```

28.3.9 NTP Logs

The NTP daemon Log contains logs generated by the *Network Time Protocol daemon* and other time-related actions.

28.3.10 Package Logs

This tab contains messages from packages which implement logging to this central location.

Tip: If logs from a specific package are not present, contact the package maintainer to see if the package can be updated to utilize this mechanism.

Some packages will log to the main system log or a related tab inside the system logs (**Status > System Logs**). Others may keep their own logs in a separate location. Some packages, such as Snort, offer configuration options to control where and how logs are made. Some logs may need to be viewed outside the GUI or via **Diagnostics > Command**.

28.3.11 PPP Logs

The **PPP** logs tab displays any events from the PPP system for WAN type connections, not locally-hosted servers. This would be for WANs that connect using PPPoE, L2TP, Cellular networks, and so on.

28.3.12 Resolver Logs

The Resolver logs are located at **Status > System Logs** on the **System/DNS Resolver** tab.

This log contains entries from DNS-related processes. These include the DNS Resolver (Unbound), DNS Forwarder (`dnsmasq`), the `filterdns` process that monitors for updates in hostnames for Aliases/IPsec/etc., and the BIND package.

28.3.13 Routing Logs

The Routing logs are located at **Status > System Logs** on the **System/Routing** tab.

This log contains entries from routing-related processes for both IPv4 and IPv6, including:

- `radvd` (IPv6 Router Advertisements)
- `zebra` (FRR)
- `ospfd` (FRR)
- `bgpd` (FRR)
- `bfdd` (FRR)
- `pimd` (PIMD)
- `igmpproxy` (IGMP Proxy)
- `miniupnpd` (UPnP IGD & PCP)

28.3.14 IPsec Logs

The IPsec log shows output from strongSwan components such as the IPsec daemon `charon`. This log contains output for successful connections, normal ongoing activity such as DPD checks, and errors.

Troubleshooting IPsec VPNs contains example entries and guidance for interpreting the meaning of log messages.

28.3.15 OpenVPN Logs

Logs for OpenVPN are located in the GUI under **Status > System Logs** on the **OpenVPN** tab.

These logs include output from the OpenVPN daemon(s) in use, both clients and servers. Log messages include entries for successful connections as well as failures and errors.

If there are no log entries for a server after the process starts, traffic likely is not reaching the OpenVPN daemon. Check the WAN-side firewall rules and the address/port used by the client.

See also:

Troubleshooting OpenVPN.

28.3.16 Captive Portal Authentication Logs

The Captive Portal Authentication Logs are available through the GUI at **Status > System Logs**, on the **Portal Auth** tab. The logs list login information from the **Captive Portal** system.

28.3.17 Wireless Logs

The Wireless logs can be found in the GUI under **Status > System Logs** on the **System/Wireless** tab.

These logs contain entries from the `hostapd` daemon which handles wireless access point connections. This process can be overly verbose when handling client traffic, logging rekey events and other information that can otherwise clutter the main system log.

28.3.18 L2TP Logs

A record of login and logout events is kept on **Status > System Logs**, on the **VPN** tab, under **L2TP Logins**.

Each login and logout is recorded with a timestamp and username, and each login will also show the IP address assigned to the L2TP client. The full log can be found on the **L2TP Raw** tab.

28.3.19 DHCP Logs

The DHCP log view at **Status > System Logs** on the **DHCP** Tab, displays messages and events from the DHCP Daemon and the DHCP client for WANs.

Each DHCP request and reply from DHCP clients is shown here, along with events and errors. IP addresses, MAC addresses, and client-supplied hostnames are all visible in the logs.

See also:

- *Troubleshooting "login on console as root" Log Messages*
- *Troubleshooting "promiscuous mode enabled" Log Messages*
- *Troubleshooting ARP Move Log Messages*

CHAPTER

TWENTYNINE

DIAGNOSTICS

These documents cover functions found under the **Diagnostics** menu in pfSense® software.

29.1 DNS Lookup

Diagnostics > DNS Lookup performs simple forward and reverse DNS queries. These queries obtain information about an IP address or hostname and also test the DNS servers configured on the firewall (*DNS Server Settings*).

To perform a DNS Lookup:

- Navigate to **Diagnostics > DNS Lookup**
- Enter a **Hostname** or IP address to query
- Click  **Lookup**

The page displays the results of the DNS query along with timing information and options.

29.1.1 DNS servers included in testing

The page will query a specific set of DNS servers. This set depends upon the *DNS Server Settings* under **System > General**.

The page will test against `127.0.0.1` if the DNS Resolver or DNS Forwarder are active and the *DNS Resolution Behavior* setting is not set to ignore local DNS.

The page will test each of the *DNS Servers* from the list at **System > General**.

The page will also test DNS servers from dynamic WANs if **DNS Server Override** is set and the firewall has obtained servers from dynamic sources.

Note: The *DNS Resolver mode* does not impact the behavior of this test. Even in resolver mode the individual DNS servers are tested as described above.

29.1.2 Results

The **Results** panel contains addresses returned by the DNS query along with the record type.

Underneath the results is a table containing the resolution **Timings** per server. This shows how fast each of the configured DNS servers responded to the specified query, or if they never responded.

The **More Information** panel contains links to ping and traceroute functions on the firewall for this host.

29.1.3 Aliases

The GUI can also create a *firewall alias* from the results of the DNS Lookup query.

Click  **Add alias** to create an alias containing the results of the query.

The name of the alias is the text entered for the DNS query but with `.` characters replaced by `_`. For example, a DNS lookup for `example.com` results in an alias named `example_com`.

If an alias already exists with that name, the button is labeled **Update Alias** instead. That version of the button will replace the contents of the existing alias with the current results of the DNS lookup.

29.2 Editing Files on the Firewall

Diagnostics > Edit File contains a file editor that allows editing and creating files on the filesystem of a device running pfSense® software.

> **Warning:** Be careful when choosing a file to edit! It is very easy to render the firewall unusable by editing the wrong file or introducing errors into the source code.
>
> Do not use this except under guidance of Netgate TAC or when there is sufficient knowledge to use it without causing unintended side effects.

29.2.1 Edit an Existing file

- Enter the full path of the filename to edit in **Save / Load from path** or click **Browse** and locate the file
- Click **Load**
- Edit the text
- Click **Save** to store the new content in the file

29.2.2 Create a new File

- Enter the new path and filename in **Save / Load from path**
- Enter the new contents of the file
- Click **Save** to create the new file with the given content

29.3 Command Prompt

The command prompt, available at **Diagnostics > Command Prompt**, executes shell commands, PHP code, and can download or upload whole files.

> **Warning:** Exercise caution using any of these utilities. Executing commands and PHP code improperly can render the firewall unusable. Use of this tool is not recommended except under the guidance of a Netgate TAC representative or if there is sufficient knowledge on the part of the user.

29.3.1 Execute Shell Commands

To execute a shell command:

- Navigate to **Diagnostics > Command Prompt**
- Enter the command into the **Command** box under **Execute Shell command**
- Click **Execute**

This page executes commands as if they were run from a console command line, and then it prints the results after the command terminates.

> **Warning: Commands must run and then stop or return**.
>
> If a command runs indefinitely, such as `ping` without a count or `tcpdump` without a limit set, it will never stop or return output. The command will run indefinitely in the background until it is manually killed.
>
> Interactive commands, such as `vi` will fail similarly, or may exit due to other issues with the terminal being non-interactive.

Previously used commands from this session can be recalled with the  and  buttons. The browser will forget the previous command list once it leaves the page.

29.3.2 Download

To download a file from the firewall filesystem:

- Navigate to **Diagnostics > Command Prompt**
- Enter the full path name in **File to download**
- Click  **Download**

29.3.3 Upload

To upload a file:

- Navigate to **Diagnostics > Command Prompt**
- Click **Browse**
- Locate and select the file on the local client computer
- Click  **Upload**

Note: Uploaded files are placed in `/tmp/` and can then be moved to alternate locations by other functions (such as the **Execute Shell Command** feature).

29.3.4 PHP Execute

This page can also execute PHP code.

- Navigate to **Diagnostics > Command Prompt**
- Type or paste PHP code into the **Execute PHP Commands** text area
- Click **Execute**

The GUI displays the output from the PHP code above the text area, or an error if it could not run the code.

29.4 Ping Host

The firewall can send ICMP echo requests, also known as "pings", to hosts over the network. These diagnostic packets test if the target host responds and measures latency between the firewall and target host. A basic ping test can be performed at the console, and a more detailed test is available in the GUI at **Diagnostics > Ping**.

29.4.1 Ping Options

When performing a ping test from the GUI, the following options are available:

Hostname
: A hostname or IP address to which the firewall will send ping requests.

IP Protocol
: The address type to ping when a hostname is entered that has both A (*IPv4*) and AAAA (*IPv6*) records.

Source Address
: The IP address on the firewall from which the ping request will be sent. This is especially important when testing LAN-to-LAN VPN connectivity. The default choice allows the operating system to automatically select the closest address to the target, based on the routing table.

Maximum Number of Pings
: The number of ping requests the firewall will send during this test. A higher count will take longer to complete and display results, especially if the target is down. Default is 3.

Seconds Between Pings
The number of seconds to wait between sending ping requests. Default is 1 second.

29.4.2 Ping from the GUI

To perform a ping test from the GUI:

- Navigate to **Diagnostics > Ping**
- Fill in the *Ping Options*

Note: At a minimum the **Hostname** is required.

- Click  **Ping** to start the test
- Wait for the GUI to display the test results

The GUI will display the results of the test automatically once complete. Do not navigate away from the page while the test is running.

29.4.3 Ping from the Console

A simple ping test may also be performed at the console menu, but without the additional options mentioned earlier. See *Ping host* for more information.

- Access the console menu locally or via SSH with an admin-level account (`admin`, `root`, or another privileged account using `sudo`).
- Enter the menu option which corresponds with **Ping Host** (e.g. `7`)
- Press `Enter`
- Enter the IP address or hostname to ping
- Press `Enter` to start the test
- Wait for the test to complete.

  The console outputs the test results in real time, and pauses afterward.
- Press `Enter` to return to the menu

29.5 Halting and Powering Off the Firewall

The firewall can be shut down safely by the **Halt** function available at **Diagnostics > Halt System** or from the console menu.

Warning: The best practice is to **never** cut power from a running system. Halting before removing power is always the safest procedure.

After the operating system halts, the device power will also be turned off if the hardware is capable.

29.5.1 Halt from the GUI

To halt the operating system from the GUI:

- Navigate to **Diagnostics > Halt System**
- Click  **Halt**
- Click **OK** to confirm the action and start the halt process

29.5.2 Halt from the Console

- Access the console menu locally or via SSH with an admin-level account (`admin`, `root`, or another privileged account using `sudo`).
- Enter the menu option which corresponds with **Halt system** (e.g. `6`)
- Press `Enter`
- Enter the `y` to confirm the action
- Press `Enter` to start the halt process

29.6 Rebooting the Firewall

pfSense® software can be rebooted safely and returned to an operational state using the page at **Diagnostics > Reboot System** or the console.

29.6.1 Reboot Methods

The following reboot methods are possible, but available options may be limited depending on the platform and installation options.

Reboot normally
: Performs a normal reboot in the traditional way. This method is always available.

Reroot
: Performs a "reroot" style reboot, which is faster than a traditional reboot but does not restart the entire operating system. All running processes are killed, all filesystems are remounted, and then the system startup sequence is run again. This type of restart is much faster as it does not reset the hardware, reload the kernel, or need to go through the hardware detection process.

Reboot into Single User Mode
: Restarts the firewall into single user mode for diagnostic purposes. The firewall cannot automatically recover from this state, console access is required to use single user mode and reboot the firewall.

  This option is not compatible with ARM-based systems.

> **Warning:** Using this option with ZFS will not automatically return to a normal boot and requires manual intervention at the console. See *Re-mount ZFS Volumes as Read/Write*.

> **Warning:** In single user mode the root filesystem defaults to read-only and other filesystems are not mounted. The firewall also does not have an active network connection. This option must only be used under the guidance of a Netgate TAC representative or a FreeBSD user with advanced knowledge.

Reboot and run a filesystem check
Reboots the firewall and forces a filesystem check using `fsck`, run five times. This operation can often correct issues with the filesystem on the firewall.

This option is not compatible with ARM-based systems or ZFS.

29.6.2 Reboot from the GUI

To reboot from the GUI:

- Navigate to **Diagnostics > Reboot System**
- Select the *Reboot Method*
- Click  **Submit** to reboot the system immediately

29.6.3 Reboot from the Console

To reboot from the console:

- Access the console menu locally or via SSH with an admin-level account (`admin`, `root`, or another privileged account using `sudo`).
- Enter the menu option which corresponds with **Reboot system** (e.g. `5`)
- Press `Enter`
- Enter the letter which corresponds with the desired *Reboot Method*
- Press `Enter`

Note: The single user mode and filesystem check options require an uppercase letter to be entered to confirm the action. This is necessary to avoid activating the options accidentally. The reboot and reroot options may be entered in upper or lower case.

29.7 Testing a TCP Port

The **Diagnostics > Test Port** page performs a simple TCP port connection test to check if the firewall can communicate with another host. This tests if a host is up and accepting connections on a given port, at least from the perspective of the firewall.

No data is transmitted to the remote host by this test. The test only attempts to open a connection and optionally displays the data sent back from the server.

In the default mode the test attempts a simple TCP handshake (SYN, SYN+ACK, ACK), and if the attempt succeeds, it reports the result.

Note: This test does not function for UDP since there is no way to reliably determine if a UDP port accepts connections in this manner.

To perform a test:

- Navigate to **Diagnostics > Test Port**
- Fill in the fields on the page. The **Hostname** and **Port** fields are required, the rest are optional.
- Click  **Test**.

The following options are available on this page:

Hostname
The IP address or hostname of the target system.

This is a required field.

Port
The TCP port on the target used by the test.

This is a required field and must be a valid port number, meaning an integer between `1` and `65535`.

Source Port
An optional specific source port for the query. This is unnecessary in most cases.

Remote Text
If checked, this option shows the text given by the server when connecting to the port. The server is given 10 seconds to respond, and this page will display all the text sent back by the server in those 10 seconds. As such, the test will run for a minimum of 10 seconds when performing this check.

Note: Not all daemons will output text to the user on connect, so this may be blank even if the service is working properly. For example, an SMTP server will respond with a welcome message, as will FTP, but an HTTP daemon will not send any text.

Source Address
A specific source IP address or IP Alias/CARP Virtual IP from which the query will be sent. The service being tested may require a specific source IP address, network, etc., in order to make a connection.

IP Protocol
This option selects either *IPv4* or *IPv6* to control which type of IP address is used when testing a hostname. If the connection is forced to IPv4 or IPv6 and the hostname does not contain a result using that protocol, the test will produce an error. For example if forced to IPv4 and given a hostname that only returns an IPv6 IP address (AAAA record), the test will fail.

29.7.1 Troubleshooting

```
nc: bind failed: Address already in use
```

The test produces this error if the **Source Port** field is set to a port currently in use by a local daemon on the firewall. Leave **Source Port** blank or pick another unused port.

See also:

To view a list of *ports currently in use*, visit **Diagnostics > Sockets**.

29.8 Traceroute

The traceroute page, located at **Diagnostics > Traceroute**, works like the *traceroute* command found on many platforms. It sends special packets which, as the name implies, trace a route across the network from this firewall to a remote host.

The results include a list of hops between hosts along with response times, as long as the intervening hosts respond to (or do not filter) traffic required for traceroute to work.

This document also includes a detailed explanation of how traceroute functions.

29.8.1 Traceroute Options

The GUI page to perform a traceroute contains the following options which control the behavior of the test:

Host
A hostname or IP address to which the firewall will trace the route.

IP Protocol
The IP address family the firewall will use when a hostname has both A (*IPv4*) and AAAA (*IPv6*) records.

Source Address
The IP address from which the firewall will send the trace.

This is especially important when testing LAN-to-LAN VPN connectivity.

Maximum number of hops
The maximum length of the path to trace.

The trace will stop if the path cannot be traced completely after this number of hops.

Reverse Address Lookup
When checked, traceroute will attempt to perform a PTR lookup to locate hostnames for hops along the path.

This option slows down the process as it has to wait for DNS replies.

Use ICMP
By default, traceroute uses UDP, but that may be blocked by some routers. Check this box to use ICMP instead, which may succeed.

The page will display output once the trace is complete. Press the **Stop** button at any time to see the current output of the trace if it is still running or stalled.

29.8.2 How Traceroute Works

Every IP packet contains a time-to-live (TTL) value. When a router passes a packet it decrements the TTL by one. When a router receives a packet with a TTL of 1 and the destination is not a locally attached network, the router returns an ICMP error message "Time-to-live exceeded" and drops the packet. This limits the impact of routing loops, which otherwise would cause each packet to loop indefinitely.

Traceroute uses this TTL to its advantage to map the path to a specific network destination. It starts by sending the first packet with a TTL of 1. The first router (usually the default gateway) will send back an ICMP time-to-live exceeded error. The program outputs the time between sending the packet and receiving the ICMP error along with the IP address that sent the error and its reverse DNS, if any. After sending three packets with a TTL of 1 and displaying their response times it increments the TTL to 2 and sends three more packets. It outputs the same types of information for the second hop. Traceroute increments the TTL and repeats the process until it reaches the specified destination or exceeds the maximum number of hops.

Traceroute functions slightly differently on Windows and Unix-like operating systems (BSD, Linux, macOS, Unix, etc.). Windows uses ICMP echo request packets (pings) while Unix-like systems use UDP packets by default. ICMP and UDP are layer 4 protocols, and traceroute is done at layer 3, so the protocol is largely irrelevant except when considering firewall rules and policy routing configurations. Traceroute from Windows clients will be policy routed based on rules which permit ICMP echo requests, while Unix-like clients will be policy routed by rules matching UDP.

In this example, `traceroute` is used to view the route to `www.google.com`:

```
# traceroute www.google.com
traceroute: Warning: www.google.com has multiple addresses; using 74.125.95.99
traceroute to www.l.google.com (74.125.95.99), 64 hops max, 40 byte packets
 1  core (172.17.23.1)  1.450 ms  1.901 ms  2.213 ms
 2  172.17.25.21 (172.17.25.21)  4.852 ms  3.698 ms  3.120 ms
 3  bb1-g4-0-2.ipltin.ameritech.net (151.164.42.156)  3.275 ms  3.210 ms  3.215 ms
 4  151.164.93.49 (151.164.93.49)  8.791 ms  8.593 ms  8.891 ms
 5  74.125.48.117 (74.125.48.117)  8.460 ms  39.941 ms  8.551 ms
 6  209.85.254.120 (209.85.254.120)  10.376 ms  8.904 ms  8.765 ms
 7  209.85.241.22 (209.85.241.22)  19.479 ms  20.058 ms  19.550 ms
 8  209.85.241.29 (209.85.241.29)  20.547 ms  19.761 ms
    209.85.241.27 (209.85.241.27)  20.131 ms
 9  209.85.240.49 (209.85.240.49)  30.184 ms
    72.14.239.189 (72.14.239.189)  21.337 ms  21.756 ms
10  iw-in-f99.google.com (74.125.95.99)  19.793 ms  19.665 ms  20.603 ms
```

The output shows that it took 10 hops to reach the destination and the latency generally increased with each hop, which is expected.

Note: When utilizing policy routing, such as with Multi-WAN, the firewall itself may not appear as a hop in traceroute. When policy routing is employed, PF does not decrement the TTL when forwarding packets, so traceroute cannot detect the firewall as an intermediate router.

29.9 Packet Capturing

29.9.1 Selecting the Proper Interface

To perform a packet capture, first determine the location from which to take the capture. A packet capture looks different depending upon the chosen interface and in certain scenarios it is better to capture on one specific interface, and in others, running multiple simultaneous captures on different interfaces is preferable.

Using `tcpdump` at the command line requires the "real" interface names that go with the friendly names shown in the firewall GUI. Visit **Interfaces > Assignments** and make a note of which OS interfaces (e.g. `igb1`), correspond with the friendly interfaces names on the firewall (e.g. `WAN`). *Real Interfaces vs. Friendly Names* lists common additional unassigned interface names that are present in many firewalls, depending on their configuration.

Table 1: Real Interfaces vs. Friendly Names

Real/Physical Name	Friendly Name
`enc0`, `ipsecX`	IPsec, encrypted traffic
`ovpncX`, `ovpnsX`	OpenVPN, encrypted traffic (Clients, Servers)
`pppoeX`, `poesX`	PPPoE WAN, PPPoE Server
`l2tpX`, `l2tpsX`	L2TP WAN, L2TP Server
`lo0`	Loopback Interface
`pfsync0`	pfsync interface – used internally
`pflog0`	PF logging – used internally

When selecting an interface, start with where the traffic flows into the firewall. For example, if a user is having trouble connecting to a port forward from outside the network, start with the WAN interface since that is where the traffic originates. If a client PC cannot reach the Internet, start with the LAN interface. When in doubt, try multiple interfaces and filter for the IP addresses or ports in question, keeping in mind when NAT will be applied.

29.9.2 Limiting capture volume

When capturing packets, limiting the volume of packets captured is important. Set limits on the capture so that it captures enough relevant traffic to troubleshoot the problem. If the limit is too low, the capture may be missing important details. If the limit is too high, there may be too much noise to sort through to find the problem.

Note: Capture files also consume disk space, which can be a factor on systems with smaller drives. Large captures will also take more time to download, which can be a concern on remote systems with slow WAN upload capacity.

When capturing without filtering on most networks, even for short time frame, huge amounts of data will end up in the capture to dig through when attempting to locate the problem. Display filters in Wireshark can limit which parts of an existing capture file are shown, but filtering appropriately at the time of capture is preferable to keep the capture file size down and to reduce processing time. Filters are discussed later in this chapter.

With an appropriate filter and packet count, capture files can be manageable and contain useful information.

29.9.3 Packet Capture GUI

The pfSense® software GUI offers an easy-to-use front end to `tcpdump` that performs packet captures which can then be viewed in the GUI or downloaded for deeper analysis using utilities such as Wireshark.

This feature is located at **Diagnostics > Packet Capture**.

See also:

If the options available in the GUI are too limiting, skip ahead to *Using tcpdump on the command line*.

Packet Capture Options

Interface
The network interface from which `tcpdump` will capture packets. Each assigned and unassigned interface on the firewall appears in the list, excluding special interfaces such as `pfsync0` and `pflog0`.

Packet Capture Filter
Selects a filter option or preset for the packet capture. Selecting a preset filter will hide the *Custom Filter Options* section.

Custom Filter
Enter custom values to limit captures.

Everything
Will capture any packet.

Only Untagged
Will exclusively capture any packets which are *not* VLAN-tagged.

Only Tagged
Will exclusively capture any packets which are VLAN-tagged.

Packet Count
Determines the total number of packets to capture before the capture stops automatically.

Captures may be "noisy" if they are not limited in some way. To get a usable result in these cases, increase this value beyond the default of `1000` to a much higher amount such as `10000`.

Packet Length
Sets the portion, in bytes, of each packet to capture.

In most cases the best practice is to capture the full packet (`0`), but for captures run over longer periods of time where the headers matter more than the payload of the packets, limiting this to `64` bytes or so will result in a much smaller capture file that may still have adequate data for troubleshooting purposes.

Promiscuous Mode
When checked, a capture includes all traffic arriving on the network interface for any destination MAC Address.

Without promiscuous mode the capture can only include traffic to/from the firewall itself as well as broadcast and multicast traffic.

> **Warning:** Some interface drivers and chipsets do not handle promiscuous mode well.

View Detail
Selects the amount of detail to display in the GUI when viewing a capture.

Note: View options (including `Name Lookup`) **do not** affect the packet capture file itself, and may be changed before or after the packet capture. Change the value and click  **View** to display the capture with the new view option.

View Type
Force the captured traffic to be interpreted as the specified type when viewed. This is particularly helpful when viewing CARP traffic.

Name Lookup
Causes `tcpdump` to perform a name lookup for the port and host address, including the MAC OUI.

Warning: Avoid using this option when possible as it will delay the output due to the extra time taken by reverse DNS lookups. Also, it is typically easier to troubleshoot when viewing IP addresses instead of hostnames, and reverse DNS can be inaccurate.

Custom Filter Options

Filter Sections

Filter options are separated into two sections, untagged and tagged. Each section may be individually included or excluded in the packet capture. This is useful in the case where both untagged and some other tagged traffic need to be captured. The available filter sections are:

Untagged Filter
Values entered in this section will only apply to packets which *do not* have a VLAN tag.

Tagged Filter
Values entered in this section will only apply to packets which have a VLAN tag set.

Multiple Values and Filter Operators

All input fields accept multiple **space-separated** values. When multiple values are specified, the drop-down menu for each option has an operator which defines the behavior:

all of
Captures only packets that match **all** of the values specified in this single filter option.

any of
Captures packets that match **any** of the values specified in this single filter option.

none of
Captures packets that **do not** match the values specified in this single filter option.

OR <type>
Places a logocal "or" between multiple separate filter options (e.g. a given host IP address **or** a given MAC address)

Filter Options

The following fields are available for filtering in the Untagged and/or Tagged filter option groups:

Include/Exclude
Defines how the GUI handles packets and options in a section (untagged, tagged):

Include any of
Enables the filter options for this section and includes packets of this type provided they match the defined criteria.

Exclude all
Disables the filter options for this section and excludes all packets of this type.

VLAN Tag
Match traffic that is tagged with the specified VLAN tags. This is particularly useful when capturing on a trunk interface and not all VLANs need to be captured.

VLAN Tag Level
Which VLAN tag stack level to apply the filter to. This is useful when filtering QinQ traffic.

Host Address or Subnet
Filters traffic going to or from specific IP addresses (`x.x.x.x`) or CIDR-masked subnets (`x.x.x.x/yy`).

Host MAC Address
Filters traffic going to or from specific MAC addresses.

Enter MAC addresses in colon-separated format, such as `xx:xx:xx:xx:xx:xx`. To match a partial address, use one (`xx`), two (`xx:xx`), or four (`xx:xx:xx:xx`) segments in the same colon-separated format.

Protocol
Limits the capture to packets for specific protocols. The drop-down option in the GUI includes common protocols such as TCP, UDP, ICMP, ICMP6, CARP, but any protocols may be specified manually by number or name.

Port Number
Limits the capture to packets with the specified source or destination port. Only effective on protocols which have ports (TCP, UDP).

Ethertype
May be used to limit the capture to only IPv4, IPv6, or ARP traffic. This is useful when not filtering by IP address (e.g. by port number or MAC address).

Performing a Packet Capture

To make a packet capture in the GUI:

- Navigate to **Diagnostics > Packet Capture**
- Configure the options on the page as described in *Packet Capture Options*
- Click  **Start** to begin capturing packets.

  The page will display a live preview of the captured packets which refreshes every few seconds. The resulting `tcpdump` command is also shown, along with the message "Running packet capture" indicating the capture is in process.

- Click  **Stop** to manually end the capture and view the output.

  If the capture has a maximum packet **Count** set it will stop automatically when it reaches that count. In this case the capture does not need a manual stop action unless it must be stopped before reaching that count.

Viewing the Captured Data

The capture output can be viewed in the GUI or downloaded for later viewing in a program such as Wireshark.

- Navigate to **Diagnostics > Packet Capture**
- Set the **View Options** (*Packet Capture Options*) to control how the GUI displays the contents of the capture.
- Click  **View**

  The page displays the output in a field titled **Packet Capture Output** in standard `tcpdump` format.

  Note: If the **View** button is not visible on the page, there is no existing capture data to view. Perform a new capture first.

- Click  **Download** to download this file for later viewing (Optional).

  See also:

  For more detail on using Wireshark to view a capture file, see *Viewing a Packet Capture File*.

29.9.4 Using tcpdump on the command line

The `tcpdump` program is a command line packet capture utility provided with most UNIX and UNIX-like operating system distributions, including FreeBSD. It is included in pfSense® software and is usable from a shell on the console or over SSH.

The `tcpdump` program is an exceptionally powerful tool, but that also makes it daunting to the uninitiated user. The `tcpdump` binary in FreeBSD has over 50 different command line flags, limitless possibilities with filter expressions, and its man page, providing only a brief overview of all its options, is nearly 1200 lines long and 67k.

After learning to use `tcpdump`, knowledge of how to interpret the data it provides is also necessary, which can require an in-depth understanding of networking protocols.

This section is intended to provide an introduction to this topic and leave the reader with enough knowledge for basic troubleshooting. A comprehensive review of packet capturing and interpretation of the results is outside the scope of this documentation.

See also:

For those with a thirst for more than basic knowledge in this area, see *Additional References* for more resources.

Common tcpdump flags

The following table shows the most common command line flags for `tcpdump`. This section contains information on each of these flags.

Table 2: Commonly Used tcpdump Flags

Flag	Description
`-i <interface>`	Listen on `<interface>`, .e.g. `-i igb0`
`-n`	Do not perform reverse DNS resolution on IP addresses
`-w <filename>`	Save capture in pcap format to `<filename>`, e.g. `-w /tmp/wan.pcap`
`-s <bytes>`	Snap length: Amount of data to be captured from each frame
`-c <packets>`	Exit after receiving a specific number of packets
`-p`	Do not put the interface in promiscuous mode
`-v`	Verbose output
`-e`	Print link-layer header on each line

-i flag

The `-i` flag specifies the interface on which `tcpdump` will listen. Use FreeBSD interface names here, such as `igb0`, `em0`, `vmx0`, etc.

-n flag

Do not resolve IP addresses using reverse DNS. When this option is *not* specified, `tcpdump` will perform a reverse DNS (PTR) lookup for each IP address. This generates a significant amount of DNS traffic in captures displaying large volumes of traffic. Disable this to avoid adding load to DNS servers.

The best practice is to always use `-n` because it eliminates the delay caused by performing the reverse lookup between when `tcpdump` captures a packet and when it can display the content. Also, IP addresses are typically easier to read and understand than their PTR records. That is a matter of personal preference, though, and in familiar environments where the PTR records are known to provide the actual host names of the devices, captures may be run without `-n` to show the hostnames.

Another reason to use `-n`, is to be "sneaky." One means of detecting packet capturing is looking for spikes and patterns in DNS PTR lookups. Skipping the DNS lookup will not cause any extra traffic to be generated in the process.

-w flag

`tcpdump` can save capture files in pcap format for later analysis or analysis on another system. This is commonly done from command line only devices like those running pfSense software so the file can be copied to a host running Wireshark or another graphical network protocol analyzer and reviewed there. When saving to a file using `-w`, the frames will not be displayed in the terminal as they otherwise are.

See also:

See *Using Wireshark* for more information about using Wireshark with pfSense software.

-s flag

By default `tcpdump` only saves the first `64` bytes of each frame when capturing to a file. This is enough to contain the IP and protocol header for most protocols, but limits the usability of capture files. By using the `-s` flag, `tcpdump` can be told how much of the frame to capture, in bytes. This is called the snap length.

Table 3: Example Uses of tcpdump -s

Flag	Description
`-s 500`	Capture the first 500 bytes of each frame
`-s 0`	Capture each frame in its entirety

In most cases, using `-s 0` is the best practice when capturing to a file for analysis on another system. The only exception to this is scenarios where a significant amount of traffic must be captured over a longer period of time. If the information being sought is known to be in the header, the default `64` bytes of each frame may be used to get the required information while significantly reducing the size of the resulting capture file.

-c flag

To capture a certain number of frames and then exit, use the `-c` flag. Example usage: `tcpdump` will exit after capturing 100 frames by specifying `-c 100`.

-p flag

Normally when capturing traffic with `tcpdump`, it puts the network interface into promiscuous mode. When not running in promiscuous mode, the interface only receives frames destined for its own MAC address as well as broadcast and multicast addresses. When switched into promiscuous mode, the interface shows every frame on the wire that arrives at the network interface. In a switched network, this generally has little impact on the capture. In networks where the device is connected to a vswitch also in promiscuous mode, or a hub, using `-p` can significantly limit noise in the capture when the only traffic of interest is to and from the system performing the capture.

-v flag

The `-v` flag controls the detail, or verbosity, of the output. Using more `v` options yields more detail, so use `-v`, `-vv`, or `-vvv` to view even more detail in the output printed to the console. This option does not affect the detail stored in a capture file when using the `-w` switch, but will instead cause the process to report the number of packets captured every `10` seconds.

-e flag

Normally `tcpdump` does not show any link layer information. Specify `-e` to display the source and destination MAC addresses, and VLAN tag information for any traffic tagged with 802.1q VLANs.

Example capture without -e

This capture shows the default output, containing no link layer information:

```
# tcpdump -ni igb1 -c 5
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes
23:18:15.830706 IP 10.0.64.210.22 > 10.0.64.15.1395: P 2023587125:2023587241(116)
   ack 2091089207 win 65535
23:18:15.830851 IP 10.0.64.210.22 > 10.0.64.15.1395: P 116:232(116) ack 1 win 65535
23:18:15.831256 IP 10.0.64.15.1395 > 10.0.64.210.22: . ack 116 win 65299
23:18:15.839834 IP 10.0.64.3 > 224.0.0.18: VRRPv2, Advertisement, vrid 4, prio 0,
   authtype none, intvl 1s, length 36
23:18:16.006407 IP 10.0.64.15.1395 > 10.0.64.210.22: . ack 232 win 65183
5 packets captured
```

Example capture using -e

Here the link layer information is included by using `-e`. Note the source and destination MAC addresses in addition to the source and destination IP addresses:

```
# tcpdump -ni igb1 -e -c 5
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes
23:30:05.914958 00:0c:29:0b:c3:ed > 00:13:d4:f7:73:d2, ethertype IPv4 (0x0800), length␣
↪170:
   10.0.64.210.22 > 10.0.64.15.1395: P 2023592509:2023592625(116) ack 2091091355 win␣
↪65535
23:30:05.915110 00:0c:29:0b:c3:ed > 00:13:d4:f7:73:d2, ethertype IPv4 (0x0800), length␣
↪170:
   10.0.64.210.22 > 10.0.64.15.1395: P 116:232(116) ack 1 win 65535
23:30:05.915396 00:13:d4:f7:73:d2 > 00:0c:29:0b:c3:ed, ethertype IPv4 (0x0800), length␣
↪60:
   10.0.64.15.1395 > 10.0.64.210.22: . ack 116 win 65299
23:30:05.973359 00:00:5e:00:01:04 > 01:00:5e:00:00:12, ethertype IPv4 (0x0800), length␣
↪70:
   10.0.64.3 > 224.0.0.18: VRRPv2, Advertisement, vrid 4, prio 0, authtype none, intvl␣
↪1s,
   length 36
23:30:06.065200 00:13:d4:f7:73:d2 > 00:0c:29:0b:c3:ed, ethertype IPv4 (0x0800), length␣
↪60:
   10.0.64.15.1395 > 10.0.64.210.22: . ack 232 win 65183
5 packets captured
```

tcpdump Filters

Running `tcpdump` without filters on most firewalls will produce so much output that it is extremely difficult to find traffic of interest. There are numerous filtering expressions available that limit traffic display and capture.

Host filters

To filter for a specific host, append `host` and the IP address to the `tcpdump` command.

To filter for host `192.168.1.100` use the following command:

```
# tcpdump -ni igb1 host 192.168.1.100
```

The previous command captures all traffic to and from the given host. To only capture traffic being *initiated* by that host, use the `src` directive:

```
# tcpdump -ni igb1 src host 192.168.1.100
```

Similarly, filtering for traffic *destined* to that IP address is possible by specifying `dst`:

```
# tcpdump -ni igb1 dst host 192.168.1.100
```

Network filters

Network filters narrow the capture to a specific subnet using the `net` expression. Following `net`, specify a CIDR-masked network (`192.168.1.0/24`), dotted quad (`192.168.1.1`), dotted triple (`192.168.1`), dotted pair (`192.168`) or simply a number (`192`). A dotted quad is equivalent to specifying `host`, a dotted triple uses a subnet mask of `255.255.255.0`, a dotted pair uses `255.255.0.0`, and a number alone uses `255.0.0.0`.

The best practice for filtering by network is to use a CIDR masked subnet prefix specification as an argument to `net`:

```
# tcpdump -ni igb1 src net 172.16.0.0/12
```

Alternately, omit parts of an address to use the assumed masks mentioned previously.

The following command displays traffic to or from any host with a `192.168.1.x` IP address:

```
# tcpdump -ni igb1 net 192.168.1
```

The next command will capture traffic to or from any host with a `10.x.x.x` IP address:

```
# tcpdump -ni igb1 net 10
```

Those examples will capture all traffic to or from the specified network. The `src` or `dst` keywords may be used the same as with `host` filters to capture only traffic initiated by or destined to the specified network:

```
# tcpdump -ni igb1 src net 10
```

Protocol and port filters

Narrowing down by host or network can be inadequate to eliminate unnecessary traffic from a capture. Or the source or destination of traffic may not be significant, and all traffic of a certain type should be captured. In other cases, filtering out all traffic of a specific type can reduce noise.

TCP and UDP port filters

To filter on TCP and UDP ports, use the `port` directive. This captures both TCP and UDP traffic using the specified port either as a source or destination port. It can be combined with `tcp` or `udp` to specify the protocol, and `src` or `dst` to specify a source or destination port.

Capture all HTTP traffic

```
# tcpdump -ni igb1 tcp port 80
```

Capture all DNS traffic

Capture all DNS traffic (Queries can use both UDP and TCP):

```
# tcpdump -ni igb1 port 53
```

Protocol filters

Specific protocols can be filtered using the `proto` directive or by using the protocol name directly. Parameters passed to the `proto` directive can be specified using the IP protocol number or one of the names `icmp`, `igmp`, `igrp`, `pim`, `ah`, `esp`, `carp`, `vrrp`, `udp`, or `tcp`. Because the normal protocol names are reserved words, they must be escaped with one or two backslashes when used with the `proto` directive, depending on the shell. The default shell available in pfSense software requires two backslashes to escape these protocol names. If the command returns a syntax error, check that the protocol name is properly escaped.

The following capture will show all ICMP traffic on the `igb1` interface:

```
# tcpdump -ni igb1 proto \\icmp
```

Specifying `carp` for the protocol will capture CARP traffic but it also needs `-T carp` in order to interpret the CARP packets correctly when viewing the output using `tcpdump`. The GUI makes this adjustment automatically when capturing CARP.

The following capture will show all CARP traffic on the `igb1` interface, which can be useful to ensure CARP traffic is being sent and received on the specified interface. It also omits the `proto` keyword, showing that it works on its own:

```
# tcpdump -i igb1 -T carp carp
```

Negating a filter match

In addition to matching specific parameters, a filter match can be negated by specifying `not` in front of the filter expression. When troubleshooting something other than CARP, and its multicast heartbeats are cluttering the capture output, exclude it as follows:

```
# tcpdump -ni igb1 not proto \\carp
```

Combining filters

Any of the aforementioned filters can be combined using `and` or `or`. The following sections provide some examples.

Display all HTTP traffic to and from a host

Display all HTTP traffic to or from `192.168.1.11`:

```
# tcpdump -ni igb1 host 192.168.1.11 and tcp port 80
```

Display all HTTP traffic to and from multiple hosts

Display all HTTP traffic from either `192.168.1.11` or `192.168.1.15`:

```
# tcpdump -ni igb1 host 192.168.1.11 or host 192.168.1.15 and tcp port 80
```

Filter expression usage

Filter expressions must come after every command line flag used. Adding any flags after a filter expression will result in a syntax error.

Incorrect ordering

```
# tcpdump -ni igb1 -T carp carp -c 2
tcpdump: syntax error
```

Correct ordering

```
# tcpdump -ni igb1 -T carp -c 2 carp
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 65535 bytes
14:50:07.426993 IP 198.51.100.12 > 224.0.0.18: CARPv2-advertise 36: vhid=11 advbase=1
   advskew=0 authlen=7 counter=5449924379588860810
14:50:08.436849 IP 198.51.100.12 > 224.0.0.18: CARPv2-advertise 36: vhid=11 advbase=1
   advskew=0 authlen=7 counter=5449924379588860810
2 packets captured
78 packets received by filter
0 packets dropped by kernel
```

More on Filters

This section covered the most commonly used `tcpdump` filter expressions, and probably covers all the syntax most users will need. However this barely scratches the surface of the possibilities. There are many documents on the web that cover `tcpdump` in general and filtering specifically. See *Additional References* at the end of this chapter for links to more resources.

Practical Troubleshooting Examples

This section details best practice approaches for troubleshooting a few specific problems. There are multiple ways to approach any problem, but packet capturing can rarely be beat for its effectiveness. Examining the traffic on the wire provides a level of visibility into what is actually happening on the network

Port forward not working

In this example, a new port forward is failing to respond to a request from a host on the Internet. The troubleshooting steps outlined in *Troubleshooting NAT Port Forwards* offers one way to approach this, but sometimes packet capturing is the only or easiest way to find the source of the problem.

Start from WAN

First, make sure the traffic is getting to the WAN interface. Start a `tcpdump` session on the WAN interface, and watch for the traffic:

```
# tcpdump -ni igb1 tcp port 5900
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes
11:14:02.444006 IP 172.17.11.9.37219 > 10.0.73.5.5900: S 3863112259:3863112259(0)
   win 65535 <mss 1260,nop,nop,sackOK>
```

In this case, a packet comes in from the WAN, so it is making it that far. Note that the first part of the TCP handshake, a packet with only SYN set (the `S` shown), is reaching the firewall. If the port forward was working, a SYN ACK (`S.`) packet would be shown in reply to the SYN. With no return traffic visible, it could be a firewall rule or the target system may be unreachable – turned off, not listening on the specified port, host firewall blocking the traffic, etc.

Check Internal Interface

The next step would be to run a `tcpdump` session on the internal interface associated with the port forward:

```
# tcpdump -ni igb0 tcp port 5900
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb0, link-type EN10MB (Ethernet), capture size 96 bytes
11:14:38.339926 IP 172.17.11.9.2302 > 192.168.30.5.5900: S 1481321921:1481321921(0)
   win 65535 <mss 1260,nop,nop,sackOK>
```

Looking at the internal traffic, the connection left the inside interface and the local IP address was translated correctly. If this local address matches what was expected, then both the port forward and the firewall rule are working properly, and connectivity to the local PC must be confirmed by other means. If no output was displayed, then there is a problem with the firewall rule or the port forward may have been incorrectly defined. For this example, the target system was unplugged.

IPsec tunnel will not connect

tcpdump has some awareness of the protocols being used, which can be very helpful in figuring out problems with IPsec tunnels. The next few examples will show how certain error conditions may present themselves when monitoring with tcpdump. The IPsec logs are usually more helpful, but this can confirm what is actually being seen by the firewall. For encrypted traffic such as IPsec, packet capturing of the traffic is of less value as the payload of the captured packets cannot be examined without additional parameters, but it is helpful to determine if traffic from the remote end is reaching the firewall and which phases complete.

This first tunnel has an unreachable peer:

```
# tcpdump -ni igb1 host 192.168.10.6
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes

19:11:11.542976 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
19:11:21.544644 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
```

This tunnel attempt has a mismatched PSK, notice how it attempts to move to phase 2, but then stops:

```
# tcpdump -ni igb1 host 192.168.10.6
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes
19:15:05.566352 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
19:15:05.623288 IP 192.168.10.6.500 > 192.168.10.5.500: isakmp: phase 1 R agg
19:15:05.653504 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 2/others I inf[E]
```

Now Phase 1 is OK but there is a mismatch in the Phase 2 information. It will repeatedly attempt phase 2 traffic but there will not be any traffic in the tunnel:

```
# tcpdump -ni igb1 host 192.168.10.6
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes
19:17:18.447952 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
19:17:18.490278 IP 192.168.10.6.500 > 192.168.10.5.500: isakmp: phase 1 R agg
19:17:18.520149 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
19:17:18.520761 IP 192.168.10.6.500 > 192.168.10.5.500: isakmp: phase 2/others R inf[E]
19:17:18.525474 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 2/others I inf[E]
19:17:19.527962 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 2/others I oakley-
↪quick[E]
```

Finally, a fully working tunnel with two-way traffic after Phase 1 and Phase 2 have completed!:

```
#  tcpdump -ni igb1 host 192.168.10.6
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on igb1, link-type EN10MB (Ethernet), capture size 96 bytes
21:50:11.238263 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
21:50:11.713364 IP 192.168.10.6.500 > 192.168.10.5.500: isakmp: phase 1 R agg
21:50:11.799162 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 1 I agg
21:50:11.801706 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 2/others I inf[E]
21:50:11.812809 IP 192.168.10.6.500 > 192.168.10.5.500: isakmp: phase 2/others R inf[E]
21:50:12.820191 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 2/others I oakley-
↪quick[E]
21:50:12.836478 IP 192.168.10.6.500 > 192.168.10.5.500: isakmp: phase 2/others R oakley-
```

(continues on next page)

(continued from previous page)

```
↪quick[E]
21:50:12.838499 IP 192.168.10.5.500 > 192.168.10.6.500: isakmp: phase 2/others I oakley-
↪quick[E]
21:50:13.168425 IP 192.168.10.5 > 192.168.10.6: ESP(spi=0x09bf945f,seq=0x1), length 132
21:50:13.171227 IP 192.168.10.6 > 192.168.10.5: ESP(spi=0x0a6f9257,seq=0x1), length 132
21:50:14.178820 IP 192.168.10.5 > 192.168.10.6: ESP(spi=0x09bf945f,seq=0x2), length 132
21:50:14.181210 IP 192.168.10.6 > 192.168.10.5: ESP(spi=0x0a6f9257,seq=0x2), length 132
21:50:15.189349 IP 192.168.10.5 > 192.168.10.6: ESP(spi=0x09bf945f,seq=0x3), length 132
21:50:15.191756 IP 192.168.10.6 > 192.168.10.5: ESP(spi=0x0a6f9257,seq=0x3), length 132
```

Traffic traversing an IPsec tunnel

Traffic can also be observed traversing IPsec tunnels by capturing on the `enc0` interface. This can help determine if traffic is attempting to reach the far end by using the tunnel. All traffic for all IPsec tunnels appears on the `enc0` interface.

In the following example, a host on one side of the tunnel is successfully sending an ICMP echo request (ping) to the far side, and receiving replies:

```
# tcpdump -ni enc0
tcpdump: WARNING: enc0: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on enc0, link-type ENC (OpenBSD encapsulated IP), capture size 65535 bytes
15:52:46.151098 (authentic,confidential): SPI 0xcd77e085: IP 10.3.0.1 > 10.7.0.1:
   ICMP echo request, id 44640, seq 0, length 64
15:52:46.151814 (authentic,confidential): SPI 0xc0afb14d: IP 10.7.0.1 > 10.3.0.1:
   ICMP echo reply, id 44640, seq 0, length 64
15:52:47.154243 (authentic,confidential): SPI 0xcd77e085: IP 10.3.0.1 > 10.7.0.1:
   ICMP echo request, id 44640, seq 1, length 64
15:52:47.154843 (authentic,confidential): SPI 0xc0afb14d: IP 10.7.0.1 > 10.3.0.1:
   ICMP echo reply, id 44640, seq 1, length 64
```

If traffic was not properly entering the tunnel, no output would be shown. If there is a firewall or internal routing issue on the far side, traffic will appear leaving but nothing will show returning.

Troubleshooting Outbound NAT

For complex environments where Manual Outbound NAT is needed, `tcpdump` can be of great assistance in troubleshooting the Outbound NAT configuration. One good capture to use is to look for traffic with private IP addresses on the WAN interface, as everything on WAN should be have NAT applied and appear to be a public IP address. The following capture will display any traffic with RFC 1918 IP addresses as the source or destination. This will show any traffic that is not matching one of the outbound NAT rules, providing information to help review the Outbound NAT configuration to find the problem:

```
# tcpdump -ni igb1 net 10 or net 192.168 or net 172.16.0.0/12
```

29.9.5 Using Wireshark

Wireshark is a GUI protocol analysis and packet capture tool that can view and capture traffic much like `tcpdump`. Wireshark is Open Source software, freely available at http://www.wireshark.org/. Wireshark can analyze capture files generated by the pfSense® software GUI, `tcpdump`, Wireshark, or any other software that writes files in the standard `pcap` file format.

Before proceeding, download and install Wireshark onto a client computer.

Viewing a Packet Capture File

To view a capture file in Wireshark, use one of the following methods:

Manually Open File
The basic way to open a file manually is:

- Start Wireshark
- Navigate to **File > Open**
- Locate the capture file and click it
- Click the **Open** button

Double Click
A file with a `.pcap` extension can be opened by double clicking on it in Windows, macOS, and many Linux distributions. This action is typically performed in a file manager such as File Explorer, Finder, Nemo, Dolphin, or similar programs.

Download and Open
Browsers may often to open a downloaded capture file directly in Wireshark. This may be an option on a file download prompt, or an option from the list of downloaded files.

Once the file is open Wireshark displays a screen similar to Figure *Wireshark Capture View* which contains data from the capture file.

This view in Wireshark has a list summarizing the packets in the capture file in the top pane, with one packet per line. If there are too many packets, the results can be filtered using the **Filter** box on the toolbar.

Select a packet by clicking it in the list and the lower frames show the details of what is contained within the packet payload. The first lower pane shows a break-down of the packet structure, and each of these items can be expanded for more detail. If the packet is part of a protocol known to Wireshark, in some cases it can interpret the data and show even more details. The bottom pane shows a hexadecimal and ASCII representation of the data contained in the packet.

Viewing the capture this way makes it easy to see the flow of traffic with as much or as little detail as needed.

Wireshark Analysis Tools

While some problems will require considerable knowledge of how the underlying protocols function, the analysis tools built into Wireshark helps lessen that need for many protocols. The **Analyze** and **Statistics** menus have a few options that automate some of the analysis and provide summarized views of capture content. The **Expert Info** options under the **Analyze** menu show a list of errors, warnings, notes and network conversations contained in the capture.

Wireshark may note errors for incorrect checksums. This is because most network interfaces handle checksums in hardware directly before putting it on the wire. This is the only exception to the earlier note saying what is shown in a packet capture is what is on the wire. Traffic sent from the system where the capture is taken will have incorrect checksums where they are performed in hardware, though traffic coming in from a remote system should always have correct checksums. Checksum offloading can be turned off to ensure traffic is shown exactly as the host is putting it on

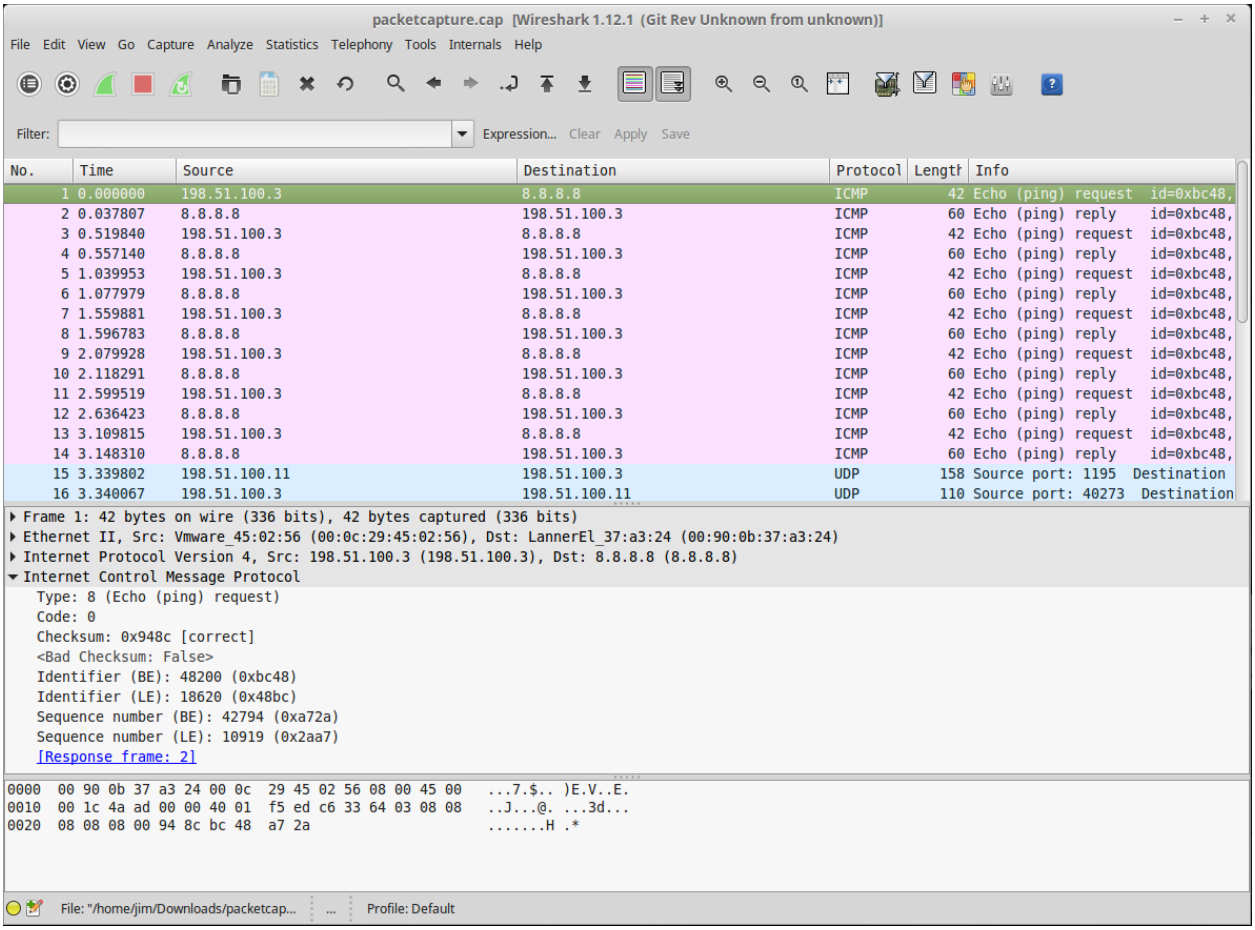


Fig. 1: Wireshark Capture View

the wire, though usually this is safe to ignore. To verify checksums, capture traffic from another system using a network tap or switch span port.

Tip: Span ports can also be setup on bridges in pfSense software, see *Span Port* for more information.

The **Telephony** menu is one example of automated analysis Wireshark can perform. These functions make it easy to diagnose VoIP problems. In the example shown in Figure *Wireshark RTP Analysis*, VoIP traffic was traversing an MPLS WAN circuit with the provider's routers attached to an OPT interface of pfSense software on both sides. A capture from the OPT interface on the initiating end showed no loss, indicating the traffic was being sent to the provider router, but the OPT interface on the opposite end showed considerable packet loss in one direction when multiple simultaneous calls were active. These packet captures helped convince the provider of a problem on their network, and they found and fixed a QoS configuration problem on their side.

Src IP addr	Src port	Dest IP addr	Dest port	SSRC	Payload	Packets	Lost	Max Delta (ms)	Max Jitter (ms)	Mean Jitter (ms)	Pb?
10	13114	192	2244	0x63B69143	ITU-T G.711 PCMU	2103	0 (0.0%)	20.07	0.12	0.01	X
192.1	2244	1	13114	0x6F1D173B	ITU-T G.711 PCMU	1646	477 (22.5%)	179.89	51.24	1.84	X
10	11224	192	2268	0x2247C8D8	ITU-T G.711 PCMU	1321	0 (0.0%)	99.99	5.03	0.07	X
192.1	2268	1	11224	0x6C5B26A1	ITU-T G.711 PCMU	879	460 (34.4%)	340.79	49.96	2.67	X
10	17924	192	2242	0x393CBA89	ITU-T G.711 PCMU	480	0 (0.0%)	20.04	0.15	0.01	X
192.1	2242	1	17924	0x6177246E	ITU-T G.711 PCMU	133	366 (73.3%)	339.79	71.38	9.17	X

Fig. 2: Wireshark RTP Analysis

When viewing a packet capture containing RTP traffic, click **Telephony > RTP > Show all streams** to see this screen.

Remote Real-time Capture

From a UNIX host that has Wireshark available, real-time remote capture is possible by redirecting the output from an SSH session. This has been tested and known to work on FreeBSD and Ubuntu-based Linux distributions.

In order to use this technique, SSH must be enabled on pfSense software and an SSH key is required (see *Secure Shell (SSH)*). The key must first be loaded into `ssh-agent` or generated without a passphrase because the redirection will not allow the user to enter a password or passphrase.

Warning: Using `ssh-agent` is the best practice as any key without a passphrase is highly insecure.

Before attempting this technique, check that the user can connect to the firewall running pfSense software using an SSH key without needing to type the passphrase. The first time the user connects, they are prompted to save the host key, so that must also be done before trying to start Wireshark. `ssh-agent` may also be started from a terminal window or shell like so:

```
# eval ssh-agent
Agent pid 29047
# ssh-add
Enter passphrase for /home/jimp/.ssh/id_rsa:
Identity added: /home/jimp/.ssh/id_rsa (/home/jimp/.ssh/id_rsa)
```

Then start an SSH session as usual:

```
# ssh root@192.168.1.1
The authenticity of host '192.168.1.1 (192.168.1.1)' can't be established.
DSA key fingerprint is 9e:c0:b0:5a:b9:9b:f4:ec:7f:1d:8a:2d:4a:49:01:1b.
```

(continues on next page)

(continued from previous page)

```
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.1.1' (DSA) to the list of known hosts.

*** Welcome to pfSense ***
[...]
```

After confirming the SSH connection works, start the remote capture as follows:

```
# wireshark -k -i <(ssh root@192.168.1.1 tcpdump -i igb1 -U -w - not tcp port 22)
```

Replace `192.168.1.1` with the IP address of the firewall running pfSense software. The `not tcp port 22` filter excludes traffic from the SSH session, which will otherwise clog the capture output. The above is written in BASH style syntax, but may work with other shells. Adjust the `tcpdump` arguments for the interface, and add additional expressions. The `-U` and `-w -` are necessary so that it writes the output to `stdout`, and writes each packet as it arrives.

See also:

The Capture Setup/Pipes page on the Wireshark wiki contains other related techniques.

29.9.6 Additional References

This chapter only scratches the surface of the possibilities with packet captures. Packet capturing is a powerful means of troubleshooting network connectivity issues, and troubleshooting skills are greatly improved when the possibilities are learned in more depth. The following links are related resources with deeper knowledge beyond the scope of this documentation.

- Computer Networking: Internet Protocols in Action by Jeanna Matthews
- Tcpdump Filters by Jamie French
- Tcpdump Advanced Filters by Sebastien Wains
- Tcpdump Filters by Marios Iliofotou
- FreeBSD Man Page for tcpdump

Capturing packets is the most effective means of troubleshooting problems with network connectivity. Packet capturing, also known as "sniffing", shows packets "on the wire" coming in and going out of an interface. Observing how traffic is sent and received by the firewall is a great help in narrowing down problems with firewall rules, NAT entries, and other networking issues. pfSense® software includes a GUI page that captures packets using `tcpdump` in an easy manner. The `tcpdump` utility can also be used at the command line in a shell, and alternately, packet capturing can be performed by Wireshark.

29.9.7 Capture frame of reference

Keep in mind that packet captures show exactly what is on the wire. A packet capture is the first process to see traffic when an interface receives a packet and it is the last to see traffic when an interface sends a packet as it flows through the firewall. It sees traffic before firewall rule, NAT rule, and all other processing on the firewall happens for traffic coming into that interface, and after all that processing occurs for traffic leaving that interface. For incoming traffic, captures shows traffic that arrives on an interface on the firewall regardless of whether that traffic will be blocked by the firewall configuration. Figure *Stack Processing Order* illustrates where `tcpdump` packet captures tie into the processing order.

See also:

Ordering of NAT and Firewall Processing explains how the firewall processes connections in greater detail.

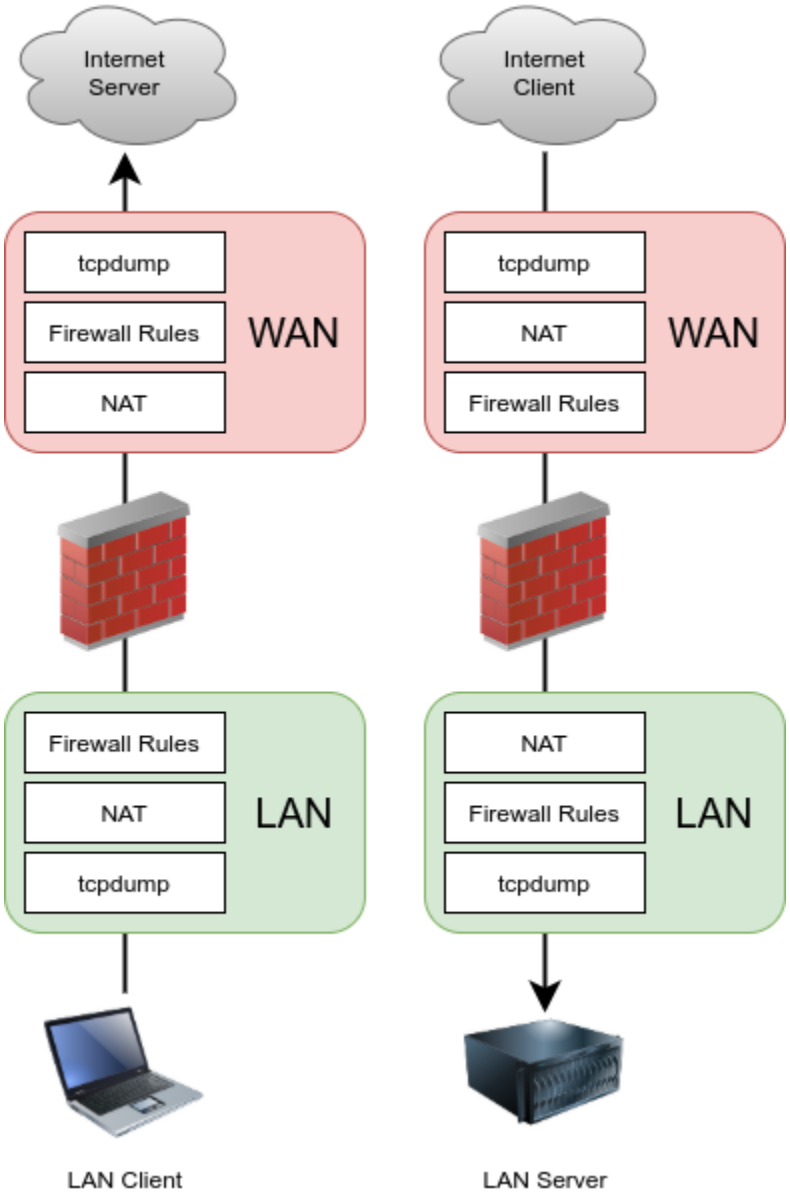


Fig. 3: Stack Processing Order

CHAPTER

THIRTY

PACKAGES

30.1 Package Manager

Packages are managed at **System > Packages** (Figure *Package Listing*). The listings there, presented in alphabetical order, show all the information about a package.

See also:

pfSense® software uses similar mechanisms to handle base system upgrades. For more information on that process, see *Upgrade Guide*.

Name
The name of the package. This is a unique, and typically short, name used to identify the package. On some packages, the name is a link to more information about the package.

Version
The version number of the package. This number is specific to the package on pfSense, and is not necessarily related to the version of the underlying software (if there is any). The version number is also a link to recent changes for the package.

Description
Longer text describing the package, its purpose, and so on. If the package depends upon other packages, the GUI lists them here denoted by  .

Warning: For security reasons, keep the installed packages to the bare minimum required for a deployment.

See also:

- *Troubleshooting Upgrades* (Packages and Updates use the same backend)
- *Troubleshooting DNS Resolution Issues*
- *Troubleshooting Routes*
- *Troubleshooting a Broken pkg Database*



Fig. 1: Package Listing

30.1.1 Installing Packages

Packages are installed as follows:

- Navigate to **System > Packages**, **Available Packages** tab
- Locate the package to install in the list

Tip: Search for a package by entering a value in the **Search term** box and clicking **Search**

- Click **Install** to the right of the package entry
- Click **Confirm** to proceed with the package installation

After confirming the installation, the GUI displays the package installation progress (Figure *Post-Install Package Screen*).

30.1.2 Reinstalling and Updating Packages

Packages are reinstalled and updated the same way they are installed:

- Navigate to **System > Packages**, **Installed Packages** tab

  The list will look like Figure *Installed Package List*
- Locate the package to reinstall or update in the list

  If there is a newer version available than is installed, the **Package Version** column will state the old and new versions with special highlighted text
- Click to update or to reinstall the package
- Click **Confirm** to proceed with the package reinstallation

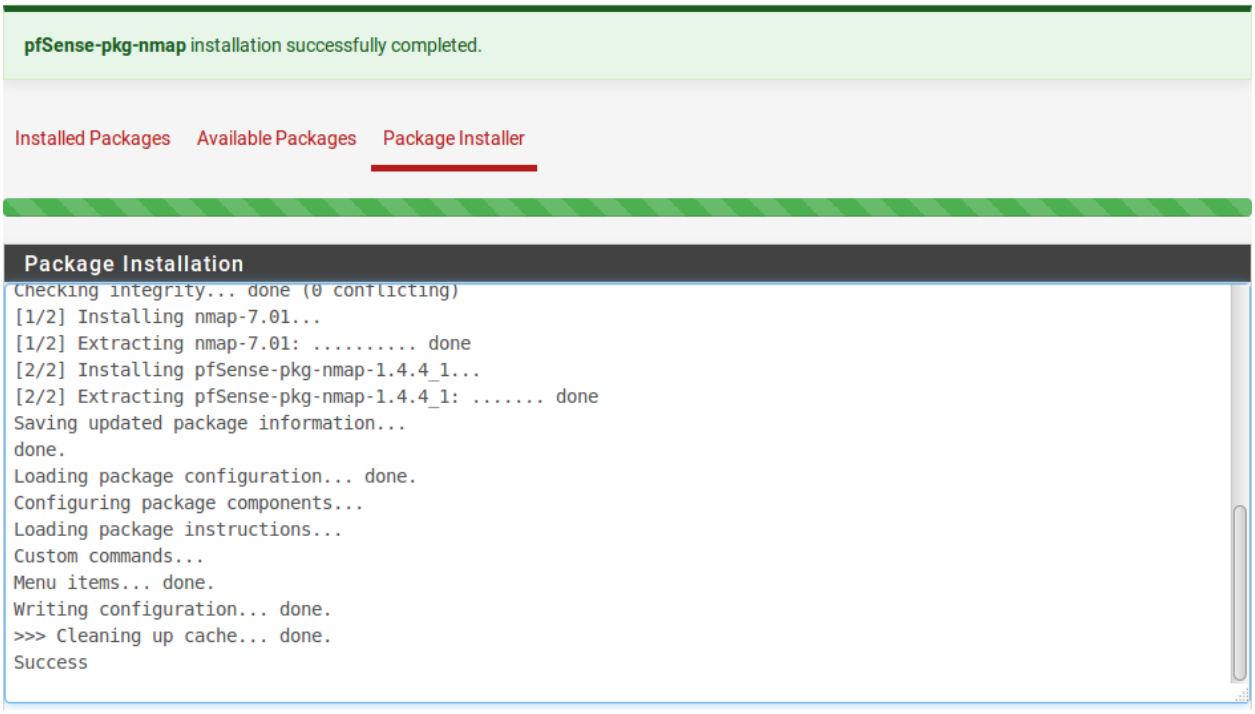


Fig. 2: Post-Install Package Screen

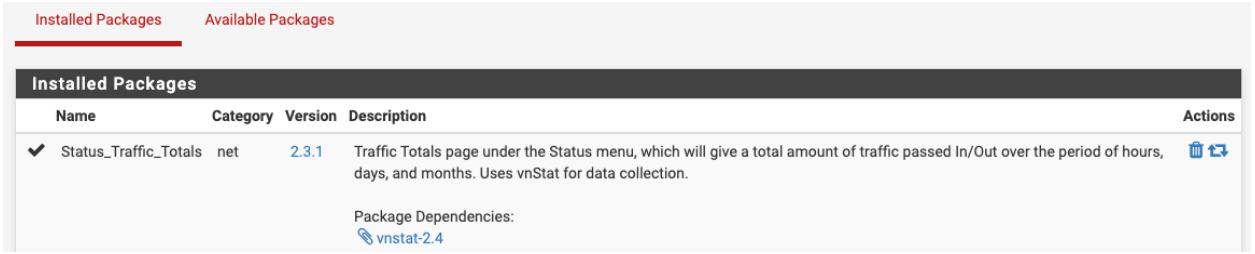


Fig. 3: Installed Package List

30.1.3 Uninstalling Packages

To uninstall a package:

- Navigate to **System > Packages**, **Installed Packages** tab
- Locate the package to uninstall in the list
- Click  to remove the package
- Click  **Confirm** to proceed with the package removal

30.2 Package List

The following packages are available from the pfSense® software package repository.

Warning: Packages availability can change over time. Check **System > Package Manager > Available Packages** for an always up-to-date list of packages.

Tip: The package name in the list below links to documentation for the package, if it exists.

ACME
: The Automated Certificate Management Environment (ACME) package manages certificates from ACME providers such as Let's Encrypt.

  See also:

  ACME package

arping
: Broadcasts a `who-has` ARP packet on the network and prints answers.

  See also:

  Arping Package

arpwatch
: Monitors devices on directly attached networks and notifies when it detects new MAC addresses.

apcupsd
: Controls all APC UPS models. It can monitor and log the current power and battery status, perform automatic shutdown, and can run in network mode to power down other hosts over the network.

aws-wizard **(pfSense Plus Only)**
: AWS VPC VPN Connection Wizard. Automatically creates a VPN tunnel and BGP configuration to communicate with an Amazon AWS VPC.

Avahi
: Facilitates service discovery on a local network via the mDNS/DNS-SD protocol suite. This enables clients to plug a laptop or computer into a network and instantly be able to view other people who they can chat with, find printers to print to or find files being shared. In addition, it Avahi can reflect mDNS across LAN segments. Compatible technology is found in Apple macOS (branded Bonjour and sometimes Zeroconf).

See also:

Avahi package

Backup
Backs up and restores arbitrary files and directories.

See also:

Backup Files and Directories with the Backup Package

bandwidthd
Tracks TCP/IP network usage and creates graphs of data consumption for individual IP addresses.

BIND
Provides a GUI for BIND DNS server.

cellular
Provides a GUI for cellular cards (e.g. 4G/LTE/5G), compatible with certain Huawei models.

Cron
Manages scheduled commands run periodically by the firewall.

Darkstat
A network statistics gatherer that offers bandwidth graphs for interfaces, as well as traffic to/from specific IP addresses. Once installed, it appears under **Services > darkstat**.

filer
Stores custom files persistently in the configuration.

FreeRADIUS
A free implementation of the RADIUS protocol, used for Authentication, Authorization, and Accounting (AAA).

See also:

FreeRADIUS package

FRR
A GUI for the FRR routing daemon which provides dynamic routing using BGP, OSPF, and OSPF6.

FTP Client Proxy
A basic FTP client proxy using `ftp-proxy` from FreeBSD.

HAproxy
A reliable, high performance TCP/HTTP(S) load balancer. This package implements the TCP, HTTP and HTTPS balancing features from haproxy and ACLs for smart backend switching. Requires SSD/HDD.

See also:

HAProxy

HAproxy-devel
The development package for *HAproxy*.

iperf
A tool for testing network throughput, loss, and jitter. Can act as a client or a server.

See also:

iperf package

ipsec-profile-wizard **(pfSense Plus Only)**
Creates IPsec configuration profiles for Apple devices (iOS and macOS) and IPsec import script bundles for Windows devices.

LADVD
Sends and decodes link layer advertisements.

Can utilize a variety of protocols, including: LLDP (Link Layer Discovery Protocol), CDP (Cisco Discovery Protocol), EDP (Extreme Discovery Protocol) and NDP (Nortel Discovery Protocol).

See also:

Using LLDP on pfSense software

LCDproc
LCD device drivers and service.

Lightsquid

> **Danger:** The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

A high performance web proxy reporting tool. Includes real-time proxy statistics (SQStat). Requires the **Squid** package. Requires SSD/HDD.

lldpd
Provides 802.1ab Link Layer Discovery Protocol (LLDP) service, as well as several proprietary discovery protocols including Cisco Discovery Protocol (CDP), Extreme Discovery Protocol (EDP), Foundry Discovery Protocol (FDP), and Nortel Discovery Protocol (NDP / SONMP).

Similar to LADVD but a more modern implementation.

Mailreport
Manages periodic e-mail reports containing command output and log file contents.

MTR
An enhanced traceroute replacement. `mtr` combines the functionality of the traceroute and ping programs in a single network diagnostic tool.

Netgate Firmware Upgrade **(pfSense Plus Only)**
Provides a mechanism to update firmware on certain Netgate hardware models. Varies by hardware and may be Coreboot, Blinkboot, or other types of firmware.

net-snmp
The NET-SNMP implementation of SNMP. It is more extensible than the built-in SNMP daemon (bsnmpd), and it includes SNMPv3 authentication and TLS encryption.

nmap
A utility for network exploration and security auditing.

Its capabilities include: scanning to determine active hosts, many port scanning techniques to determine services offered by hosts, version detection to determine what application/service is running on a port, and TCP/IP fingerprinting to identify the OS on remote hosts. It also offers flexible target and port specification, decoy/stealth scanning, SunRPC scanning, and more.

See also:

Nmap package

node_exporter
Prometheus exporter for machine metrics.

Notes
Maintains a list of noteworthy items for the system.

NRPE
Provides a GUI for Nagios NRPE. It executes Nagios plugins on remote hosts and report the results to the main Nagios server.

It also allows Nagios to execute plugins like `check_disk`, `check_procs`, etc. on remote hosts.

ntopng
A network probe that shows network usage in a way similar to what top does for processes. In interactive mode, it displays the network status on the user's terminal. In Web mode it acts as a Web server, creating an HTML dump of the network status. It sports a NetFlow/sFlow emitter/collector, an HTTP-based client interface for creating ntop monitoring applications, and RRD for persistently storing traffic statistics. Requires SSD/HDD.

Network UPS Tools (NUT)
Monitors Uninterruptible Power Supplies. It is compatible with UPS units attached locally via USB or serial, and remote units via the SNMP protocol, the APCUPSD protocol or the NUT protocol.

See also:

NUT package

Open-VM-Tools
A suite of open source utilities which enhance the performance of VMware virtual machine guest operating systems and improve management of virtual machines.

See also:

Open VM Tools package

OpenVPN Client Export
Generates pre-configured OpenVPN configuration files for clients, Windows Client installers with configurations bundled, and macOS Viscosity configuration bundles, among others.

See also:

OpenVPN Client Export Package

OpenVPN Client Import **(pfSense Plus Only)**
Imports a unified OpenVPN client configuration file as exported by an OpenVPN server, allowing clients to be easily configured without creating a client instance and adding settings manually.

pfBlockerNG
Utility for controlling connections through the firewall based on more general criteria than firewall rules (e.g. by country, by domain name, etc.). Manages IPv4/v6 List Sources into 'Deny, Permit or Match' formats. GeoIP database by MaxMind Inc. (GeoLite2 Free version). De-Duplication, Suppression, and Reputation enhancements. Provision to download from diverse List formats. Advanced Integration for Proofpoint ET IQRisk IP Reputation Threat Sources. Domain Name (DNSBL) blocking via Unbound DNS Resolver.

See also:

pfBlocker-NG Package

pfBlockerNG-devel
The development version of pfBlockerNG

See also:

pfBlocker-NG Package

PIMD
A GUI for `pimd`, a multicast routing daemon. Primarily replaces the role of the built-in IGMP Proxy function to allow routing multicast traffic across multiple interfaces. Not a replacement for Avahi.

RRD Summary
Gives a total amount of traffic passed In/Out during this and the previous month. Set to be replaced by the **Traffic totals** package.

Service Watchdog
Monitors for stopped services and restarts them.

Shellcmd
Manages boot-time commands.

See also:

Executing Commands at Boot

Siproxd
A proxy for handling multiple SIP devices using a single public IP address.

See also:

Siproxd package

snmptt
SNMP Trap Translator for use with the Net-SNMP. Easy to set up and use.

Snort
An open source network intrusion detection and prevention system (IDS/IPS). Combining the benefits of signature, protocol, and anomaly-based inspection. SSD/HDD is strongly recommended.

See also:

IDS / IPS

Softflowd
A flow-based network traffic analyzer capable of Cisco NetFlow data export. Tracks traffic flows and reports via NetFlow to a collecting host.

See also:

- *Firewall Packet Flow Data* (Plus only)
- *Exporting NetFlow with softflowd*

Squid

> **Danger:** The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

A high performance web proxy cache. It combines Squid as a proxy server with its capabilities of acting as an HTTP/HTTPS reverse proxy. It includes an Exchange-Web-Access (OWA) Assistant, SSL filtering and antivirus integration via C-ICAP. SSD/HDD recommended.

See also:

Squid

SquidGuard

> **Danger:** The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

A high performance web proxy URL filter. SSD/HDD recommended.

See also:

Configuring the SquidGuard Package

Status Traffic Totals
Calculates a total amount of traffic passed In/Out over the period of hours, days, and months. Uses vnStat for data collection. It shows up in the menu under **Status > Traffic Totals**.

See also:

Status Traffic Totals

Stunnel
A TLS encryption wrapper between a remote client and local or remote servers.

See also:

Stunnel package

Sudo
Delegates privileges to users in the shell, so commands can be run as other users, such as `root`.

See also:

Sudo Package

Suricata
A high performance network IDS/IPS and security monitoring engine by OISF. SSD/HDD strongly recommended.

Syslog-ng
A modern syslog server capable of using TCP and TLS encryption, among other features.

Note: This service is not intended to replace the default syslog server on the firewall but rather acts as an independent syslog server.

System Patches
Manages custom code patches to be applied and maintained to the system. These can be commits from GitHub, manual diffs, or loaded from URLs.

See also:

System Patches Package

Telegraf
An agent written in Go for collecting, processing, aggregating, and writing metrics.

TFTPD
GUI for a TFTP server, using the versatile **tftp-hpa** daemon.

Tinc
A Virtual Private Network (VPN) daemon that uses tunneling and encryption to create a secure private network between hosts on the Internet. A single tinc daemon can accept more than one connection at a time, thus making it possible to create larger virtual networks, because some limitations are circumvented. Instead of most other

VPN implementations, tinc encapsulates each network packet in its own UDP packet, instead of encapsulating all into one TCP or even PPP over TCP stream. This results in lower latency, less overhead, and in general better responsiveness and throughput.

UDP Broadcast Relay
A GUI for UDP Broadcast Relay. This program listens for UDP broadcast packets and retransmits them on additional interfaces.

WireGuard®
WireGuard is a new VPN Layer 3 protocol designed for speed and simplicity. It performs nearly as fast as hardware-accelerated IPsec and has only a few options in its configuration.

Zabbix-agent
Zabbix Monitoring agent. The agent gathers operational information locally and reports data to Zabbix server for further processing. The agent can also generate alerts in case of failures. Available in multiple versions.

Zabbix-proxy
Zabbix Agent proxy. Collects performance and availability data on behalf of the Zabbix server, lowering the burden on the server. Available in multiple versions.

The package system in pfSense® software provides the ability to extend the functionality of the software without adding bloat and potential security vulnerabilities to the base distribution. To see the packages available for the current firewall platform being utilized, browse to **System > Packages**, on the **Available Packages** tab.

Some packages have been written by the pfSense community, and others directly developed by Netgate. The available packages vary quite widely, and some are more mature and well-maintained than others. There are packages which install and provide a GUI interface for third-party software, such as HAProxy, and others which extend the functionality of pfSense software, like the OpenVPN Client Export Utility package which automatically creates OpenVPN configuration files.

Some other examples of available packages are:

- Additional filtering functionality (pfBlockerNG)
- IDS/IPS software (Snort and Suricata)
- Additional VPN technologies (WireGuard, Tinc)
- Bandwidth monitors that show traffic by IP address such as ntopng, and Darkstat.
- Extra services such as FreeRADIUS and BIND.
- Reverse proxies for HTTP and HTTPS such as HAProxy
- Proxies for other services such as SIP and FTP
- System utilities such as NUT for monitoring a UPS.
- Popular third-party utilities such as nmap, iperf, and arping.
- BGP Routing, OSPF routing, Cron editing, Zabbix agent, and many others.
- Features that were formerly in the base system but were moved to packages, such as RIP (routed)

As of this writing there are more than 60 different packages available; too many to cover them all in this documentation! The full list of packages that can be installed on a particular system is available from within the GUI at **System > Packages** on the **Available Packages** tab.

The packages screen takes longer to load than other pages in the web interface because the firewall fetches package information from the Netgate package repository servers before the page is rendered to provide the most up-to-date information. If the firewall does not have a functional Internet connection including DNS resolution, this will fail. This is usually caused by a missing or incorrect DNS server configuration, or a missing default gateway.

See also:

- *Troubleshooting Upgrades* (Packages and upgrades both use `pkg`)
- *Troubleshooting DNS Resolution Issues*
- *Troubleshooting Routes*
- *Troubleshooting a Broken pkg Database*
- *Developing Packages*

Danger: Installing packages from repositories other than those provided by Netgate is **unsupported** and carries a significant risk of breaking the installation.

Packages installed from manually configured non-Netgate repositories may result in failed package installations, upgrades, or unbootable systems. Even if the third party repository configurations and packages are removed, packages from those repositories may have already damaged the base system in unknown ways. If an administrator chooses to ignore these warnings and use such packages, they must take extra caution and have installation media ready in case the device is unable to install packages, upgrade packages, or even reboot successfully.

30.3 ACME package

The ACME package on pfSense software enables users to obtain certificates from providers who run servers compatible with the Automatic Certificate Management Environment (ACME) protocol.

The most common ACME service is provided by Let's Encrypt. Let's Encrypt is an open, free, and completely automated Certificate Authority from the non-profit Internet Security Research Group (ISRG). The goal of Let's Encrypt is to encrypt the web by removing the cost barrier and some of the technical barriers that discourage server administrators and organizations from obtaining certificates for use on Internet servers, primarily web servers.

Most browsers trust certificates from Let's Encrypt and some other common ACME providers. These certificates can be used for web servers (HTTPS), SMTP servers, IMAP/POP3 servers, and other similar roles which utilize the same type of certificates. By using a certificate from a trusted provider such as Let's Encrypt for a web server, including a firewall running pfSense software, client browsers and other software will trust the certificate and show a green check mark, padlock, or similar indication. The connection will be encrypted without the need for a client to manually trust an invalid or self-signed certificate.

The ACME Package for pfSense® software interfaces with ACME servers run by Let's Encrypt and other similar providers to handle the certificate generation, validation, and renewal processes.

Note: Though the ACME package works with other compatible providers, the information here is primarily focused around the Let's Encrypt service as it is the most common. Other ACME servers may operate similarly, but they may not offer the same features provided by Let's Encrypt or be trusted as widely. Some also require additional registration procedures which must be completed before attempting to issue a certificate.

Check with the provider and test thoroughly before using any certificate in a production role.

ACME providers such as Let's Encrypt perform validation before issuing certificates, and this validation ensures that the system requesting the certificate has authority over the hostname or domain in question. This validation can be performed in a number of ways, such as by proving ownership of the domain's DNS records or publishing a file on a web server for a hostname. Check with the ACME provider for specific information on available validation methods.

ACME providers generally limit certificate validity to a short time frame, so they must be renewed periodically. For example, certificates signed by Let's Encrypt are valid for 90 days. The ACME package automates the renewal process by using a scheduled job to check once per day for certificates nearing expiration.

30.3.1 ACME Overview

Rate Limits

ACME providers such as Let's Encrypt often enforce rate limitations when using production validation systems.

Specifically, rate limits imposed by Let's Encrypt include:

- Five validation failures per account, per hostname, per hour
- Each certificate may have at most 100 SAN entries
- Only 50 certificates may be created per domain per week

Some providers offer a test validation system for administrators testing their settings or developers who are programming clients. Test systems typically have higher limits, which aid testing and development, but the test system does not produce certificates which are trusted publicly.

Note: Other providers may have different rate limits. Some providers do not offer staging servers, only production servers.

Security Limitations

When ACME validates a certificate using a method such as *webroot* or *standalone*, the service must be available to the Internet on its standard port: 80 for HTTP or 443 for TLS-ALPN. This is a security limitation to prevent a user from running an alternate web server on a high-numbered port and obtaining a certificate for a server they do not normally control.

Validation Process

When creating a certificate, one or more fully qualified domain names (FQDNs) are listed on the certificate in the SAN list. Let's Encrypt will query each of these domain names in DNS in different ways depending on the validation method.

When a validation method starts, the client obtains an authorization value from the server ("authz").

For DNS-based methods, Let's Encrypt checks for a TXT record in the form of `_acme-challenge.<domain name>` which must contain the authorization value. This proves that the person or system requesting the certificate controls DNS records for the domain.

For File-based methods such as *webroot* or *standalone*, Let's Encrypt connects to an IP address obtained by resolving the A record for the FQDN and requests a file from the web server at `.well-known/acme-challenge/` underneath the `webroot` directory. This file contains the authorization value. This proves that the person or system requesting the certificate controls web server for the domain name.

Note: ACME providers often do not publish which IP addresses or address families they use when performing validation. As such, the port must be accessible to any client on the Internet which can resolve the hostname during the validation process.

30.3.2 Obtaining a Certificate

These instructions cover the general process of obtaining a certificate. Specific settings will vary by deployment, and each section below links to the settings for each area.

Generate an Account Key

Before a certificate can be created by the firewall, the firewall must first obtain an account key. This key is typically unique for each server, but can be shared.

For users unfamiliar with the Let's Encrypt service, the first key should be for the staging system which has no rate limits but is not valid for public use. Once a certificate is successfully issued by the staging system, create an account key for the production system and then issue the certificate again using that key.

To create and register an account key:

- Navigate to **Services > ACME Certificates**, **Account Keys** tab
- Click  **Add**
- Fill in the info as described in *Account Key Settings*
- Click  **Create new account key**
- Click  **Register ACME account key**
- Click **Save**

Create a certificate

The next step is to create a certificate entry.

- Navigate to **Services > ACME Certificates**, **Certificates** tab
- Click  **Add**
- Fill in the info as described in *Certificate Settings*
- Add one or more **Domain SAN List entries** (*Certificate Settings*) with appropriate validation settings (*Validation Methods*)
- Add one or more **Actions list** entries (*Certificate Settings*)
- Click **Save**

Configure General Settings

The last configuration step is to enable at least the **Cron Entry** to ensure that the ACME package will automatically renew certificates before they expire. See *General Settings* for detailed descriptions of the options.

- Navigate to **Services > ACME Certificates**, **General Settings** tab
- Check **Cron Entry**
- Check **Write Certificates** (optional)
- Click **Save**

Issue a Certificate

With the settings in place, the final task is to issue the certificate:

- Navigate to **Services > ACME Certificates**, **Certificates** tab
- Find the certificate entry in the list
- Click  **Issue/Renew**
- Wait for the process to complete

This step may take a few minutes to complete as it sets up the validation, contacts the ACME server, then waits for the server to complete validation.

Once complete, the page will display the result at the top. Check the output to ensure that it was successful and included a certificate.

If the process was successful, the package will import the certificate along with its associated certificate authorities in the GUI Certificate Manager. Check at **System > Certificates** to ensure the entries are present.

30.3.3 ACME Package Settings

These sections describe the settings for each tab in the ACME package.

Account Key Settings

An ACME account key has the following settings:

Name
A short name for the key

Description
A longer string describing the key

ACME Server
The ACME server to which this key will be registered by the package.

Available servers include:

Let's Encrypt Staging ACMEv2
Use this server when testing the certificate validation process. Does not produce publicly trusted certificates.

Let's Encrypt Production ACMEv2
Use this server for trusted production certificates.

ZeroSSL
Certificates issued by the ZeroSSL service.

SSL.com RSA/ECC
Certificates issued by SSL.com either with RSA or ECC keys.

Google ACMEv2
Certificates issued by Google, either for staging/testing or production.

The selection list also includes *Custom ACME servers*.

E-Mail Address
An e-mail address to which the service might send important communications about the account or certificates.

Note: The Let's Encrypt service no longer sends expiration notices for certificates via e-mail. They still send announcements to this address, such as when certificate features are added or deprecated.

Account Key

The RSA private key for this entry. To create a new key, click  **Create new account key**.

Key Registration

Click  **Register ACME account key** to register the account key with the selected ACME server. This registration may happen automatically while issuing or renewing a certificate, but doing so manually allows the user to ensure the registration process is working before attempting validation.

Warning: The Let's Encrypt service does not require any special registration steps, the above procedure will work as written with that service.

Some ACME servers have manual key validation and/or registration procedures which must be performed before attempting to issue a certificate. Check with the provider before creating a key.

Certificate Settings

Certificate entries have the following settings:

Name
A short name for the certificate

Description
A longer string describing the certificate

Status
Whether this entry is active

Active
This entry will be processed manually and by the Cron job (*General Settings*)

Disabled
This entry will be ignored

Acme Account
The account key ACME will use when requesting the certificate (see *Generate an Account Key*)

Private Key
The key length of the private key for this certificate. Can be either RSA or ECDSA in several predefined sizes. Select *Custom* to manually enter a private key generated elsewhere

2048-bit RSA is an acceptable default choice, but larger keys are more secure

Preferred Chain
If the ACME CA provides multiple trust chains, this field allows the user to choose an alternate preferred chain. This utilizes a case-insensitive substring match to locate a chain.

Domain SAN List
A list of all domain names which will be included in this certificate as Subject Alternative Name (SAN) entries.

SAN entries can use the same or different update methods. Each SAN must be individually validated by the ACME server before it will issue or renew a certificate.

Note: The ACME server may limit the number of SAN entries. For example, Let's Encrypt limits a certificate to at most 100 SAN entries.

Mode
Whether this SAN is active in the certificate

Domain Name
The domain name for a SAN entry in this certificate (e.g. `www.example.com`)

Method
The method used by ACME to validate ownership of this domain. Method settings are described in (*Validation Methods*)

Click  **Add** for additional SAN entries

DNS Providers also have some common settings which appear for all types:

DNS Alias
An alternative domain name used by the validation process. Instead of updating the DNS record for **Domain Name** directly, the package uses this domain name is used instead. See *DNS Alias Mode* for details.

DNS Alias Mode
When set, controls whether the DNS alias mode used is Challenge Alias (Unchecked, Default) or Domain Alias (Checked). See *DNS Alias Mode* for details.

DNS-Sleep
The amount of time the ACME validation process will wait after making DNS changes before attempting to validate. Some DNS services take a few minutes to propagate entries after making backend changes.

The default settings are typically sufficient, but slower providers may require a longer sleep time.

Actions List
Commands to run after the package renews a certificate.

Mode
Whether this action is active.

Command
Full path to command and arguments, service name, or name of script.

Method
Defines how the **Command** is executed by the package.

Shell Command
The **Command** is a full path to a shell command and its arguments.

PHP Command Script
The **Command** value is run as PHP code.

Restart Local Service
The name of a local service to restart.

Restart Remote Service
The name of a remote service to restart via XMLRPC. This utilizes the system *XMLRPC sync configuration*.

The GUI help text for this option includes several examples of common actions.

Certificate Renewal After
When the package will attempt a renewal for the certificate. Default is `60` days (2 months).

Note: Certificates issued by the Let's Encrypt service are valid for a maximum of `90` days. When using other ACME servers, check with the provider to find its default and maximum validity intervals.

Tip: The best practice is to attempt renewal at around two-thirds of the certificate's total lifetime. That way, any problems renewing the certificate can be resolved before it expires.

Validation Methods

ACME providers can validate by checking the contents of a TXT record in DNS, or by fetching a file in a known location from a web server.

The ACME package can validate directly with standalone methods or webroot, but those options are less secure than DNS-based options. The ACME package can also use numerous methods to validate through DNS providers. Wildcard certificates can only be obtained through DNS-based methods (*Wildcard Certificates*)

Tip: DNS-based update methods are the best practice as they do not require external inbound access. They can be used for internal systems that do not allow or cannot receive Internet traffic.

The following list is only a portion of the validation methods available in the package.

See also:

DNS methods also have a common option for DNS Alias mode. See *Certificate Settings* and *DNS Alias Mode* for details.

nsupdate

The `nsupdate` method uses RFC 2136 style DNS updates to populate a TXT record in DNS.

Before starting, an appropriate DNS key and settings must be in place in the DNS infrastructure for the domain to allow the host to update a TXT DNS record for `_acme-challenge.<domain name>`.

This method has the following options:

Server
: The IP address or hostname of the DNS server to which the client sends updates

Key Name
: The name of the update key

  Leave this blank unless it is different from `_acme-challenge.<domain name>`

Key Algorithm
: The algorithm used for the key, which must match the key and the server

Key
: The update key for this record

Zone
: Sets the zone name the package sends to the DNS server in the update request

DNS-Manual

The manual DNS method can be utilized when a firewall cannot receive inbound traffic, and it does not have access to any automatic DNS-based method.

The **manual** in the name indicates that the process **must be performed by hand** both initially and when it is time to renew the certificate. The firewall obtains the authorization value and then the TXT record must be manually created or updated with this value.

> **Warning:** Avoid using this method unless no other method is available.

To use this method:

- Add an entry to the **Domain SAN list**
- **Mode**: Enabled
- Enter domain name (e.g. `myhost.example.com`)
- Set **Method** to *DNS-Manual*
- Click **Save**
- Click **Issue**
- Locate the record info in the output:

```
[Mon Feb 6 14:49:23 EST 2017] Add the following TXT record:
[Mon Feb 6 14:49:23 EST 2017] Domain: '_acme-challenge.www.example.com'
[Mon Feb 6 14:49:23 EST 2017] TXT value: 'xPrykHSri5epT5yrJJWyY536Z1T51r_Ef4LkWJry-
↪iw'
[Mon Feb 6 14:49:23 EST 2017] Please be aware that you prepend _acme-challenge.␣
↪before your domain
```

(continues on next page)

(continued from previous page)

```
[Mon Feb 6 14:49:23 EST 2017] so the resulting subdomain will be: _acme-challenge.
↪www.example.com
[Mon Feb 6 14:49:23 EST 2017] Please add the TXT records to the domains, and retry␣
↪again.
```

- Add or update the TXT record in the domain's DNS server for `_acme-challenge.<domain name>` with the TXT value from the output
- Wait approximately 2 minutes, or longer, for DNS to propagate
- Click **Renew**

Namecheap API

For certain accounts with Namecheap, API access may be obtained that allows remote manipulation of DNS records. This can be used with the ACME package to validate certificates for domains with DNS hosted at Namecheap using their BasicDNS servers.

Warning: The Namecheap DNS API requires that the client read all records and then write them all back when making any change. This is potentially dangerous. Take a backup of all DNS records on the domain before attempting to use the API.

The first step is to request API access:

- Login to a Namecheap account
- Navigate to **Profile > Tools** under the account
- Look for **Namecheap API Access** under **Business & Dev Tools**
- If the status does not say **On**, then click **Manage** and change the slider to **On**.

Note: API access must be approved by Namecheap. There are qualifications to meet, such as a specific number of domains or a balance on the account. Check the Namecheap API documentation for more information. The process is documented as taking 2 days, but may take longer. If API access is not enabled after several days, contact Namecheap support.

Once the API is enabled, then perform the following steps:

- Login to a Namecheap account
- Navigate to **Profile > Tools** under the account
- Look for **Namecheap API Access** under **Business & Dev Tools**
- Click **Manage**
- Note the API key for use in the ACME package
- Click **Edit** and add whitelisted IP addresses that can contact the API using this API key.

Now set up the account in the ACME package:

- Add an entry to the **Domain SAN list**
- **Mode**: Enabled

- Enter domain name (e.g. `myhost.example.com`)
- Set **Method** to *DNS-Namecheap*
- Click + to expand the method-specific settings
- Fill in the info

  API Key
  The API Key displayed in the Namecheap API Access manager, as described previously.

  Username
  The Namecheap account username associated with the API Key.
- Ensure the other options are set properly, per *Create a certificate*.
- Click **Save**
- Click **Issue/Renew**

Other DNS Methods

The package contains several additional DNS-based methods for other providers. These work similar to the nsupdate method above, but have configuration values specific to each provider. Contact the DNS provider or server administrator to obtain the necessary settings or credentials.

FTP Webroot

The **FTP webroot** method is useful when the firewall is performing NAT (port forward or 1:1) or reverse proxy duty for handling traffic for the domain. The firewall can use SFTP or FTPS to store the domain validation files on a web server behind the firewall so it does not have to host the files itself.

This method has the following options:

Server
URL for the server where the package will send the challenge response files, e.g. `sftp://x.x.x.x`

Note: This method can communicate with servers using `sftp://` and `ftps://` URLs.

Username/password
Credentials for the SFTP/FTPS account

Folder
Full path to the target directory including `/.well-known/acme-challenge` at the end

> **Warning:** Make sure the specified user has write permissions to the directory!

Webroot Local Folder

This method works similar to FTP Webroot but with the files hosted on the firewall itself. This method cannot be utilized by the GUI web server as that would mean exposing the GUI to the Internet, which is a major security issue.

This method can, however, be used in conjunction with the HAProxy package to host the files on the firewall itself in some circumstances. See https://forum.netgate.com/post/677786 for details.

Standalone (HTTP/TLS-ALPN)

The **Standalone** methods for HTTP and TLS-ALPN run a small web server natively that is active only while the validation process is running. The TLS-ALPN method is more secure as it encrypts communication with the ACME provider.

Warning: These methods are not best practices as they expose a service on the firewall to the Internet. Only use these methods if no other method is available.

Warning: The service **must** be accessible using port `80` (HTTP) or `443` (TLS-ALPN)!

If the firewall is using port `80` (HTTP) or `443` (TLS-ALPN) for another service, such as the firewall GUI or its redirect, then this method may not be viable. If the service on the port is public, then it cannot be used. If the service is private, then it may be possible to relocate the existing service or bind the update method to an alternate port, then port forward on the WAN interface.

A firewall rule must allow traffic to the target port at all times, it cannot be automatically enabled and disabled in the current package.

Note: The standalone binding should only be changed if the port is forwarded via NAT to a different port (e.g. 80 forwarded to 8080)

Standalone HTTP

Standalone HTTP Server has the following options:

Port
: Port to which the package will bind listening for HTTP requests with a standalone server. Must be `80` or port `80` must be forwarded to this port on the default gateway WAN.

  Warning: If port `80` is used by the standalone service, the GUI redirect must be disabled on **System > Advanced** using the **Disable webConfigurator redirect rule** option. If the redirect is active when standalone mode attempts to use the port, it will print an error message stating that `socat` is unable to bind to the port.

Bind to IPv6 instead of IPv4
: If the domain name for the firewall has both an A and AAAA DNS record, check this option so that validation can occur over IPv6.

Standalone TLS-ALPN

Standalone TLS-ALPN Server has the following options:

Port
Port to which the package will bind listening for TLS-ALPN requests with a standalone server. Must be `443` or port `443` must be forwarded to this port on the default gateway WAN.

> **Warning:** If port `443` is used by the standalone service, the GUI must be moved to an alternate port on **System > Advanced** using the **TCP Port** option for the GUI. If the redirect is active when standalone mode attempts to use the port, the standalone service will fail.

DNS Alias Mode

DNS Alias mode allows a DNS update method to update an alternate domain name instead of updating a record for the domain name directly.

If the main DNS provider does not allow updating TXT records, a CNAME record can point to an alternative domain which does.

Challenge Alias

In Challenge Alias mode (default), the ACME package still automatically prepends `_acme-challenge.` to both the **Domain Name** and the **DNS Alias** domain.

In the certificate entry, set:

Domain Name
`company.example` which cannot use automatic updates

DNS Alias Domain
`dynamic.example` which is the alternative domain in a dynamic zone

DNS Domain Alias mode
Leave unchecked

On the DNS server, add a CNAME record pointing to the **DNS Alias** hostname with `_acme-challenge.` prepended:

```
_acme-challenge.company.example      IN      CNAME   _acme-challenge.dynamic.example.
```

When updating, the package will update `_acme-challenge.dynamic.example` in DNS while sending `company.example` in the certificate request to the ACME provider.

Domain Alias

Domain Alias mode works similar to Challenge Alias mode, but it **does not** prepend `_acme-challenge.` to the **DNS Alias** domain. Some administrators prefer this when using many hostnames in a single dynamic zone, or for working around limitations in DNS providers or platforms.

In the certificate entry, set:

Domain Name
`company.example` which cannot use automatic updates

DNS Alias Domain
`checkme.dynamic.example` which is the alternative domain in a dynamic zone

DNS Domain Alias mode
Checked

On the DNS server, add a CNAME record pointing directly to the **DNS Alias** hostname:

```
_acme-challenge.company.example      IN      CNAME   checkme.dynamic.example.
```

When updating, the package will update `checkme.dynamic.example` in DNS while sending `company.example` in the certificate request to the ACME provider.

General Settings

These settings control the general behavior of the ACME package and are not specific to any single certificate or key.

Cron Entry
A checkbox which enables the ACME renewal cron job. When set, the ACME package will check all certificates each night and if any are up for renewal, it will attempt to renew them.

Write Certificates
When set, the ACME package will write the certificate files out in `/conf/acme`. From there, other scripts or processes can read the certificate.

Custom ACME servers

This section allows users to define ACME services which are not available in the ACME package by default. These custom servers are available to select when creating *ACME Account Keys*.

> **Warning:** There is no way for the ACME package to determine if the service is valid or which validation methods it provides. Use this feature only after testing it thoroughly.

To define a custom ACME server, fill in the values for an entry:

Internal ID
An internal identifier used by the ACME package when referring to this server.

This value must be under 32 characters in length and use only lowercase letters, numbers, and dash, e.g. `my-acme-server`.

Display Name
The name which will appear in the ACME package GUI to select this server, e.g. `My ACME Server`.

Server URL
The full URL to the ACME server, e.g. `https://acme.example.com:4433/acme/acme/directory`.

30.3.4 Wildcard Certificates

Let's Encrypt offers wildcard certificates (e.g. `*.example.com`) through their ACMEv2 infrastructure. A wildcard certificate will work for any hostname inside a given domain, which helps with handling certificates for multiple domains. Check with other ACME providers to find out if, and how, they allow wildcard certificates.

Note: This is not related to ACME specifically, but any wildcard certificate: A wildcard certificate is only valid for **one level** of subdomains.

For example, a wildcard certificate for `*.example.com` will work for `host.example.com` and `host-sub.example.com`, but it **will NOT work** for `host.sub.example.com`. If hosts are structured in this way, a wildcard certificate is required for each sub-zone, e.g. `*.sub.example.com`.

Wildcard validation **requires a DNS-based method** and works similar to validating a regular domain. For example, to get a certificate for `*.example.com`, the package updates a TXT record in DNS the same as it would for `example.com`, which means the DNS record (and potentially key name) would be for `_acme-challenge.example.com`.

To obtain a wildcard certificate, follow the same procedures as other DNS validation methods, with the following differences:

- The **Account Key** must be registered with an ACME v2 server (staging for testing, or production)
- The **Domain SAN list** should contain entries for the base domain (e.g. `example.com`) and the wildcard version of the same domain (e.g. `*.example.com`). The settings will be the same for both entries.
- For *DNS-NSupdate / RFC 2136*: Set the **Key Name** to the base domain (`example.com`) for both entries.

30.4 Arping Package

arping is a utility to test the reachability and responsiveness of hosts to ARP. It is effectively like ICMP ping, except using ARP instead. This is beneficial in circumstances where the host has a firewall enabled (every host, even one with a firewall, will respond to ARP), or there is no layer 3 connectivity on the IP subnet of the host and hence cannot ping, but do have layer 2 connectivity.

The **arping** package can be very useful when trying to pick an unused IP address for a subnet to which there is not yet a route or link, but is connected at Layer 2.

See also:

Visit the arping website for more information.

30.4.1 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.5 Avahi package

The Avahi package used in pfSense® software is a system which facilitates service discovery on a local network. This means that a laptop or computer may be connected into a network and instantly be able to view other people to chat with, find printers to print to or find files being shared.

This kind of technology is also found in Apple macOS (branded Rendezvous, Bonjour, and sometimes Zeroconf) and is very convenient.

30.5.1 Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.5.2 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.6 AWS VPC Wizard

This guide explains how to use the AWS VPC Wizard, available in pfSense Plus software, to simplify the configuration of a VPN to a remote VPC. The administrator is asked for the minimum amount of basic information required to establish the VPN. The configurations, both on the AWS VPC side and on the pfSense Plus software side are then automatically created. When the wizard is finished executing, a functioning VPN connection to a VPC should be established.

30.6.1 AWS Access Keys

In order to connect to the AWS API to make certain required configuration changes, the AWS VPC Wizard requires Access Keys to retrieve and modify VPC configurations.

See also:

Find more information about AWS Security Credential, including Access Keys by reading AWS Security Credentials.

Access keys consist of two parts:

1. An access key ID
    - For example, `AKIAIOSFODNN7EXAMPLE`.
2. A secret access key
    - For example, `wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY`.

Access keys are like a username/password and needed for programmatic requests to AWS, including the AWS VPC Wizard. Use both the access key ID and secret access key together to authenticate requests.

Important: Manage access keys as securely as a username and password.

Managing Access Keys

To create, modify, or delete IAM user access keys, do the following:

1. Sign in to the IAM console.
2. In the navigation bar on the upper right, choose a username, and then choose **My Security Credentials**.
3. On the **AWS IAM Credentials** tab, in the **Access keys for CLI, SDK, and API access** section, choose **Create access key**.
4. Choose **Download .csv file** to save the access key ID and secret access key to a .csv file on the client computer. Store the file in a secure location. There is no access to the secret access key again after this dialog box closes. After the .csv file has been downloaded, choose **Close**. When an access key is created, the key pair is active by default, and the pair can be used right away.
    - To disable an active access key, choose **Make inactive**.
    - To re-enable an inactive access key, choose **Make active**.
    - To delete an access key, choose its **X** button at the far right of the row. Then choose **Delete** to confirm. When an access key is deleted, it's gone forever and cannot be retrieved. However, new keys can always be created.

To create, modify, or delete another IAM user's access keys, do the following:

1. Sign in to the IAM console.
2. In the navigation pane, choose **Users**.
3. Choose the name of the user whose access keys to manage, and then choose the **Security credentials** tab.
4. In the **Access keys** section, choose **Create access key**.
5. Choose **Download .csv file** to save the access key ID and secret access key to a CSV file on your computer.

Rotating Access Keys

As a security best practice, regularly rotate (change) IAM user access keys. Rotating access keys can be done from the AWS Management Console.

To rotate access keys for an IAM user without interrupting the applications (console), create a second access key while the first access key is still active:

1. Sign in to the IAM console.
2. In the navigation pane, choose **Users**.
3. Choose the name of the user whose access keys to manage, and then choose the **Security credentials** tab.
4. In the **Access keys** section, choose **Create access key**.
5. Choose **Download .csv file** to save the access key ID and secret access key to a CSV file on the client computer.
6. The new access key is active by default. At this point, the user has two active access keys.

After waiting some period of time to ensure that all applications and tools have been updated, delete the first access key:

1. Sign in to the IAM console.
2. In the navigation pane, choose **Users**.
3. Choose the name of the user whose access keys to manage, and then choose the **Security credentials** tab.
4. Locate the access key to delete and choose its **X** button at the far right of the row. Then choose **Delete** to confirm.

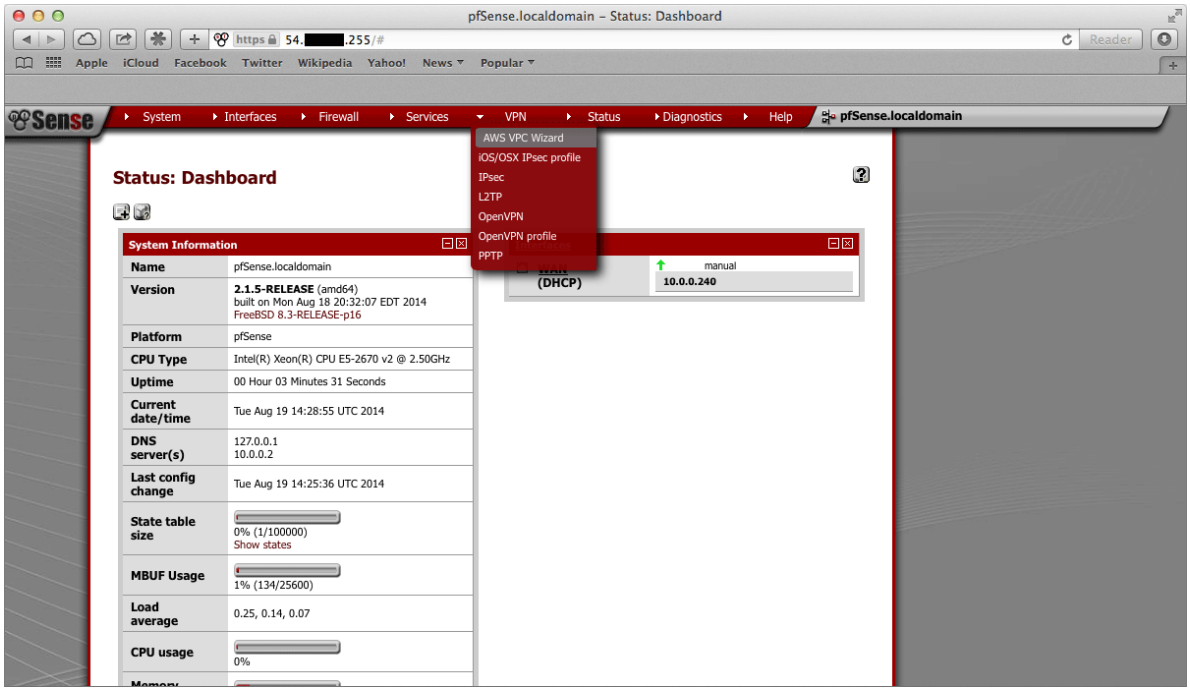
Determining When Access Keys Need Rotating

To determine when access keys need rotating (console), do the following:

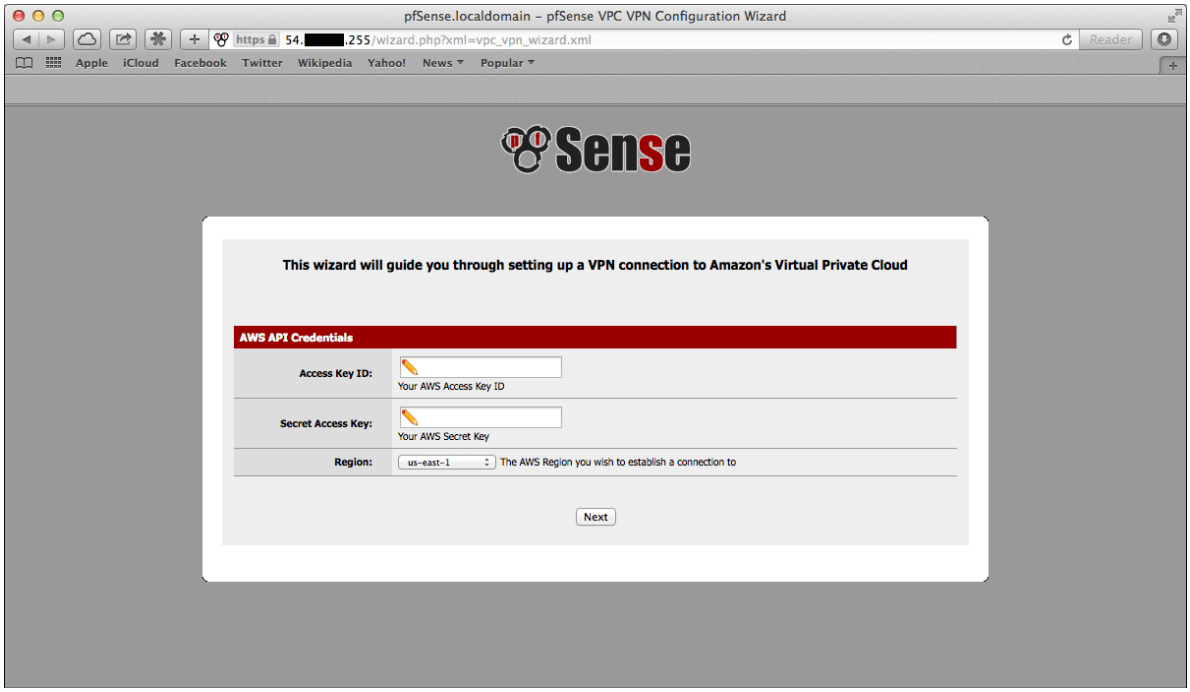
1. Sign in to the IAM console.
2. In the navigation pane, choose **Users**.
3. If necessary, add the **Access key age** column to the users table by completing the following steps:
    1. Above the table on the far right, click the settings icon.
    2. In **Manage columns**, select **Access key age**.
    3. Choose **Close** to return to the list of users.
4. The **Access key age** column shows the number of days since the oldest active access key was created. Use this information to find users with access keys that need rotating. The column displays **None** for users with no access key.

30.6.2 Using the Wizard

1. Gather the **AWS Access Key ID** and **Secret Key**.
2. Navigate to **VPN > AWS VPC Wizard** in the main menu in the Netgate® pfSense Plus GUI.



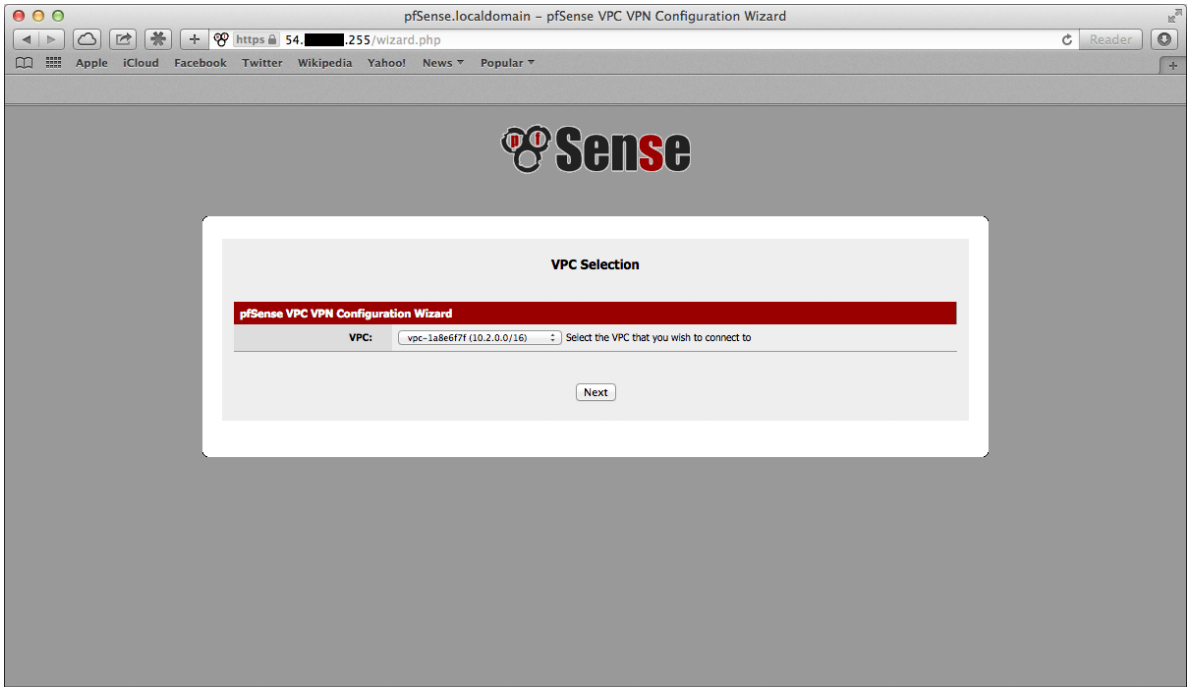
3. The first screen of the wizard prompts for the AWS Credentials. Enter the **Access Key ID** and **Secret Access Key** in the appropriate text fields and select the **Region** to connect to in the dropdown menu, then click **Next**.



The wizard will then query the AWS API using those credentials to find which VPCs exist in the selected region. If the credentials were rejected, an error message will be displayed and return to the first screen.

4. The next screen will prompt to select from the available VPC's in the selected region. Select the one to connect to from the dropdown menu. The wizard will not create a new VPC, it will only connect to an existing VPC.

   Click **Next** after selecting the desired VPC.

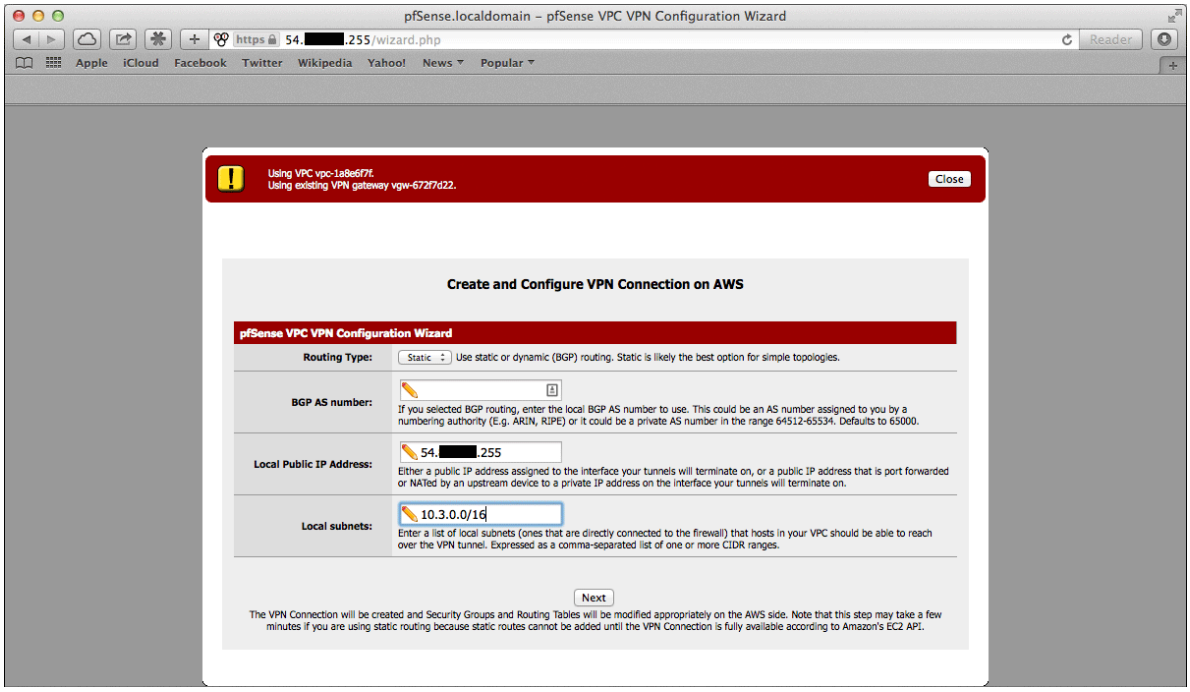


Note: If the desired VPC's to connect to isn't available, create one via the **AWS Management Console** and try

again.

The wizard will then query the AWS API to check whether there is a VPN Gateway attached to the selected VPC. If none exists, one will be created via the API. Then the next screen will be displayed.

5. On the next screen, specify routing and network data.



A description of what should be entered for each of these fields follows.

Routing Type
: AWS offers either **static** routing or **BGP** routing. Select the appropriate type from the dropdown menu. If unsure, static routing is likely to be adequate.

BGP AS Number
: If static routing was chosen, leave this field blank. Otherwise, it is possible to specify an AS number to use. If left blank, the value will default to `65000`.

Local Public IP Address
: On an AWS Netgate appliance instance, this should be the public IP address of the Elastic IP associated with the instance. If configuring a hardware device running pfSense Plus, this could be the public address assigned to the WAN (or other) interface of the device.

Local subnets
: The subnets connected to the pfSense Plus instance that should be routed over the tunnel from hosts in the remote VPC. As an example, if connecting a pfSense Plus instance to a remote VPC in the AWS **us-east-1** region, enter the subnets (or a single subnet) that are local to the pfSense Plus instance and when hosts in your VPC in **us-east-1** attempted to reach addresses within those subnets, the traffic will be sent through the VPN tunnel that is being configured.

Note: When selecting static routing as the routing type, there will be a delay that is typically between 2-5 minutes before the next screen is displayed. This is because static routes must be added to the VPN Connection via the AWS API. This operation fails until the VPN Connection reaches the "available" state. This can take a few minutes to occur.

Click **Next** when done.

The wizard will then query the AWS API to find whether a Customer Gateway is configured with the selected Public IP Address. If none exists, one will be created. If one already exists, the ID will be retrieved, and it will be used.

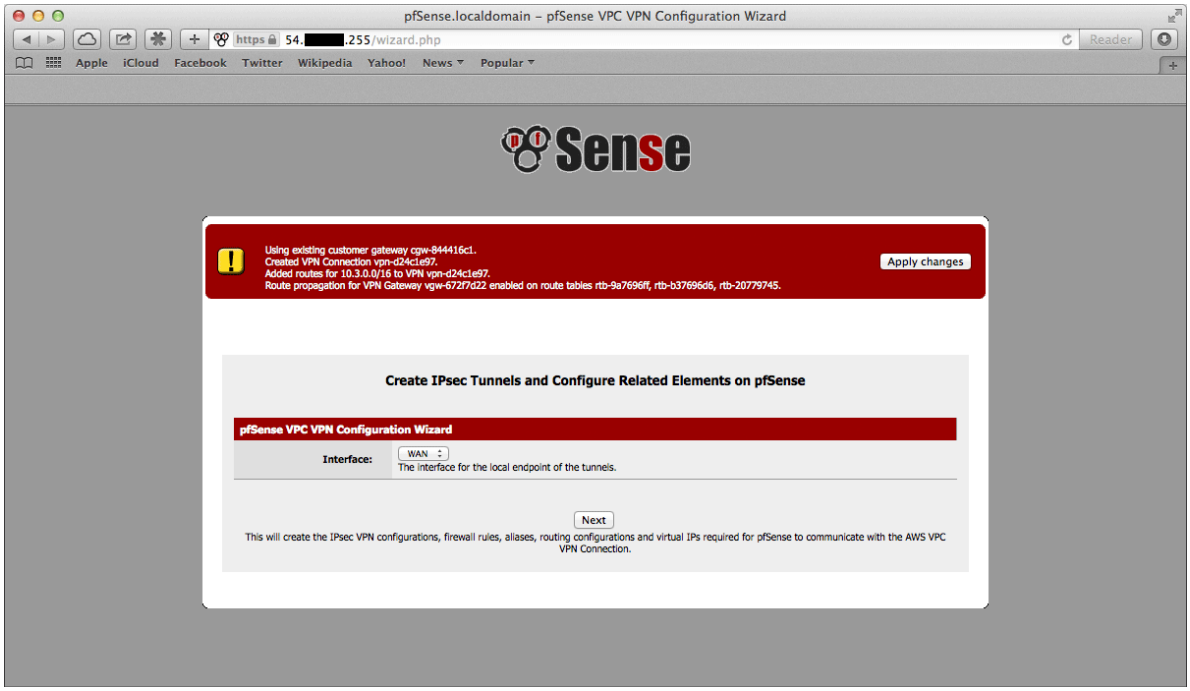
The wizard will then query the AWS API to see if a VPC Connection already exists that matches the data entered. If one exists, it will be used. Otherwise, one will be created.

If static routing was selected, static routes will be added to the VPN Connection for the Local subnets entered.

Route propagation will be enabled for the VPN Gateway in each of the Route Tables that are associated with the VPC. All of these configurations are carried out in the AWS API, nothing has been changed in the pfSense Plus VPN configurations yet.

> **Warning:** **Important Note on Billing:** Once this step is carried out and the VPN Connection is created, AWS will start billing the AWS account the hourly rate for a VPN Connection. This is $0.05 as of this writing, and that is a charge that goes entirely to AWS itself. They will do this until the VPN Connection is deleted via the **VPC Management Console**.
>
> Nothing in pfSense Plus will ever cause AWS to stop billing for this VPN Connection. Whether it works or not, whether the pfSense Plus instance is up or down, whether the IPsec tunnels have been deleted or reconfigured, AWS will continue to bill the hourly fee for a VPN Connection if the creation of it succeeds until it is deleted through their web interface.
>
> The wizard helps establish an initial configuration that works and configures the appropriate elements in AWS's API to facilitate this.
>
> The account holder is responsible for understanding billable charges and disabling any functions, including VPN Connections, that are no longer necessary.

6. If the operations of the previous step succeeded, the next screen will appear.



Select an **Interface** to act as the local endpoint of the VPN tunnels that will be created. In most cases, this should

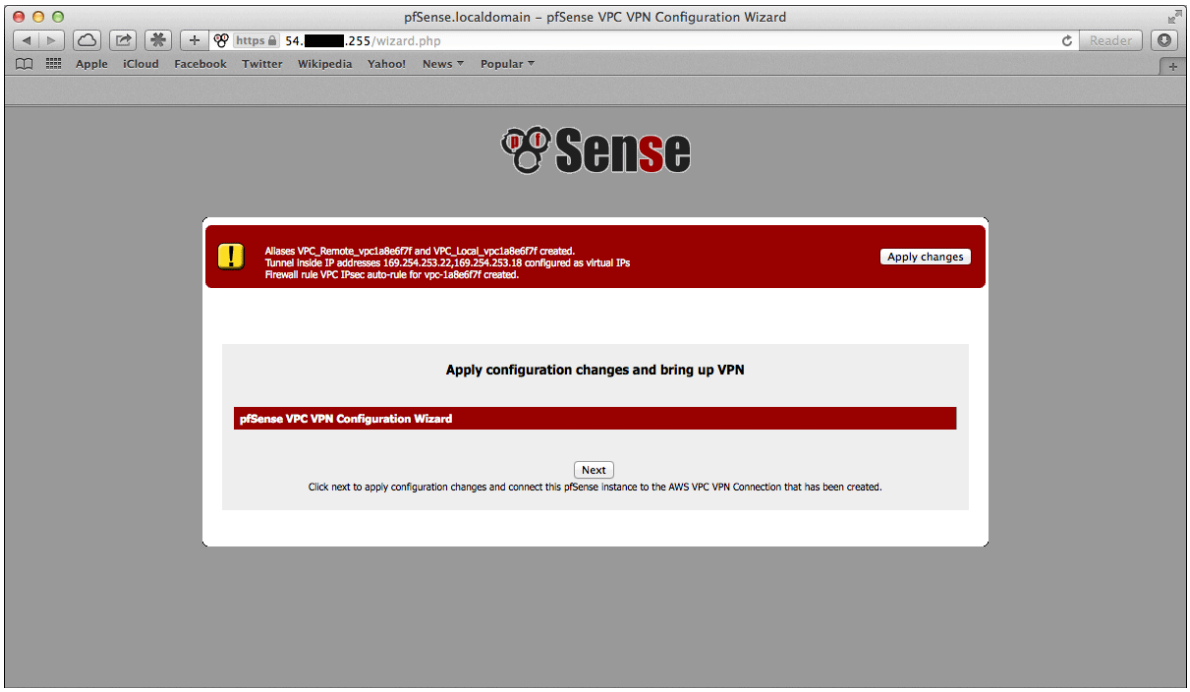
be the WAN interface. It should generally be whatever interface is associated with the Local Public IP Address entered in the previous step.

On an AWS Netgate pfSense Plus instance, this will be whatever interface the Elastic IP is associated with. On a hardware device running pfSense Plus that has the Local Public IP Address directly configured on an interface, this will be the interface that the Local Public IP Address is configured on.

After clicking **Next**, the wizard will configure the VPN and associated settings within pfSense Plus itself using data returned by the AWS API in the previous step.

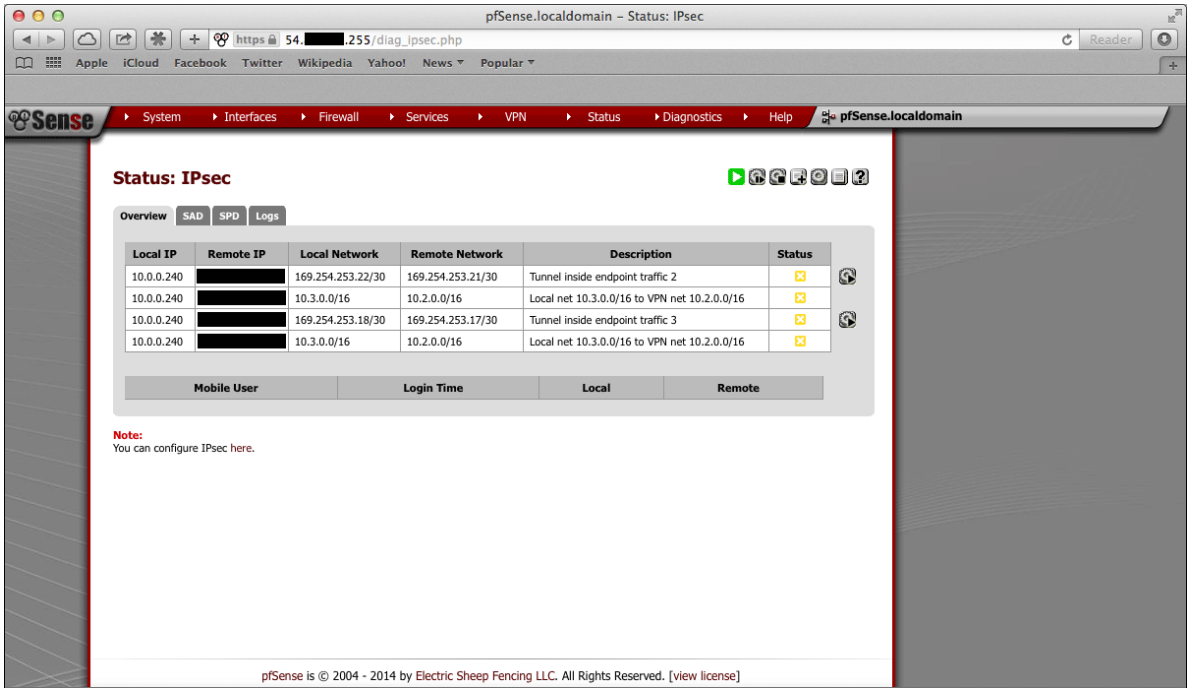
It will configure 2 IPsec tunnels, a firewall rule, 2 Aliases (referenced by the firewall rule), and 2 Virtual IP Addresses. If **BGP** was selected as the **Routing Type** in the previous step, it will install the **FRR** package automatically and configure it appropriately.

7. The next screen will appear and prompt to apply the configuration changes that have been made.



After clicking **Next**, all the configuration changes that were made will be applied.

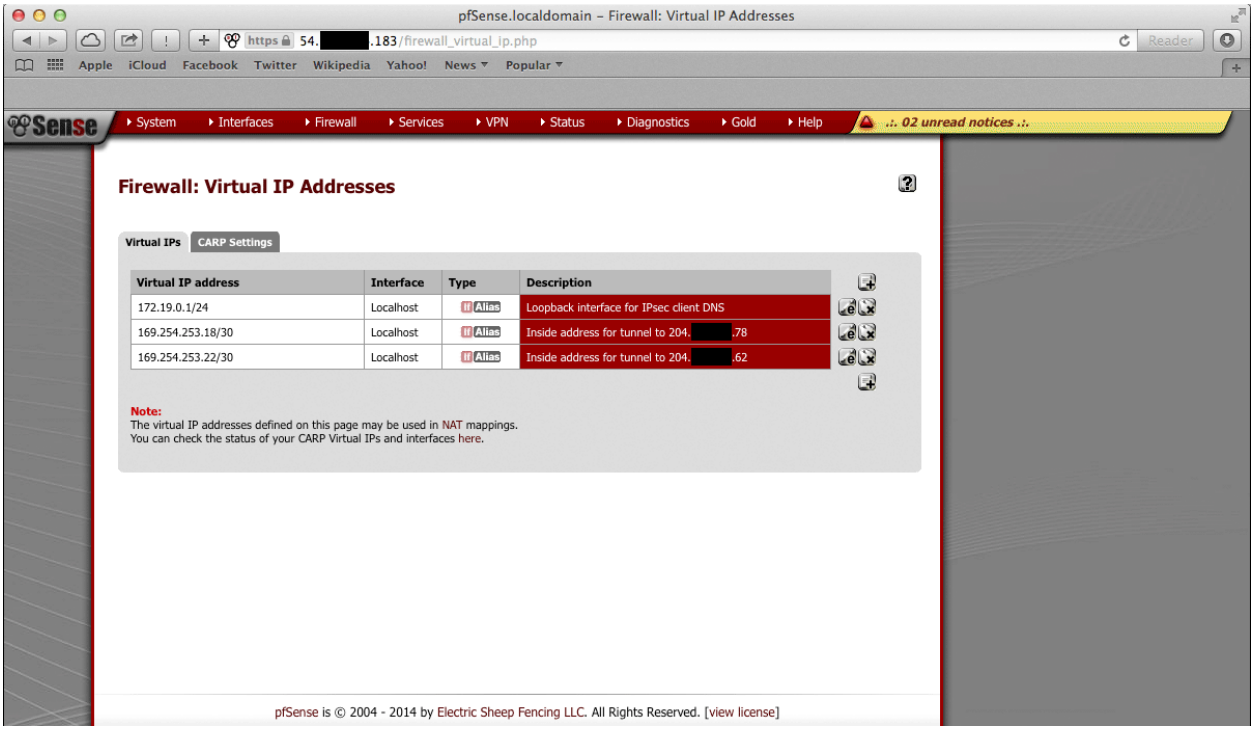
8. The wizard will be completed, and the browser will be redirected to the IPsec status page. The VPN to the VPC should now be fully configured.



Note: Sometimes there is a delay of 5-10 minutes before the tunnels are fully functional and passing traffic. This has been observed particularly often during the setup of tunnels using BGP routing.

30.6.3 Testing Connectivity

Verify that the IPsec tunnels are functioning by attempting to ping the "inside tunnel addresses" of the VPC side of the tunnel by navigating to **Firewall > Virtual IP**. There should be two virtual IP addresses configured that have Descriptions like "Inside address for tunnel to <remote IP address>".



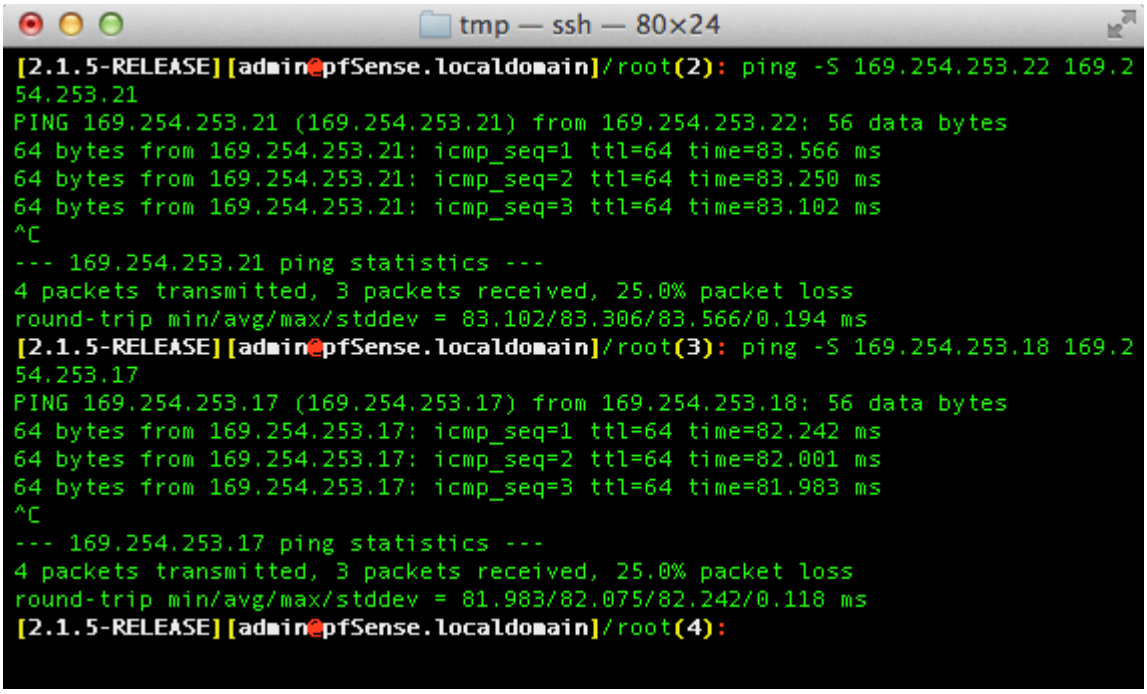
Amazon provides inside addresses for each end of the tunnel in a `/30` subnet in IPv4 link local address space (`169.254.x.y`). Typically, the first usable address in the `/30` is the inside address for the VPC end of the tunnel and the other usable address is the inside address for the local end of the tunnel. Ping from one of the virtual IP addresses configured on the pfSense Plus instance to the IP address that is one less (for example, if the virtual IP address were `169.254.253.22`, ping from `169.254.253.22` to `169.254.253.21`), that checks whether the other end of the tunnel is responding and whether the tunnel is functioning properly.

Note: It sometimes takes a few minutes for the tunnels to begin working after the configuration wizard completes.

To ping, log into the pfSense Plus instance via SSH and execute `ping` from a shell prompt. For the previous example of `169.254.253.22`, the proper syntax is:

```
ping -S 169.254.253.22 169.254.253.21
```

When pinging from a shell prompt, it is possible to leave the command running indefinitely and interrupt it there is a response.



It is also possible to ping via the **Diagnostics > Ping** page by selecting the appropriate source address and entering the remote tunnel inside address in the **Host** field. This will only send a limited number of ping packets, so it may be necessary to repeat this a few times.

30.6.4 VPC Configuration Details

The AWS documentation for connecting a hardware device to a VPC provides a great amount of detail on configuring VPN Connections to VPC.

The main configuration elements that exist in Amazon's data model are the Virtual Private Gateway, the Customer Gateway, and the VPN Connection. In order to configure a VPN connection to a VPC, all 3 of these need to exist. The Virtual Private Gateway is the VPN endpoint on Amazon's side. The customer gateway is the VPN endpoint on the Netgate® pfSense Plus instance/device being configured. The VPN Connection is the IPsec VPN between Amazon and the pfSense Plus instance/device.

A Virtual Private Gateway needs to exist and be associated to the VPC. A particular VPN Gateway can only be associated to one VPC at a time. Once the VPC to connect to has been selected, the wizard invokes the AWS API call **DescribeVpnGateways** to determine if a VPN Gateway already exists that is attached to the VPC. If none exists, it creates one with the **CreateVpnGateway** API call and attaches it to the VPC with the **AttachVpnGateway** call.

A Customer Gateway needs to be created for the public IP address of the device or virtual machine that will be used to connect to the VPC. The wizard invokes the AWS API call **DescribeCustomerGateways** to determine if a Customer Gateway already exists. If none exists, one is created with the **CreateCustomerGateway** API call.

The VPN Connection connects the Virtual Private Gateway and Customer Gateway. The wizard checks to see if there is an existing VPN Connection configured that connects those endpoints by invoking the AWS API call **DescribeVpnConnections**. If none exists, one is created using the **CreateVpnConnection** call. One of the fields returned by this call is a block of XML configuration data that contains configuration data assigned by AWS for use with configuring the VPN Connection. This data is stored and used in subsequent steps to make the required configuration changes within pfSense Plus.

Any objects created through API calls by the wizard will be tagged with names like `auto-created by pfSense VPC <your_VPC_ID>`. This is accomplished by calling the AWS API call **CreateTags** and using `Name` as the key for the

tag.

In addition to creation of the items mentioned above, required adjustments are made to Security Groups and Route Tables to facilitate communication over the VPN. The Security Groups associated with the VPC are updated to allow inbound access from the local subnets on the pfSense Plus end of the VPN. They are checked first via the **DescribeSecurityGroups** AWS API call to determine if the access is already allowed. Any of the subnets that is not already allowed has inbound access added via the **AuthorizeSecurityGroupIngress** AWS API call.

Route Tables associated with the VPC are updated to receive routes from the VPN Gateway used by the VPN Connection. They are checked first via the **DescribeRouteTables** AWS API call. If the VPN Gateway ID is not included in the list of VPN Gateways propagating routes, route propagation for the VPN Gateway is enabled on that table using the **EnableVgwRoutePropagation** AWS API call.

For VPN Connections using static routing, static routes for the specified subnets are added to the VPN Connection. This is done via the **CreateVpnConnectionRoute** AWS API call.

30.6.5 pfSense Plus Software Configuration Details

On the pfSense Plus software side, the wizard adds several items to configure the VPN to the VPC.

Aliases

First, it creates aliases for use in a firewall rule. These aliases are intended to contain the subnets that traffic should be allowed to ingress over the IPsec tunnel. One alias represents the local subnets on the pfSense Plus side and is given a name like **VPC_Local_vpc_12345678** and the other represents the remote subnets on the VPC side and is given a name like **VPC_Remote_vpc_12345678**.

Virtual IP addresses

Next, it adds virtual IP addresses on the `lo0` (loopback) interface. These virtual IP addresses are the local "inside addresses" of the IPsec tunnels. These addresses are used as the local address for BGP communication when BGP routing is selected. These addresses are IPv4 link local addresses (see RFC 3297). AWS assigns `/30` subnets out of the network `169.254.0.0/16` for this purpose.

Note: These addresses are also useful as a ping target to execute a basic test of whether the tunnel is functioning properly. Executing a ping from a source address of one of these IP addresses to the corresponding inside address of the other end of the tunnel helps determine whether the tunnel negotiation is completing properly.

Firewall rules

Next, it adds a firewall rule on the IPsec interface which allows traffic from the VPC networks to the local subnets. This rule uses the previously created Aliases as source/destination targets.

IPsec

Then, it sets up IPsec phase 1 and phase 2 associations. Most of the required settings are extracted from a block of XML data returned by the **CreateVpnConnection** call during the AWS configuration step. This includes parameters such as the endpoint IP addresses, encryption ciphers, timer values, and so on.

If BGP routing was selected, the wizard creates configuration entries for the FRR BGP daemon. The required settings are determined using the AS number entered into the wizard and the parameters returned by the **CreateVpnConnection** call made during the AWS configuration step.

30.6.6 Tested Configurations

Various topologies may be possible to establish using the AWS VPC Configuration Wizard. This section enumerates some of the configurations that were successfully tested.

Platforms

- Various hardware platforms.
- 64-bit virtual machines in VMware vSphere/ESX.
- Amazon EC2 instances (Xen virtual machines) of the Netgate® pfSense Plus Router/Firewall/VPN AWS AMI.

Local network/routing configurations

- pfSense Plus with a public address configured on the WAN interface.
- pfSense Plus with a private address configured on the WAN interface behind a 1:1 NAT.
- pfSense Plus with a private address configured on the WAN interface behind a PAT (1:many NAT).

VPC Topologies

- pfSense Plus connected to a single VPC.
- pfSense Plus connected to multiple VPCs in different regions.
- Amazon EC2 instance of the Netgate pfSense Plus Router/Firewall/VPN AWS AMI connected to a VPC belonging to the same AWS account in a different region.
- Amazon EC2 instance of the Netgate pfSense Plus Router/Firewall/VPN AWS AMI connected to a VPC belonging to a different AWS account in the same region.

The configuration recommended for the greatest amount of stability is to have a public IPv4 address directly configured on the WAN interface of the firewall, but VPNs have been successfully established under all the conditions listed above.

Whether any of these solutions is appropriate should be evaluated in the context of personal needs and existing infrastructure. Other configurations not listed above may be possible as well.

30.6.7 AWS VPC Wizard FAQ

1. What level of redundancy is provided by the two tunnels?

Amazon provides two tunnel endpoints that will allow traffic to be sent between local networks and the remote VPC to which the firewall connects. The IPsec daemon in pfSense Plus is only capable of establishing an active phase 2 association for a particular source/destination pair on a single tunnel.

Phase 2 associations between the local subnets and the remote VPC subnet are configured in the pfSense Plus software GUI for both tunnels, but IPsec will only actually establish an association for the first tunnel. This means that the IPsec daemon will only ever try to send traffic destined for the remote VPC subnet over the first tunnel.

If that tunnel goes down, the second tunnel may be up and inbound traffic from the remote VPC may be sent to the local networks over that tunnel automatically. But outbound traffic to the remote VPC would not automatically fail over to the second tunnel. To send outbound traffic over the second tunnel, disable the phase 2 associations for the first tunnel and apply the changes.

2. I quit the wizard before finishing. Now what?

To finish setting up the VPN, go back to the wizard and run through it again. It should reuse any partial configurations that were generated before it was stopped and create the new elements that are required.

3. What are the AWS charges for this?

AWS determines their own pricing and provides details for EC2 pricing and VPC pricing. There are many types of charges that may be incurred for operating instances on AWS (e.g. charges related to running an instance, bandwidth, storage, elastic IPs, etc.).

The charge of specific interest in this case is the hourly charge for a VPN Connection. As of this writing, it costs $0.05 (USD) per hour in most regions to have a VPN Connection configured and available. AWS will charge whether the VPN Connection is being used or not as long as it is configured. This will be configured by the third step of the wizard and will never be removed by pfSense Plus software.

If the VPN Connection is no longer needed and billing for it needs to be stopped, visit the AWS **VPC Management Console** and delete the VPN Connection manually.

4. Can I use the wizard to connect to the GovCloud region?

This hasn't been officially tested, but at least one user has reported that they were able to successfully connect to the GovCloud region. They manually added the region **us-gov-west-1** to the list of regions in the first step of the wizard and were able to successfully connect to their VPC in that region. This may be possible in a future build, but to try without official support, do the following:

1. Under the **System > Advanced** menu, make sure the **Enable Secure Shell** box is checked. This is already done by default on AWS instances, but is off by default on Netgate hardware devices with pfSense Plus software.
2. Log into the instance via SSH.
3. Make sure the root filesystem is mounted as read/write. On an AWS instance or a hardware device running on an SSD, this should be true. On a hardware device using Compact Flash or an SD card for storage, it will probably be necessary to remount the root filesystem in read/write mode by running:

```
mount -uw /
```

4. Edit the file `/usr/local/www/wizards/vpc_vpn_wizard.xml` using `vi`. Look for a section of the file that looks like this:

```
<option>
    <name>sa-east-1</name>
    <value>sa-east-1</value>
</option>
```

That should appear directly after several similar `<option>` specifications containing all the other available regions. Underneath that section, add the following:

```
<option>
    <name>us-gov-west-1</name>
    <value>us-gov-west-1</value>
</option>
```

Then save the file and exit vi.

5. If the filesystem had to be remounted in read/write mode earlier, remount it in read-only mode by running:

```
mount -ur /
```

The GovCloud region should now appear as a choice in the first step of the wizard.

30.7 Backup Files and Directories with the Backup Package

The **Backup** package allows any given set of files/folders on the system to be backed up and restored. For most, this is not necessary, but it can be useful for backing up RRD data or for packages that may have customized files that are not kept in config.xml.

To install the package:

- Navigate to **System > Packages**
- Locate **Backup** in the list
- Click **Install** at the end of its entry
- Click **Confirm** to begin the installation

Once installed, the package is available at **Diagnostics > Backup Files/Dir**. It is fairly simple to use, as shown in the following example.

30.7.1 Backing up RRD Data

Using this Backup package it is quite easy to make a backup of RRD graph data outside the config.xml method.

See also:

Monitoring Graphs

- Navigate to **Diagnostics > Backup Files/Dir**
- Click **Add** to add a new location to the backup set
- Enter `RRD Files` in the **Name** field

- Enter `/var/db/rrd` in the **Path** field
- Set **Enabled** to *True*
- Enter `RRD Graph Data Files` in the **Description**
- Click **Save**
- Click the **Backup** button to download the backup archive, which contains the configured files and directories for the backup set.
- Save the file in a safe location and consider keeping multiple copies if the data is important.

30.7.2 Restoring RRD Data

- Navigate to **Diagnostics > Backup Files/Dir**
- Click **Browse**
- Locate and select the backup archive file downloaded previously
- Click **Upload** to restore the files

For this example, because the RRD files are only touched when updated once every 60 seconds, it is not necessary to reboot or restart any services once the files are restored.

30.8 Cache / Proxy

Proxies are intermediaries that sit between clients and servers. A client connects to a proxy, and then the proxy decides if the client can receive content from a server. If so, the proxy makes its own connection to the server and then passes back data to the client.

There are two major types of proxies:

Forward Proxy
: Typically sits between local clients and remote Internet servers. It can be used to control which websites that clients are allowed to load, or log servers and URLs clients are visiting. These mostly work with HTTP, but in special cases can also work with HTTPS.

Reverse Proxy
: Typically sits between remote clients and local servers. These allow for load balancing, failover, or other intelligent connection routing for public services such as web servers.

30.8.1 Squid

Danger: The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

Squid is primarily a forward proxy used for client access control. It can, however, be used in a reverse proxy role if needed. The reverse proxy capabilities are inferior to HAProxy, however.

Tuning the Squid Package

Danger: The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

Performance Tweaks

Some users have reported better performance by using the *ufs* cache filesystem setting. When using *ufs* filesystem, `vfs.read_max=32` may be increased to `vfs.read_max=128` in **System > Advanced**, **Sytem Tunables** tab.

Compact swap.state

Squid keeps a cache index journal called `swap.state` in the top level of the squid cache folder, typically `/var/squid/cache/swap.sate`. **This file can grow very large and consume all hard drive space**. To ensure this does not happen, set a **Log Rotate** value in the squid configuration.

By setting a number of days to retain the logs, the squid package will activate a nightly cron job which runs:

```
squid -k rotate
```

Part of this rotation process includes compacting the `swap.state` file, keeping it from getting too large.

If this file is too large and needs to be removed, this may be done while squid is running. After the file is removed, run:

```
squid -k rotate
```

This will cause it to be written out again (but compacted). Alternately, tell squid to perform a clean shutdown with:

```
squid -k shutdown
```

This will also write the `swap.state` file out again, but squid will stop after this and must be restarted, so it is a less desirable option.

If the `swap.state` file is removed while squid is *not* running, it will have to completely rescan the cache folder to rebuild it once squid restarts. This can be a lengthy and time-consuming process. It may be better to remove the contents of the existing cache folder, and rebuild the structure again by running:

```
squid -z
```

See the Squid FAQ entry for more details.

Random Tips/Tricks

> **Warning:** There could be any number of reasons not to do the following things. Be careful, and test any changes.

Caching Windows Updates

> **Warning:** These settings could break things, but when it works, it works beautifully!

If there are a bunch of local PCs that need Windows Updates, but a WSUS server is not an option, squid can cache them.

- On the **General Settings** tab of the squid configuration (**Services > Proxy Server**), place the following patterns recommended by Squid in the **Custom Options** box at the bottom:

```
refresh_pattern -i windowsupdate.com/.*\.
↪(cab|exe|ms[i|u|f|p]|[ap]sf|wm[v|a]|dat|zip|psf) 43200 80% 129600 reload-into-ims
refresh_pattern -i microsoft.com/.*\.
↪(cab|exe|ms[i|u|f|p]|[ap]sf|wm[v|a]|dat|zip|psf) 43200 80% 129600 reload-into-ims
refresh_pattern -i windows.com/.*\.(cab|exe|ms[i|u|f|p]|[ap]sf|wm[v|a]|dat|zip|psf)␣
↪43200 80% 129600 reload-into-ims
refresh_pattern -i microsoft.com.akadns.net/.*\.
↪(cab|exe|ms[i|u|f|p]|[ap]sf|wm[v|a]|dat|zip|psf) 43200 80% 129600 reload-into-ims
refresh_pattern -i deploy.akamaitechnologies.com/.*\.
↪(cab|exe|ms[i|u|f|p]|[ap]sf|wm[v|a]|dat|zip|psf) 43200 80% 129600 reload-into-ims
range_offset_limit none
```

- Click **Save**
- On the **Cache Management** tab of the Squid configuration:
  - Change **Hard Disk Cache size** to something large, say `3000` or `4000` (3GB or 4GB), to accommodate the updates.
  - Change the **Maximum object size** to something big, such as `512000` for 512MB. Going bigger may be needed if any updates larger than that size are released.
- Click **Save**

Caching Mac Updates

```
refresh_pattern ([^.]+.|)(download|adcdownload).(apple.|)com/.*\.(pkg|dmg) 4320 100%␣
↪43200 reload-into-ims
```

Caching AVG and other Updates

As above, but add this before the other options. The same warnings apply:

```
refresh_pattern ([^.]+.|)avg.com/.*\.(bin) 4320 100% 43200 reload-into-ims
refresh_pattern ([^.]+.|)spywareblaster.net/.*\.(dtb) 4320 100% 64800 reload-into-ims
refresh_pattern ([^.]+.|)symantecliveupdate.com/.*\.(zip|exe) 43200 100% 43200 reload-
↪into-ims
refresh_pattern ([^.]+.|)avast.com/.*\.(vpu|vpaa) 4320 100% 43200 reload-into-ims
```

Tweaking Update Caching / Squid seems to download on its own

Change This:

```
range_offset_limit none
```

To:

```
range_offset_limit 0
```

As an alternative, also try:

```
quick_abort_min 0 KB
quick_abort_max 0 KB
```

Or:

```
quick_abort_pct 70
```

To ensure that a file is only downloaded if a user actually receives 70% or more of it. Otherwise, if a user requests a file and then aborts, it will download the whole file.

Note: If `range_offset_limit` is set to `-1` the quick abort options will NOT work

---

Parent proxy

Setting parent proxy available at the **Proxy server: Upstream proxy settings** tab. In most cases, these settings work if the parent proxy also squid.

To use a parent proxy on another server (not squid), it is necessary to disable **Upstream proxy settings**, and use the **Custom options** in the **Proxy server: General settings** tab.

Configuring the SquidGuard Package

Danger: The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

squidGuard is a URL re-director used to integrate blacklists with the Squid proxy software. There are two big advantages to squidGuard: it is fast and free. squidGuard is published under the GNU Public License.

squidGuard can be used to:

- Limit the web access for some users to a list of accepted/well known web servers and/or URLs only.
- Block access to some listed or blacklisted web servers and/or URLs for some users.
- Block access to URLs matching a list of regular expressions or words for some users.
- Enforce the use of domain names/prohibit the use of IP addresses in URLs.
- Redirect blocked URLs to an info page.
- Redirect banners to an empty GIF.
- Have different access rules based on time of day, day of the week, date etc.

Installing Squid and squidGuard

1. From the pfSense® software GUI, navigate to **System > Packages**, **Available Packages** tab
2. Install the **Squid** package if it is not already installed.
3. Install the **squidGuard** package
4. Configure **Squid** package.
5. Configure **squidGuard** package.

Configure the squidGuard Package

Basic configuration

Here describes how to enable and configure squidGuard, and common users access.

1. Open **General settings** tab.
    1. Check the **Enable** box to activate the package.
    2. Set **Blacklist** options to use blacklist categories. (See above, optional)
    3. Click **Save** button.
2. Open **Common ACL** page.
    1. Click **Target Rules List** to show defined blacklists and target categories
        1. Define default user access: select **Default access [all]** as *allow* or *deny*.
        2. Define other category actions:

1. Select **—**, to ignore a category.
2. Select **allow**, to allow this category for clients.
3. Select **deny**, to deny this category for clients.
4. Select **white**, to allow this category without any restrictions. This option is used for exceptions to prohibited categories.

3. To prohibit clients from using IP addresses in URLs, check **Do Not Allow IP Addresses in URL**.
4. Select **Redirect mode**:
   1. *Int error page*: Use the built-in error page. A custom message may be entered in the **Redirect info** box below.
   2. *Int blank page*: Redirect to a blank page
   3. The other options are various redirects to external error pages, and a URL must be entered in the **Redirect info** box if they are chosen.
5. **Use safe search engine**: Protect customers from unwanted search results. It is compatible with *Google, Yandex, Yahoo, MSN, Live Search*. Make sure that these search engines are available. If this protection should be strictly enforced, disable access to all other search engines.

3. After settings are complete, return to the **General Settings** tab and press **Apply**.

Blacklist

Blacklists are optional, but often useful for allowing access to certain types of sites.

squidGuard comes with a small blacklist basically for testing purposes. They should not be used in production. A better way is to start with one of the blacklist collections recommended by squidGuard.

Downloading blacklists:

1. Open **General Settings** tab in squidGuard package GUI, found at **Services > Proxy Filter**.
2. Check **Blacklist** to enable the use of blacklists.
3. Enter blacklist URL in the field **Blacklist URL**.
4. If the firewall is itself behind a proxy, enter the proxy information in **Blacklist proxy** (this step is not necessary for most people).
5. Click **Save**.
6. Navigate to the **Blacklist** tab inside squidGuard.
7. Click the **Download** button.
8. Wait while blacklist is downloaded and prepared to use (10-35 min). Progress will be displayed on that page as the list is downloaded and processed.

How-tos

Exclude domain/URL from blacklist

In the squidGuard GUI (**Services > Proxy Filter**):

1. Open the **Target categories** page
2. Click  to add a new item
3. Enter a name for the category - `myWhitelist` for example.
4. Add domains and/or URLs to the lists as needed. Entries should be separated by a space. The examples on the page show how entries should be formatted.
5. As with the Common ACL discussed previously, redirect and logging options specific to this category may be set.
6. Click **Save**.
7. Open **Common ACL** or **Groups ACL** page (whichever should have an exclusion).
8. Click **Target Rule List** to expand the list of categories. The newly created category should show alphabetically in the list, above any blacklist categories. Find the **MyWhiteList** entry in the list and select **whitelist**.
9. Click **Save**.
10. Return to the **General Settings** tab and press **Apply**.

Block download by Extension

In the squidGuard GUI (**Services > Proxy Filter**):

1. Open the **Target categories** page.
2. Click  to add a new item.
3. Enter a name for the category - `myBlockExt` for example.
4. Add Expressions (for example for ASF, ZIP, EXE, etc. files):

```
(.*\/.*\.(asf|wm|wma|wmv|zip|rar|cab|mp3|avi|mpg|swf|exe|mpeg|mp.|mpv|mp3|wm.|vpu))
```

5. Click **Save**.
6. Open **Common ACL** or **Groups ACL** page (whichever should have an exclusion).
7. Click **Target Rule List** to expand the list of categories. The newly created category should show alphabetically in the list, above any blacklist categories. Find the **myBlockExt** entry in the list and select **deny**.
8. Click **Save**.
9. Return to the **General Settings** tab and press **Apply**.

Troubleshooting

Netflix

If Netflix does not load while squidGuard is active, it is likely because Netflix requires accessing URLs by IP address. Ensure that ACLs matching clients allowed to reach Netflix also have **Do not allow IP-Addresses in URL** unchecked.

Service Does Not Start

If the squidGuard service does not start, there are a few possible explanations:

- On all versions of Squid, if **only** blacklists have been configured, then at startup some important files/directories may not be set properly.
  - Add at least one **Custom Target Category** with a site to pass or block and use it along with the blacklist entries to work around the problem.
- On squid 3.x, the squidGuard service will only start when traffic requires it to run, so it can appear to be stopped even when working properly.
  - Only worry about the service if it appears to not work, don't count on the service status alone.

Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

See also:

- *Troubleshooting the Squid Package*

30.8.2 HAProxy

HAProxy is a powerful reverse proxy that can handle many different types of tasks and scales well for large deployments.

- *HAProxy package*
- *Troubleshooting the HAProxy Package*

30.9 FreeRADIUS package

FreeRADIUS is a free implementation of the RADIUS protocol capable of working with MySQL, PostgreSQL, LDAP, Kerberos.

Refer to the following articles for more information on the listed topics:

30.9.1 Testing the FreeRADIUS Package

Testing the *FreeRADIUS Package* on a firewall running pfSense® software.

Test Configuration

At a minimum, testing FreeRADIUS requires A **User**, an **Interface**, and a **NAS/Client**.

- Add a **User** with the following configuration:

  Username
  `testuser`

  Password
  `testpassword`

- Add a **Client/NAS** with the following configuration:

  IP Address
  `127.0.0.1`

  Shared Secret
  `testing123`

- Add an **Interface** with the following configuration:

  IP Address
  `127.0.0.1`

  Interface Type
  Auth

  Port
  `1812`

GUI Test

The easiest way to test is by using **Diagnostics > Authentication** in the GUI.

First, add a RADIUS server entry to the user manager as described in *Authentication Servers*.

- Navigate to **System > User Manager**, **Authentication Servers** tab
- Fill in the settings to match the entry in FreeRADIUS:

  Descriptive Name
  `FreeRADIUS`

  Type
  RADIUS

  Hostname or IP Address
  `127.0.0.1`

Shared Secret
`testing123`

Services Offered
Authentication

Authentication Port
`1812`

- Click **Save**

Next, perform the GUI test:

- Navigate to **Diagnostics > Authentication**
- Set **Authentication Server** to the RADIUS server in the user manager
- Fill in the **Username** and **Password**
- Click  **Test**

If the test succeeds, the GUI prints a success message:

User testuser authenticated successfully.

The system log will also contain a message indicating a successful login:

```
radiusd[44793]: Login OK: [testuser/testpassword] (from client testing port 0)
```

If the test fails, the GUI prints a failure message:

Authentication failed.

The system log will also contain a message indicating failure:

```
radiusd[44793]: Login incorrect: [testser/testpassword] (from client testing port 0)
```

CLI Test

FreeRADIUS offers an easy-to-use command line tool to check if the server is running and listening to incoming requests.

SSH to the firewall, start a shell, and type in the following command:

```
radtest testuser testpassword 127.0.0.1:1812 0 testing123
```

The following output will appear if the test succeeds:

```
: radtest testuser testpassword 127.0.0.1:1812 10 testing123
Sending Access-Request of id 1 to 127.0.0.1 port 1812
       User-Name = "testuser"
       User-Password = "testpassword"
       NAS-IP-Address = 192.168.0.22
       NAS-Port = 10
       Message-Authenticator = 0x00000000000000000000000000000000
rad_recv: Access-Accept packet from host 127.0.0.1 port 1812, id=1, length=20
```

The `Access-Accept` portion of the output is the most relevant.

Check the system log for the following output:

```
radiusd[44793]: Login OK: [testuser/testpassword] (from client testing port 10)
```

If a part of the test fails, such as incorrect username, then the test command output will look like the following:

```
: radtest testser testpassword 127.0.0.1:1812 10 testing123
Sending Access-Request of id 104 to 127.0.0.1 port 1812
       User-Name = "testser"
       User-Password = "testpassword"
       NAS-IP-Address = 192.168.0.22
       NAS-Port = 10
       Message-Authenticator = 0x00000000000000000000000000000000
rad_recv: Access-Reject packet from host 127.0.0.1 port 1812, id=104, length=20
```

The `Accesss-Reject` packet indicates that the server rejected the attempt, and the system log will contain the following output:

```
radiusd[44793]: Login incorrect: [testser/testpassword] (from client testing port 10)
```

See also:

- *Using Mobile One-Time Passwords with FreeRADIUS*
- *Using EAP and PEAP with FreeRADIUS*
- *Zone Configuration Options*
- *Authenticating OpenVPN Users with FreeRADIUS*

30.9.2 Features

- Authentication with Captive-Portal
- Pre-defined user attributes and custom check-items and reply-items
- NAS/Clients running on IPv4 and IPv6
- Interfaces can listen on IPv4 and IPv6
- OpenVPN + Username + RADIUS and OpenVPN + Username + Cert + RADIUS
- Authentication with PAP, CHAP, MS-CHAP, MS-CHAPv2
- Authentication with EAP-MD5 + dynamic VLAN assignment
- Authentication with PEAP + dynamic VLAN assignment
- Authentication with EAP-TLS/EAP-TTLS + dynamic VLAN assignment

```
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense
   port 0 cli 00-04-23-5C-9D-19)
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense
   port 0 cli 00-04-23-5C-9D-19)
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense
   port 0 via TLS tunnel)
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense
   port 0 via TLS tunnel)
```

- Simultaneous-Use - The following will be present in the system log

```
radiusd[3206]: Multiple logins (max 1) : [testuser/testpw] (from client testing␣
↪port 10)
```

- A certain amount of time per day/week/month/forever (`CHECK-ITEM: Max-Daily-Session := 60`) The user will be disconnected and cannot re-login after the amount of time is reached:

```
radiusd[3206]: Invalid user (rlm_counter: Maximum daily usage time reached):
   [testuser/<via Auth-Type = EAP>] (from client pfsense port 0 cli 00-04-23-5C-9D-
↪19)
```

- A certain amount of traffic per day/week/month/forever. The user will be disconnected and cannot re-login after the amount of traffic is reached. The syslog output looks like this:

```
root: FreeRADIUS: Used amount of daily upload and download traffic by testuser is 0␣
↪of
   100 MB! The user was accepted!!!
root: FreeRADIUS: Credentials are probably correct but the user testuser has␣
↪reached the
   daily amount of upload and download traffic which is 243 of 100 MB! The user was␣
↪rejected!!!
```

- MySQL
- LDAP/Active Directory (connecting to MS AD with PAP)
- One-Time-Password

30.9.3 Installation and Configuration

- Navigate to **System > Packages**, **Available Packages** tab.
- Click  at the end of the row for **freeradius3**.
- Confirm the installation.
- Monitor the progress as it installs.

After Installation, the service may be configured at **Services > FreeRADIUS**.

- Configure the **Interface(s)** on which the RADIUS server should listen.
- Configure the **NAS / Client(s)** from which the RADIUS server should accept packets.
- Add the **User(s)** who should have access.

After this, take a look at the system log. There should be the following:

```
radiusd[16634]: Ready to process requests.
radiusd[16627]: Loaded virtual server
```

30.9.4 Troubleshooting RADIUS Authentication

When attempting to authenticate against a RADIUS server, errors may be encountered that prevent it from working properly. Here are some errors found in the logs and how to resolve them:

```
mpd: [pt0] RADIUS: RadiusSendRequest: rad_init_send_request failed: -1
```

- This appears to happen when the RADIUS shared secret contains special characters. Try again with an alphanumeric shared secret.

30.9.5 Get FreeRADIUS Status Server Updates

The status server provides detailed information about the FreeRADIUS server. The status data includes Accounting-Packets, dropped packets and much more.

To enable status server:

- Set up an interface with **Interface-Type**: *status* and a free port.

  The default port for RADIUS status is `18121`.

- Set up a NAS/Client with **IP-Address**: `127.0.0.1` and a password if one does not already exist.

  This example uses a password of `testing123`.

To request information from the status server:

- SSH to the firewall and enter the following command on the command line:

```
$ echo "Message-Authenticator = 0x00, FreeRADIUS-Statistics-Type = All" | \
  radclient -x localhost:18121 status testing123
```

The output should look like this:

```
Received response ID 223, code 3, length = 140
       FreeRADIUS-Total-Access-Requests = 1
       FreeRADIUS-Total-Access-Accepts = 0
       FreeRADIUS-Total-Access-Rejects = 14
       FreeRADIUS-Total-Access-Challenges = 0
       FreeRADIUS-Total-Auth-Responses = 14
       FreeRADIUS-Total-Auth-Duplicate-Requests = 0
       FreeRADIUS-Total-Auth-Malformed-Requests = 0
       FreeRADIUS-Total-Auth-Invalid-Requests = 0
       FreeRADIUS-Total-Auth-Dropped-Requests = 0
       FreeRADIUS-Total-Auth-Unknown-Types = 0
       [...]
```

To request specific subsets of the status data, replace `All` in `FreeRADIUS-Statistics-Type = All` from the command above with another valid name or value.

A few common names and values are:

- Authentication packets: `Authentication` / `1`
- Accounting packets: `Accounting` / `2`
- Internal server statistics: `Internal` / `0x10`

Tip: The status server accepts either the name **or** its corresponding value a parameter.

More name/value pairs for `FreeRADIUS-Statistics-Type` are listed in the FreeRADIUS dictionary file on the firewall:

```
/usr/local/share/freeradius/dictionary.freeradius
```

30.10 FRR Package

The FRR package manages dynamic routing for the firewall. Dynamic routing refers to routes that are capable of changing, generally due to routing protocols exchanging routing information with neighboring routers.

Unlike static routes, dynamic routing does not require remote network destinations and gateways to be hardcoded in the configuration. Routes and gateways are automatically determined by the protocol instead.

The FRR package is capable of using the following dynamic routing protocols:

Border Gateway Protocol (BGP)
: BGP routes between autonomous systems, connecting to defined neighbors to exchange routing and path information. BGP can use IPv4 and IPv6.

Open Shortest Path First v2 (OSPF)
: OSPF is a link-state routing protocol capable of automatically locating neighboring IPv4 routers within an autonomous system, typically with multicast, and exchanges routing information for networks reachable through each neighbor. OSPF v2 can only use IPv4.

Open Shortest Path First v3 (OSPF6)
: Similar to OSPF v2, but only for IPv6 networks.

30.10.1 FRR Global Settings

The **Global Settings** area in FRR defines options which control the behavior of FRR in general, options which can be utilized by multiple routing protocols, and options which are not specific to a routing protocol.

These options include **Access Lists**, **Prefix Lists**, and **Route Maps**. These mechanisms allow for fine-tuning dynamic routing behavior.

The options here primarily govern the behavior of the `zebra` and `staticd` daemons in FRR.

FRR Global Settings Configuration

Configuration of FRR Global Settings is performed at **Services > FRR Global/Zebra**, on the **Global Settings** tab in FRR. From other areas of FRR, these settings can also be reached by the **[Global Settings]** tab.

General Options

Enable
Master enable option for FRR. When unchecked, all of FRR is disabled, including individual routing daemons.

Default Router ID
The **Router ID** is an IPv4-address formatted string which uniquely identifies this router. Typically, it is set to the LAN address of a router or another unique interface address.

This default router ID is used by FRR when a per-protocol router ID is not set.

Master Password
The password used by FRR for accessing the management daemons internally. This is not typically used by humans interacting with the daemons in the shell (e.g. via `vtysh`) but only internally in FRR.

Encrypt Password
When set, encrypts passwords in output from FRR.

Ignore IPsec Restart
When checked, IPsec restarts cause no action to be taken by FRR. When unchecked, IPsec VTI interfaces will be reset in FRR when IPsec restarts. This reset can prevent routes from becoming inactive in the routing table after IPsec VTI interface events.

CARP Status IP
Used to determine the CARP status when using FRR with certain high availability setups. When the selected CARP VHID is in BACKUP status, FRR will not be started. This check is also made when a CARP VIP transitions to a new status, and the FRR daemons will be stopped or started appropriately to match the VIP status.

Logging

The dynamic routing manager daemon can send log messages to a file, via syslog, or both.

Syslog Logging
Instructs FRR to send its log messages to syslog.

Package Logging Level
Controls the verboseness of FRR package scripts

Normal
Typical log messages.

Extended
Detailed log messages, which may include debugging information.

Modules

The **Enable SNMP AgentX** option in this section controls whether data from FRR will be available through the NET-SNMP package.

Note: This feature is not compatible with the `bsnmp` daemon included with the firewall, only the NET-SNMP package.

Route Handling

The options in this section influence FRR global routing behavior. For example, it can set up special automatic lists to control route acceptance and also to set up FRR-based static routes (e.g. `staticd`). These are different from static routes managed in the firewall GUI directly (*Static Routes*).

Do Not Accept
When set, routes matching the **Subnet** exactly will not be accepted from routing protocols.

Null Route
When set, traffic from hosts inside the defined **Subnet** will never be routed. The traffic will be dropped when it arrives at the firewall.

This option takes precedence over other routing options.

Subnet
An IPv4 subnet or IPv6 prefix for this entry.

Static Route Target
A list of available system gateways, BGP neighbors, and interfaces which can be used as a destination for this route entry. Selecting an entry from the drop-down turns this entry into an FRR static route.

Force Service Restart

By default, FRR attempts to stay running and enact changes in a dynamic way so that there is no loss of service when possible.

Certain changes may necessitate a full restart of FRR, which can be done with the **Force Service Restart** button in this section.

Access Lists

Access list entries determine if networks are allowed or denied in specific contexts used in various routing daemons. For example, an access list may be used to determine if a route is accepted or rejected, or for limiting routes distributed to neighbors.

Access lists are managed on the **Access Lists** tab under **Services > FRR Global/Zebra**.

Access List Configuration

To create a new access list, click  **Add** from the **Access Lists** tab.

The top section of the page sets data about the access list itself:

Type
The type of access list, can be one of:

Standard
A standard access list can match source addresses only.

Extended
An extended access list can match source or destination addresses.

Zebra
A Zebra access list is similar to an **Extended** list, but can contained IPv6 entries.

IP Version
The IP version to match using this access list, either `IPv4` or `IPv6`.

Name
The name of this access list, which will be visible in drop-down lists throughout FRR where access lists can be selected.

The allowed names depend upon the chosen type, and are limited to:

- `1-99` or `1300-1999` for standard access lists.
- `100-199` or `2000-2699` for extended access lists.
- Text names for zebra access lists.

Description
A text comment to describe this access list.

Access List Entries

The **Access list entries** list contains rules which govern the behavior of the list. An access list can have multiple rules.

To add more entries to the list, click  **Add**.

Sequence
The order of entries inside access lists is important, and the order is determined by this sequence number.

Each rule in an access list must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use `10`, `20`, `30`, rather than `1`, `2`, `3`.

> **Warning:** The order of rules displayed in the GUI may be different from the order set by the **sequence** numbers. The **sequence** number order is the true order in which rules are evaluated.

Action
The action to take for this rule, either `permit` or `deny`.

Source Network
The source IP prefix to match for this rule, given in network/prefix notation. For example, `192.168.0.0/16`.

Source Any
When set, the **Source Network** is ignored and any source will match the rule.

Destination Network
The destination IP prefix to match for this rule, given in network/prefix notation. For example, `192.168.0.0/16`.

Destination Any
When set, the **Destination Network** is ignored and any destination will match the rule.

Exact
Will only match if a network prefix matches exactly, rather than matching networks contained within the specified prefix.

Prefix Lists

Prefix List entries determine parts of networks which can be allowed or denied in specific contexts used in routing daemons. For example, a prefix list may be used to match specific routes in a route map.

Prefix lists are managed on the **Prefix Lists** tab under **Services > FRR Global/Zebra**.

Prefix List Configuration

To create a new prefix list, click  **Add** from the **Prefix Lists** tab.

The top section of the page sets data about the prefix list itself:

IP Type
The IP version to match using this access list, either `IPv4` or `IPv6`.

Name
The name of this prefix list, which will be visible in drop-down lists throughout FRR where prefix lists can be selected.

Description
A text comment to describe this prefix list.

Prefix List Entries

The **Prefix list entries** list contains rules which govern the behavior of the list. A prefix list can have multiple rules.

To add more entries to the list, click  **Add**.

Sequence
The order of entries inside prefix lists is important, and this order is determined by a sequence number.

Each rule in a prefix list must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use `10`, `20`, `30`, rather than `1`, `2`, `3`.

> **Warning:** The order of rules displayed in the GUI may be different from the order set by the **sequence** numbers. The **sequence** number order is the true order in which rules are evaluated.

Action
The action to take for this rule, either `permit` or `deny`.

Network
The network prefix to match. This may optionally be bound by **Minimum Prefix** (lower bound) or **Maximum Prefix** (upper bound) size limit. When no upper or lower bound is set, the prefix will be matched only exactly as given. Setting bounds allows a prefix list to also match more specific routes which are a contained within the given prefix.

Any
When set, matches any prefix.

Minimum Prefix
Also known as `ge`. Sets a lower bound for the prefix length. This must be greater than the prefix length given in **Network**, and less than or equal to the value of **Maximum Prefix**, if present.

Maximum Prefix
Also known as `le`. Sets an upper bound for the prefix length. This must be greater than the prefix length given in **Network**, and greater than or equal to the value of **Minimum Prefix**, if present.

Prefix List Examples

For example, the following prefix list will match any of the RFC1918 networks:

- Sequence: 10, Action: Permit, Network: `10.0.0.0/8`, Maximum Prefix: `32`
- Sequence: 20, Action: Permit, Network: `172.16.0.0/12`, Maximum Prefix: `32`
- Sequence: 30, Action: Permit, Network: `192.168.0.0/16`, Maximum Prefix: `32`

For each of these entries, the prefix list will match based on the bits specified in the prefix. A match will occur for any network included in the specified range. For example, `10.0.0.0/8` with a **Maximum Prefix** of `32` means a route for any smaller network inside `10.0.0.0/8` will also match, so long as the prefix length is less than `32`. So `10.2.0.0/16` will also match this entry, as will `10.34.157.82/32`. Taken as a whole, this prefix list will match not only the list of RFC1918 networks exactly, but any smaller network wholly contained inside.

As another example, consider this rule instead:

- Sequence: 10, Action: Deny, Network: `10.0.0.0/8`, Minimum Prefix: `24`, Maximum Prefix: `32`

This matches routes for networks inside `10.0.0.0/8` with a prefix length greater than or equal to `24` but less than or equal to `32`. Meaning it will **not** match larger networks such as `10.2.0.0/16`, but it will match more specific networks such as `10.2.56.128/29` anywhere inside the `10.0.0.0/8` address space. This type of rule can be used to exclude small prefixes from being matched by a route map, for example.

Dynamic Routing Route Maps

Route maps are a powerful mechanism which can match or set various values for use by routing daemons, especially BGP. A route map can match based on criteria such as those set by *Access Lists* and *Prefix Lists*, among others. Route maps can control, for example, whether specific routes are accepted from neighbors, or whether specific routes are distributed to neighbors. They can also adjust various properties of routes, which largely depends upon the context in which they are used, such as for BGP or OSPF.

Route maps are managed on the **Route Maps** tab under **Services > FRR Global/Zebra**.

Route map entries are complex, and multiple entries can be combined by using the same name on more than one entry, but with different sequence numbers to control the order in which the route map entries are processed by FRR.

Route Map Configuration

To create a new route map, click  **Add** from the **Route Maps** tab.

The **General Options** section of the page sets data about this route map entry:

Name
The name of this route map entry.

Note: The same name can be used for multiple entries, but each entry using the same name must use a unique sequence number.

Description
A text description of this route map

Action
The action taken by this route map, either `permit` or `deny`.

permit
When an entry is matched and permitted, the "set" actions of a route map are carried out, if present, and then *Logic Control* entries, if present, are performed. The route will be allowed unless the control flow ultimately prevents that from happening.

deny
When an entry is matched and denied, the route is not allowed.

Sequence
The sequence number of this route map. Must be in the range `1-65535`.

The order of entries inside route maps is important, and this order is determined by a sequence number.

Each entry in a route map must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding entries in the future. For example, use `10`, `20`, `30`, rather than `1`, `2`, `3`.

Warning: The order of entries displayed in the GUI may be different from the order set by the **sequence** numbers. The **sequence** number order is the true order in which rules are evaluated.

Route Map Contents

The remaining sections on the page control what this route map entry will do. There are numerous options available, from control and logic flow, to matching, setting, and altering routes.

Generally speaking, when an option in the remainder of the page is set to *None*, it will be ignored or have no effect.

Due to complexity, these options are broken up until multiple sections.

Logic Control

Call Route Map
Will immediately process the selected route map. If the called route map returns `deny`, then processing is stopped and the route is denied.

Exit Action

next
Proceeds to the next rule in the route-map

<sequence> number
Skips to the rule with the given sequence number in this route map.

Access Lists

Sets an *Access List* used to match this route map entry.

Prefix Lists

Sets a *Prefix List* used to match this route map entry.

Next Hop

Controls operations matching or setting the next hop for a route.

Next Hop Action
Chooses between actions to take for the next hop of a route.

Match Peer
Matches a specific next hop peer (e.g. BGP Neighbor).

Match ACL
Matches based on a specific *Access List*.

Match Prefix List
Matches based on a specific *Prefix List*.

Set (Peer Only)
Changes the next hop on a route to the specified peer (e.g. BGP Neighbor).

Peer
Specifies a peer for **Next Hop Action** when it is set to **Match Peer** or **Set (Peer Only)**.

Local (match only)
Matches a route when its next hop is this firewall.

Unchanged (set only)
Leaves the next hop unchanged.

Peer Address (set only)
For inbound IPv4 routes received from a neighbor, sets the next-hop to the address of the neighbor. For outgoing routes this is the local address used to establish an adjacency with the neighbor.

<Neighbor>
A list of available peers fills out the list. Selecting an entry uses that specific peer to match or set.

ACL
Specifies an *Access List* used to match the next hop value when **Next Hop Action** is set to **Match ACL**.

Prefix List
Specifies an *Prefix List* used to match the next hop value when **Next Hop Action** is set to **Match Prefix List**.

Metric

Match or set the metric of a route.

Metric Action
Chooses between actions to take for the metric of a route.

Match
Matches the given metric value.

Set
Sets the MED value for routes. When this router has multiple links to the same AS, the MED value influences which path the router will prefer. The router will prefer to use links with a lower MED value. Adding a + before the metric value will result in a relative adjustment instead of setting an absolute value.

Set OSPF6 External Type 1 Metric
Similar to above, but only operates on the OSPF6 External Type 1 Metric.

Set OSPF6 External Type 2 Metric
Similar to above, but only operates on the OSPF6 External Type 2 Metric.

Metric Value
The metric value to match or apply. When setting a metric, the value may be `+rtt`, `-rtt`, + or - value offset, or a specific metric.

Weight

Sets the weight of the route to the supplied value. When a remote AS is reachable via multiple paths through other intermediate AS neighbors, the router will prefer to use a higher weight path to reach it.

Local Preference

The options in this section will either match or set the BGP local preference value of a route using the given **Local Preference** value.

BGP AS Paths

Matches or sets a *BGP AS Path*.

AS Path Action

Match AS Path
Match based on the *BGP AS Paths* selected in **Match AS Path** below.

Set Exclude
Excludes the AS numbers specified in **Set AS List** from the path of the route.

Set Prepend
Prepends the AS numbers specified in **Set AS List** to the AS path.

Set Prepend Last-AS
Prepends the last AS the number of times specified in **Set AS List** to the leftmost end of the path.

> **Warning:** Do not select `Set Prepend Last-AS` in an outbound route-map. The `set` statement would be executed *before* the local AS number is prepended to the AS-path.

Match AS Path
The specific *BGP AS Path* to match.

Set AS List
A list of *BGP AS Path* entries to apply to the route.

BGP Communities

Matches or sets *BGP community* values in routes.

Community Action

Match
Match based on community value in **Match Community**.

Match Exact
Match, but only if the community value matches exactly, rather than being part of a list.

Set
Sets the BGP community value to the list in **Set Community**.

Match Community

internet, no-export, no-advertise, local-as
Match one of the well-known communities.

<Community Name>
Match a community defined at *BGP Community Lists*.

Set Community
When setting a community, this is a **space-separated** list of communities in `AS:VAL` format, or a well-known community: `internet`, `no-export`, `no-advertise`, or `local-as`. Can also be set to `none` to remove BGP community values entirely.

Additive
Adds the specified community value to the route without replacing the existing values.

Origin

Origin Action
Match or set based on the origin (source) of the route.

Origin Name

Remote EGP
Routes from Exterior Gateway Protocols (e.g. BGP).

Local IGP
Routes from Interior Gateway Protocols (e.g. OSPF).

Unknown Heritage (Incomplete)
Routes from unknown sources.

Source Protocol

Matched based on a specific route source protocol from a list of possible options.

Note: Not all options in the list are implemented in the FRR package currently.

Tags

Tag Action

Match
Match a tag value set by another route map rule.

Set
Set a tag value to be matched by another route map rule.

Tag Value
The specific tag value to match or set. This value is an integer from `1-4294967295`.

RPKI

Matches based on the RPKI state.

Prefix Not Found
The prefix is not present in the configuration.

Invalid Prefix
The prefix is known but failed validation.

Valid Prefix
The prefix is known and passed validation.

Route Map Examples

This example creates a route map to control which routes will be sent to peers via BGP. The first rule prevents any route from sending if it matches entries in the `RFC1918` prefix list. The second rule allows routes that match networks listed in the `MY-ROUTES` prefix list. This ensures that even if other mechanisms would try to export routes to peers, that no routes to private networks are leaked.

- Name: EBGP-OUT, Sequence: 10, Action: Deny, Match Prefix List: RFC1918
- Name: EBGP-OUT, Sequence: 20, Action: Permit, Match Prefix List: MY-ROUTES

FRR Status

The status page for FRR can be found at **Status > FRR** or on the **Status** tab.

Accessing the **Status** tab from **Global Settings** shows an overview of FRR and basic information from all active routing protocols.

The protocol and daemon-specific tabs (Zebra, BGP, OSPF, OSPF6, BFD) display more detailed information related to each area.

The **Configuration** tab shows the current configuration file (`frr.conf`).

See also:

- *BGP Status*
- *OSPF Status*

Output Limits and Filtering

Most sections on the status page limit the amount of data shown to ensure that the GUI does not get overwhelmed when dealing with large amounts of data, such as from full BGP feeds.

There are two options which govern the output shown in each area of the status page:

Display XX of YY items
This option controls the number of items the GUI will display for that area. There are several values to choose from as well as an **all** selection which displays everything.

> **Warning:** Do not select **all** when using full BGP feeds or other sources of large routing data. Attempting to display that much data will cause the status page to fail, or will drastically reduce performance.

Filter expression
This option filters the output so it only includes lines which match the given expression. This can be a simple string to match or a regular expression pattern.

Zebra Status

The `zebra` daemon orchestrates all the other protocol and related daemons in FRR. As such, it has general information about the operation of FRR in general, plus information which spans across multiple routing protocols.

Zebra Routes
The complete IPv4 routing table known to `zebra`, including notations for routes sent from and received through each protocol.

The first column of output contains codes which are explained in the output header. These codes can also explain which routes are associated with each type of :ref:`dynamicrouting-protocol-lists`.

Note: The output here may differ from the OS routing table visible at **Diagnostics > Routes** as not every route in FRR may be set to be active there.

Example entries:

- `B>* 10.7.0.0/16 [20/0] via 10.6.106.2, ipsec4000, weight 1, 4d01h00m`
  - This route was received from BGP (`B`), selected for use (>) and is in the firewall routing table (FIB, *).
- `O 10.3.0.0/24 [110/10] is directly connected, vmx1, weight 1, 5d15h55m`
  - This route is being advertised to OSPF neighbors (`O`, but directly connected).
- `O>* 10.20.0.0/16 [110/20] via 10.3.111.2, ipsec4000 onlink, weight 1, 4d01h39m`
  - This route was received from OSPF (`O`), selected for use (>) and is in the firewall routing table (FIB, *).

Zebra IPv6 Routes
Same as **Zebra Routes** but for IPv6 routes.

Zebra CPU
CPU usage statistics for the `zebra` daemon and its various threads.

Zebra Interfaces
Information about system interfaces as seen by `zebra`. This may also include information relevant to interface participation in specific routing protocols.

Zebra Memory
Memory usage statistics for various data used by `zebra`.

30.10.2 Border Gateway Protocol

Border Gateway Protocol (BGP) is a dynamic routing protocol used between network hosts. BGP routes between autonomous systems, connecting to defined neighbors to exchange routing information.

BGP can be used for exterior routing (eBGP) or interior routing (iBGP), routing across Internet circuits, private links, or segments of local networks.

BGP Required Information

Before starting, take the time to gather all the information required to form a BGP adjacency to a neighbor. At a minimum, FRR will need to know these items:

Local AS Number
: The autonomous system (AS) number for this firewall. This is typically assigned by an upstream source, an RIR, or mutually agreed upon by internal neighbors.

Local Router ID
: Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

Local Network(s)
: The list of networks that are advertised over BGP as belonging to the Local AS. For external BGP, this is typically the IP address block allocated by the RIR. For internal BGP, this may be a list of local networks or a summarized block.

Neighbor AS Number
: The autonomous system number of the neighbor.

Neighbor IP Address
: The IP address of the neighboring router.

The example in this section uses the following values:

Table 1: Example BGP Configuration

Item	Value
Local AS Number	65014
Local Router ID	10.14.0.1
Local Network(s)	10.14.0.0/16
Neighbor AS Number	65002
Neighbor IP Address	203.0.113.2

BGP Example Configuration

The following example configures a BGP adjacency to a neighbor using the settings from *Example BGP Configuration*.

Assumptions

This example makes a few assumptions for brevity and to keep the example simple, including:

- The remote peer is already configured for BGP with equivalent settings.
- Transit to the peer across a directly attached shared network is already configured, for example over a VPN, shared network segment, or peer-to-peer link.
- Firewall rules pass BGP traffic on TCP port 179 between the peers.

Example Configuration

Route Map for Peer Filtering

Before configuring BGP, add a route map to match any routes so it can be used by FRR to allow exchanging all routes with the peer.

Warning: This basic example replicates previous FRR behavior which allowed any routes to be exchanged with a peer. This is convenient, but not secure. For increased security, create a set of route map entries which ensure that only expected routes are sent and received where possible.

- Navigate to **Services > FRR Global/Zebra**, **Route Maps** tab
- Click  **Add**
- Set the following options:

  Name
  `ALLOW-ALL`

  Description
  Match any route

  Action
  Permit

  Sequence
  100

FRR BGP Configuration

- Navigate to **Services > FRR BGP**
- Set the following options:

  Enable
  Checked

  Local AS
  `65014`

  Router ID
  `10.14.0.1`

  Networks to Distribute
  `10.14.0.0/16`

- Click **Save**
- Navigate to the **Neighbors** tab
- Click  **Add**
- Set the following options:

  Name/Address
  `203.0.113.2`

Remote AS
`65002`

Route Map Filters
Set both **Inbound** and **Outbound** to `ALLOW-ALL`

FRR Global Configuration

- Navigate to the **[Global Settings]** tab
- Set the following options:

  Enable
  Checked

  Master Password
  Create a random string to use

- Click **Save**
- Navigate to the **Status** tab
- Confirm that the BGP neighbor is present, and its routes are in the table

FRR BGP Configuration

The FRR BGP service contains numerous methods to configure and fine-tune BGP routing behavior. Due to this complexity, the topic has been split into several sections. Read through each section before attempting to create a new BGP configuration.

BGP Tab Configuration

BGP Router Options

Enable
Master enable switch for BGP routing. When checked, FRR will start the BGP routing daemon and attempt to use the BGP settings in this section.

Log Adjacency Changes
When set, BGP neighbor adjacency changes will be written via syslog.

Local AS
Required. The autonomous system (AS) number for this firewall. This is typically assigned by an upstream source, an RIR, or mutually agreed upon by internal neighbors.

Router ID
Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

Timers

Keep Alive Interval
Configures the intervals between keep alive messages.

Hold Time
How long to wait for a response before considering the peer unreachable.

Update Delay
Keeps BGP in a read-only mode for the specified time after the daemon restarts or peers are cleared.

Peer Wait
The amount of time to wait for peers to reach an established state. This starts the same time as the **Update Delay** and allows FRR to end the update delay early if peers are available within the given time period.

Disable Default IPv4 Unicast
When unchecked, FRR assumes the peer is capable of IPv4 unicast in all cases, even when the neighbor is connected over IPv6.

Modules

Enable SNMP AgentX
Enable agentx for accessing FRR BGP data via SNMP with the NET-SNMP package.

Enable BGP RPKI
Enable BGP Resource Public Key Infrastructure.

Global Neighbor Shutdown

Global Neighbor Shutdown
When checked, **all** neighbors are placed into an administrative shutdown state.

Message
An optional message sent to BGP peers when in this shutdown state.

Graceful Restart/Shutdown

Disable BGP Graceful Restart
Globally disable graceful restart functionality in both restart and helper mode.

Preserve FW State
If checked, sets the forwarding state (F) bit indication that the FIB is preserved while performing a graceful restart.

Timers

Stale Path Time
The time (in seconds) FRR will to retain stale paths from a restarting peer.

Restart Time
The time (in seconds) to wait before deleting stale routes unless a BGP open message is received.

Select Defer Time
The time (in seconds) FRR defers the route selection process after it restarts.

RIB Stale Time
The time (in seconds) stale routes are retained in the RIB.

Enable BGP Graceful Shutdown
When set, BGP graceful shutdown is enabled.

RPKI Timers

Configures timers for *BGP RPKI*.

Polling Period
The time (in seconds) FRR waits until it queries the cache for updated data.

Expire Interval
The time (in seconds) after which FRR will expire RPKI cache data.

Retry Interval
The time (in seconds) at which FRR will retry connecting to an RPKI cache server after a connection failure.

Network Distribution

These options control networks for which FRR will distribute or redistribute routes to peers. Peers will be informed to reach these networks through this router.

Redistribute Option Choices
Each option in this section may be set to one of the following choices:

No
Does not distribute routes from this source.

IPv4
Distributes only IPv4 routes from this source.

IPv6
Distributes only IPv6 routes from this source.

IPv4+IPv6
Distributes both IPv4 and IPv6 routes from this source.

<Route Map Name>
Filters distribution of routes by use of the named *route map*.

Redistribute Local
These networks are considered local to the router.

Redistribute Connected Networks
Redistributes routes for networks which are attached to and present on interfaces of this firewall.

Redistribute FRR Static Routes
Redistributes routes for networks defined as *FRR static routes* in **Global Settings**.

Redistribute Kernel Routing Table
Redistributes routes for other networks found in the kernel routing table. This includes static routes defined in pfSense (*Static Routes*), as well as automatic static routes setup for other purposes.

Redistribute OSPF
These networks can be reached through OSPF neighbors, not directly on this firewall.

Redistribute OSPF Routes to BGP Neighbors
Redistributes routes for networks reachable through IPv4 OSPF neighbors.

Redistribute OSPFv3 Routes to BGP Neighbors
Redistributes routes for networks reachable through IPv6 OSPF6 (OSPFv3) neighbors.

Networks To Distribute
A manual list of networks that are advertised over BGP as belonging to the Local AS. For external BGP, this is typically the IP address block allocated by the RIR. For internal BGP, this may be a list of local networks or a summarized block.

Subnet to Route
An IPv4 subnet or IPv6 prefix to advertise to peers.

Note: If this subnet is not in the routing table (e.g. it is a summary or aggregation) then the **Network Import Check** option in *Advanced BGP Configuration* must be *unchecked*.

Route Map
A route map to apply to messages advertising this network.

BGP Neighbor Configuration

BGP Neighbors are managed at **Services > FRR BGP** on the **Neighbors** tab.

The **Neighbors** tab contains a list of current neighbors, if any, and controls to manage the entries (e.g. edit, delete).

The  **Add** button creates a new neighbor.

The remaining sections on this page cover the various options available when creating or editing a neighbor entry.

General Options

Name/Address
The name of a peer group or IP address of a neighbor.

Enter a text name to define a Peer Group. Enter an IP Address to define a Peer.

Peer groups allow common options to be defined which may then be applied to multiple neighbors without manually placing the options on each neighbor.

Description
A text description about this neighbor.

Peer Group
A list of existing peer groups to which this neighbor can be added. Can only be used when defining a neighbor by IP address.

Password
Sets a password used to secure communication with this neighbor using TCP MD5.

The operating system of the neighbor and its BGP capabilities may restrict which of the password types can be used.

FRR and setkey Bidirectional
Configures the password in FRR and at the operating system level in security policies for both inbound and outbound packets.

This is the best option to use when possible as it fully implements TCP MD5 at the operating system level and in FRR. Use this when both neighbors are fully compatible with TCP MD5 for sending and receiving.

FRR and setkey Outbound
Configures the password in FRR and at the operating system level in security policies, but only in the outbound direction.

This option does not validate TCP MD5 on inbound packets, but will add TCP MD5 information to packets sent to the neighbor. Use this if the neighbor is unable to properly send TCP MD5 which can be validated by this firewall.

FRR Only
Configures the password only in FRR, not at the operating system level.

setkey Only Outbound
Configures the password only at the operating system level in security policies, but only in the outbound direction.

setkey Only Bidirectional
Configures the password at the operating system level in security policies for both inbound and outbound packets, but not in FRR.

Shutdown

Neighbor Administrative Shutdown
When checked, the neighbor will be put into, and kept in, an administrative shutdown state.

Shutdown Message
A text message sent to the neighbor while it is administratively shut down.

Auto-Shutdown

RTT
If the round-trip time to the neighbor exceeds this value it will be automatically shut down.

Keep alive Count
If the neighbor fails to respond to this number of keep alive messages, it will be automatically shut down.

Basic Options

Remote AS
Autonomous System (AS) Number for this neighbor. Can be an integer from `1-4294967295`, `external`, or `internal`.

Update Source
These options control how FRR will communicate with the neighbor.

IP Type
Sets the address family of the IP address to which FRR will bind for communicating to this neighbor, either IPv4 or IPv6.

Local Source
Sets the specific IP address to use when communicating with the neighbor. This can be an interface address, an IP alias VIP, or a CARP VIP.

Address Family
When set, the neighbor is allowed to advertise routes for both IPv4 and IPv6. Otherwise, the type of routes will be restricted to whichever IP type is set for the **Update Source**.

Default Originate
These options control whether FRR will advertise itself as the default route for this neighbor.

Originate Default to Neighbor
Sets the address family for which a default route will be sent to the neighbor, either IPv4, IPv6, or both.

Route Map
A route map used to restrict default origination.

Send Community
Sends the community attribute to this peer, limited to the specified types.

Next Hop Self
Disables next hop calculation for this neighbor and uses the address of this router instead.

Enabled
Uses the address of this router as the next hop in routes announced to this peer if they are learned via eBGP.

Force
When set, also sets the next hop to the address of this router on reflected routes.

Inbound Soft Reconfiguration
Allows the peer to send requests for soft reconfiguration, to apply changes to routes or new attributes without the need for a session reset.

Timers

Keep Alive Interval
Configures the interval between keep alive messages to wait for a response from this neighbor before considering the peer unreachable. This overrides the default values set on the BGP server itself.

Hold Time
Configures how long to wait for a response from this neighbor before considering the peer unreachable. This overrides the default values set on the BGP server itself.

Connect Timer
The amount of time, in seconds from `1-65535`, in which a connection to this peer must be established or else it is considered unsuccessful.

Peer Filtering

These options control which routes may be sent to, or received from, this neighbor.

Note: The current FRR package does not exchange routes with BGP peers by default without being explicitly allowed to do so by a filter. This is secure behavior but requires manually specifying a filter to allow routes to be exchanged.

To replicate the behavior of older FRR versions, add a *route map* to permit all routes (Name: `allow-all`, Action: *Permit*, Sequence: `100`), then set that route map on BGP neighbors for inbound and outbound peer filtering. For increased security, utilize route maps which filter incoming and outgoing routes so they match more strictly.

Distribute List Filter
Defines an *access list* which is used by BGP to filter route updates for this peer, in either the inbound or outbound direction.

Prefix List Filter
Defines a prefix list which is used by BGP to filter route updates for this peer, in either the inbound or outbound direction.

AS Path Filter
Defines an *AS path list* which is used by BGP to filter route updates by AS path in either the inbound or outbound direction.

Route Map Filters
Defines a route map which is used by BGP to filter route updates for this peer, in either the inbound or outbound direction.

Unsuppress Route Map
Configures a route map which BGP can use to unsuppress routes that would otherwise be suppressed by other configuration settings.

BFD

Configures *Bidirectional Forwarding Detection* (BFD) options for this peer.

BFD Enable
Listen for BFD events registered on the same target as this BGP neighbor.

BFD Check Control Plane Failure
Allow FRR to write CBIT independence in outgoing BFD packets. Also allow FRR to read both the CBIT value of BFD and lookup BGP peer status. This option allows BFD to ignore down events during a graceful restart of the remote peer if graceful restarts are enabled in BGP. When enabled, if BFD catches a down event it first checks if the BGP peer has requested that local the BGP daemon keep the remote BGP entries marked as stale. In that case it can safely ignore the event to allow the restart to happen gracefully (RFC 4724).

BFD Peer
Selects a *BFD peer* to associate with this neighbor.

Graceful Restart

Graceful restart mode for this neighbor, may be one of:

Default
Will use the default value for BGP graceful restart from *Graceful Restart/Shutdown*.

Restart
Enables BGP graceful restart functionality for this peer.

Helper
Enables BGP graceful restart helper only functionality for this peer.

Disable
Disables all BGP graceful restart functionality for this peer.

Advanced Options

Weight
Applies the given weight to routes received from this peer.

Passive
When set, this router will not issue open requests to the neighbor on its own. The BGP daemon will only respond to remote open requests from this neighbor.

Path Advertise

All Paths
Advertise all known paths to this peer, instead of only advertising the base path.

Best Path
Advertise only the best known base paths for each AS.

Advertisement Interval
Minimal time (in seconds) between sending BGP routing updates to this neighbor.

Allow AS Inbound
Allows routes to be received from this peer which are from the same AS of this router, but through a different path.

Enabled
Always allow.

Only if Origin
Accept the AS of this router in an AS path if the route originated in the AS of this router.

Allow <number>x
Allowed number of AS occurrences, from `1-10`.

AS Override
Override ASNs in outbound updates to this peer if the AS path is identical to the remote AS.

Attribute Unchanged
Propagates route attributes to this peer unchanged. This behavior can be optionally restricted to only specific attributes.

Advertise Capability
Advertises the selected capabilities to this neighbor, may be one of:

Dynamic
Enables negotiation of the `dynamic` capability with this neighbor or peer group.

Extended Next-Hop
Enables negotiation of the `extended-nexthop` capability with this neighbor or peer group. This capability can set IPv6 next-hops for IPv4 routes when peering with IPv6 neighbors on interfaces without IPv4 connectivity. This is automatically enabled when peering with IPv6 link-local addresses.

ORF
Advertise outbound route filtering capability to this peer.

Disable Capability Negotiation
Disables dynamic capability negotiation with the peer. When set, the router does not advertise capabilities, nor does it accept them. This results in using only locally configured capabilities.

Override Capability Negotiation
Ignores capabilities sent by the peer during negotiation and uses locally configured capabilities instead.

TTL Security Hops
Sets a specific hop count at which neighbors must be reached, rather than the maximum value set by **eBGP multi-hop**.

This cannot be set if **eBGP multi-hop** is set.

Disable Connected Check
Disables a check that normally prevents peering with eBGP neighbors which are not directly connected. This enables using loopback interfaces to establish adjacency with peers.

eBGP Multi-Hop
The maximum allowed hops between this router and the neighbor, in the range `1-255`. When enabled without a specific value, the default is `1`.

This value cannot be set if **TTL Security Hops** is set.

Enforce eBGP Multi-Hop
When set, enforces that neighbors perform multi-hop.

Local AS

Local AS Number
Sets the local AS number sent to this neighbor, which replaces the AS number configured on the BGP server itself. By default, this value is prepended to the AS path for routes received from this neighbor or peer group, and is added to the AS path for routes sent to this neighbor or peer group after the AS number from the BGP sever.

Do not prepend eBGP
Suppresses prepending this AS number to the AS path for received routes from eBGP.

Do not prepend iBGP
Suppresses prepending this AS number to the AS path for received routes from iBGP.

Maximum Prefix
Defines the maximum number of prefixes this router will accept from the peer before tearing down the BGP session.

Note: This action is considered harsh, and the best practice is to filter received prefixes by other mechanisms such as a `prefix-list` rather than to abruptly break contact in this way.

Maximum Prefix
The maximum number of prefixes to allow from the peer, from `1-4294967295`.

Warn Percentage
Warning message threshold, from `1-100` percent.

Warn Only
Warn the peer when the limit is exceeded, rather than disconnecting.

Restart Interval
Restarts the connection after warning limits are exceeded. The restart is performed at the defined interval, in minutes, from `1-65535`.

Maximum Prefix Out
Limits the number of prefixes which will be sent to the neighbor by FRR.

Remove Private AS

Remove Outbound
Prevents the BGP daemon from sending routes with private AS numbers to this peer.

Apply to All
When present, this action applies to all ASNs.

Replace with Local
When present, replaces private AS numbers with the AS number of this router.

Route Client

Route Reflector Client
Configures this peer as a route reflector client. This allows routes received from peers in the same AS or using iBGP to be reflected to other peers, avoiding the need for a full mesh configuration between all routing peers.

Route Server Client
Configures this peer as a route server client. This enables transparent mode, which retains attributes unmodified, and maintains a local RIB for this peer.

Solo Peer
Instructs the router to prevent reflection of routes received from this neighbor back to this neighbor. This option is not useful in peer groups with multiple members.

Advanced BGP Configuration

Advanced Options

Default Local Preference
Configure default Local Preference value (0-4294967295, higher=more preferred)

Table Map
Uses the specified route map to control how routes received from BGP peers are passed to the dynamic routing manager process, and thus, into routing tables.

Advanced Timers

Coalesce Timer
Configures the Subgroup coalesce timer, in milliseconds (`1-4294967295`).

Route Map Delay
Time to wait (in seconds) before processing route map changes. A value of `0` disables the timer and stops route updates from happening when route maps change.

Dampening
BGP route flap dampening (RFC 2439) prevents unstable routers from adversely affecting routing behavior.

Time Penalty Half Life
The time duration during which the stability value will be reduced by half if the route is unreachable.

When to Reuse a Route
Stability threshold that must be crossed for a route to be reused.

Start Suppressing Route
Stability threshold that, when crossed, a route will be suppressed.

Max Time to Suppress
Maximum time to suppress a route considered stable.

Advanced Routing Behavior

Disable Fast External Failover
Do not immediately reset session if a link to a directly connected external peer goes down.

Network Import Check
Checks if a BGP network route exists in IGP before creating BGP table entries.

Note: This should be disabled if this router wants to advertise manually summarized routes that do not exist in the routing table, such as collecting all local routes into a larger network in which they can be contained (e.g. `10.14.0.0/16`).

Reject AS_SET/AS_CONFED_SET Routes
Reject incoming and outgoing routes with AS_SET or AS_CONFED_SET type.

Route Reflecting

Route Reflector Outbound
Allows attributes modified by route maps to be reflected.

Cluster ID
Route reflector cluster ID.

Disable Client-to-Client
Disables reflection of routes from one client to another client.

Aggregate Behavior

Configures route aggregation using a list of prefixes. More specific routes contained within the specified prefixes will be aggregated into the larger prefix, minimizing the set of networks advertised to peers.

Aggregate Address
The prefix to aggregate.

Generate AS Set
When present, routes for the specified prefix will include an AS set. An AS set is a collection of AS numbers for which routes have been aggregated. This allows peers to detect routing loops, duplicate routes, and so on.

Summary Only
When present, aggregated routes for this prefix will not be announced, so peers only see the aggregate prefix and not the component networks.

Multi-Exit Discriminator

Deterministic MED
Determine route selection locally, even when MED values are present. Picks the best MED path from neighbor advertisements.

Always Compare MED
Instructs the BGP daemon to always consult MED values in routes, no matter which AS the routes were received through.

Max MED

Administratively applied max MED (Indefinite)
Sends a MED value of `4294967294` at all times.

Definite admin max MED
Sends this value at all times instead of the default `4294967294`.

Time Period for Max MED on startup
Sends a MED value of `4294967294` at startup for this number of seconds.

Max MED used during startup
Sends this value at startup instead of the default `4294967294`.

Confederation

AS Confederation
Configures an AS number for the entire group of iBGP routers participating in confederation.

Confederation Peers
Configures the sub-AS number for the subset of peers inside a group of iBGP routers participating in confederation.

Distance

Administrative Distance
Manually configures the administrative distance for a given prefix

Define
The administrative distance for this prefix, from `1-255`.

IP Source Prefix
The IP prefix to which this distance will be applied.

Access List
An access list which can be used to apply the distance to only a subset of the configured prefix.

BGP Distance
Configures distance values which control how BGP will treat routes based on the length of their AS path.

AS External Routes
The distance at which routes are considered external, from `1-255`.

AS Internal Routes
The distance at which routes are considered internal, from `1-255`.

Local Routes
The distance at which routes are considered local, from `1-255`.

Best Path Selection

Controls how the BGP daemon determines the best path to a destination.

Compare Path with Confederation
Considers the length of confederation path sets and sequences.

Ignore AS Path
Ignores AS path lengths when computing the route to a destination.

Multipath Relax

Allow Load Sharing
Consider paths of equal length when choosing between multiple paths to a destination, rather than looking for an exact match. This allows load sharing across different AS paths, so long as they are of equal length.

Generate an AS_SET
Adds AS set information for aggregate routes.

Compare Router ID
Uses the router ID of peers (or originator ID, if present) to break ties when computing paths to a destination based on other information. A lower router ID will win in a tie.

MED Confederation

Compare MED among confederation paths
Compare confederation path MEDs

Treat missing MED as least preferred path
If a route is missing MED information, it will be considered least preferred.

eBGP

eBGP Nexthop Connected
Disable checking if nexthop is an eBGP session.

Enforce First AS
Enforce the first AS for eBGP routes

Disable eBGP Require Policy
Disable the requirement to apply incoming and outgoing filter to eBGP sessions.

Networking Behavior

Subgroup Packet Queue
Maximum size of the subgroup packet queue.

Write Quanta
Controls the size of peer update transmissions.

BGP AS Paths

AS Path access lists entries determine if networks are allowed or denied in specific BGP configuration contexts. They are primarily used in BGP route maps, but also can be used in other areas of BGP configuration which accept AS Path lists as parameters.

The order of entries inside an AS Path list is important, and this order is determined by a sequence number. As with other access lists, AS Path access lists implicitly deny anything not matched.

AS Path lists are managed at **Services > FRR BGP** on the **AS Paths** tab.

The order of entries inside an AS Path list is important, and this order is determined by a sequence number.

BGP AS Path Configuration

The **AS Paths** tab contains a list of current AS Path lists, if any, and controls to manage the entries (e.g. edit, delete).

The  **Add** button creates a new AS Path list.

When creating or editing an AS Path list, the following options are available:

Name
The name of this BGP AS Path list.

Description
A text description of this list (e.g. its purpose)

AS Path Entries

A list of AS paths to match in this list. Click  **Add** to create additional rules on the same list.

Sequence
The sequence number for this rule, which controls the order in which rules are matched inside this AS Path list. Each rule must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use `10`, `20`, `30`, rather than `1`, `2`, `3`.

Action
The action taken when this AS Path rule is matched, either *permit* or *deny*.

Regular Expression
A regular expression pattern which will match on the AS number.

Regular expression patterns can use common pattern special characters for matching, but also a special _ character. The _ character matches common AS delimiters such as start of line, end of line, space, comma, braces, and parenthesis. The _ character can be used on either side of an AS number to match it exactly, such as `_65534_`.

BGP Community Lists

A BGP community, as defined in RFC 1997, is a group of destinations which share common properties. Community Lists define sets of community attributes which the BGP daemon can use to match or set community values in routing updates. BGP communities determine AS membership and priority values in BGP-specific contexts such as route-maps.

BGP community lists are managed at **Services > FRR BGP** on the **Communities** tab.

The order of entries inside a community list is important, and this order is determined by a sequence number.

BGP Well-Known Communities

There are several "well-known" communities available for use in Community Lists. Each of these communities have special meanings:

internet
: A community value of `0`, indicating the Internet as a destination.

no-export
: Routes received carrying this attribute value must not be exported to routers outside the current confederation.

no-advertise
: Routes received carrying this attribute value must not be advertised to any other BGP peer.

local-as
: Also known as "No Export Subconfed". Routes received carrying this attribute value must not be advertised to any external BGP peer, even those in the same confederation.

blackhole
: Routes received carrying this attribute should not be routed (e.g. null routed).

graceful-shutdown
: Indicates RFC 8326 Graceful Shutdown capability, which allows BGP routers to inform peers that specific paths can be gracefully shut down rather than abruptly terminated when performing an intentional shutdown.

no-peer
: Indicates that routes with this community value should not be readvertised to peers (RFC 3765).

BGP Community List Configuration

The **Communities** tab contains a list of current community lists, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new community list.

When creating or editing a community list, the following options are available:

Name
: The name of this BGP community list.

Description
: A text description of this list (e.g. its purpose)

Community List Type
: The type of community list to use for this entry, which controls how values are matched.

Standard
Matches based on specific values for community attributes, either AS:Value pairs or names of *well-known communities*.

Expanded
Matches based on an ordered list using a regular expression. Due to the use of regular expression evaluation, these lists incur a performance penalty.

Community List Entries

A list of rules for this community list. Click  **Add** to create additional rules on the same list.

Sequence
The sequence number for this rule, which controls the order in which rules are matched inside this community list. Each rule must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use `10`, `20`, `30`, rather than `1`, `2`, `3`.

Action
The action taken when this Community List rule is matched, either *permit* or *deny*.

Community
The value of the community to match.

Standard Community Lists
This is a space-separated list of communities in AS:Value format, or from the *well-known communities* list.

Expanded Community Lists
A string containing a regular expression to match against.

Regular expression patterns can use common pattern special characters for matching, but also a special _ character. The _ character matches common AS delimiters such as start of line, end of line, space, comma, braces, and parenthesis.

BGP RPKI Cache Servers

Resource Public Key Infrastructure (RPKI) is a means by which FRR can enact Prefix Origin Validation (POV) to ensure that it is talking to the correct origin for a given AS.

This validation is **not** performed by FRR or other routers directly, but by trusted servers which cache the information.

Note: For more details, see RFC 6810 for the protocol and RFC 6811 for validation.

RPKI happens over a plain TCP connection but FRR can protect this by performing the validation over SSH.

Route maps can be used to filter routes based on a validated origin.

RPKI Cache Servers are managed at **Services > FRR BGP** on the **RPKI Cache Servers** tab.

RPKI Cache Server Configuration

The **RPKI Cache Servers** tab contains a list of current RPKI Cache Servers, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new RPKI Cache Server.

When creating or editing an RPKI Cache Server, the following options are available:

Address
: **Required**. The IP **Address** or hostname of the RPKI Cache Server, and the **Port** number upon which the service is listening.

Preference
: **Required**. A preference value FRR can use to decide between multiple RPKI Cache Servers.

SSH Options
: The best practice is to encrypt communication with the RPKI Cache Server using SSH. The remaining options set up an SSH session, and all are optional.

  Username
  : The username to use when connecting to the server via SSH.

  Private Key Path
  : Full filesystem path to the private key for this router.

  > **Warning:** This must not have a passphrase as there is no way to securely store and use a passphrase. Protect the private key file appropriately, but it must also be accessible to FRR.

  Public Key Path
  : Full filesystem path to the public key for this router.

  Known hosts Path
  : Full filesystem path to a file containing valid public keys for RPKI Cache Servers in SSH `known_hosts` format.

BGP Status

The status page for BGP can be found at **Status > FRR**, **BGP** tab. It can also be accessed by visiting the **Status** tab while configuring BGP options.

See also:

For general FRR status information, see *FRR Status*.

BGP Routes
: A list of routes being advertised to BGP and received from BGP. The list includes information about the status of the route, its origin, AS path, and metrics.

BGP IPv6 Routes
: Same as **BGP Routes** but for IPv6 routes.

BGP Summary
: A high-level overview of BGP operations, including a brief list of neighbors and a summary of their activity and statistics.

BGP Neighbors
Detailed information about BGP peers and their status, including the state of the neighbor (e.g. `Established`), its capabilities, and much more.

BGP Peer Groups
Information about BGP Peer Groups, if any are in use.

BGP Next Hops
A list of known route targets kept in a cache by FRR.

BGP Memory
A summary of memory (RAM) consumed by BGP.

30.10.3 Open Shortest Path First v2 (OSPF)

Open Shortest Path First v2 (OSPF) is a link-state routing protocol defined by RFC 2328. OSPF automatically locates neighboring IPv4 routers within an autonomous system, typically with multicast, and exchanges IPv4 routing information for networks reachable through each neighbor.

OSPF is an interior routing protocol (IGP), and facilitates routing between private links or segments of local networks.

OSPF Terminology

OSPF has common terms used throughout this section which can be confusing for those unfamiliar with the protocol.

Area
A collection of routers inside an AS, each sharing the same area ID. An Area ID is typically formatted like an IP address in dotted quad notation, `nnn.nnn.nnn.nnn`, but can also be expressed as an unsigned 32-bit integer.

Area Border Router (ABR)
A router connected to multiple areas.

Autonomous System Boundary Router (ASBR)
A router connected to external networks (outside the area).

Backbone
The central area of an AS, typically area `0.0.0.0`. All areas in the AS connect to the backbone through ABRs.

Cost
A numeric value assigned to a link between networks, used by OSPF to calculate optimal paths to a destination. Typically, higher bandwidth or higher quality circuits will be assigned a low cost, while circuits that are undesirable will be given a high cost. OSPF will prefer to use a route when it has the lowest total cost from a source to a destination.

Designated Router (DR)
In a network with multiple routers, one of them will be elected as a designated router (DR) using `Hello` messages. The DR takes on the task of generating LSA messages for the network, among other special duties.

Flooding
The mechanism by which OSPF routers distribute link state database information to neighbors.

Hello
Special OSPF messages which introduce neighbors to each other. Using these messages, neighbors can discover each other and begin to form routing relationships.

Interior Gateway Protocol (IGP)
A routing protocol, such as OSPF, which exchanges information about how to reach networks *inside* an autonomous system.

Link State Advertisement (LSA)
Messages sent by OSPF routers which describe the state of network links, or the router itself, including information about its interfaces and other neighbors.

Link State Database (LSDB)
A database containing the collected LSA messages of all routers and networks in the domain.

Link State Advertisement Message Types

LSA messages each have a type, indicating the information carried within. These types may be referenced throughout this section when describing routing behaviors.

Type 1 - Router LSA
Sent by every router in an area. Contains a description of all links on the router, including their state and costs.

Type 2 - Network LSA
Sent by the DR for a network. Contains a description of every router attached to the network, including the DR.

Type 3 - Network Summary-LSA
Sent by ABRs. Contains a description of destinations outside the current area (inter-area) when the destination is an IP network.

Type 4 - ASBR Summary-LSA
Similar to Type 3, but when sent when the destination is an ASBR.

Type 5 - AS-external LSA
Sent by ASBRs. Contains a description of destinations outside of this AS. Typically, each message only contains information about a single destination.

Type 6 - Multicast Group Membership LSA
Not used.

Type 7 - NSSA External Link-State Advertisements
Similar to Type 5, but are only exchanged inside an NSSA.

Type 8 - External attribute LSA
Carry information from external routing protocols, such as BGP, when such destinations are announced with Type 5 LSAs.

Type 9 - Link Scope Opaque LSA
Carries information intended for uses other than OSPF, such as available bandwidth. It is carried through to other routers without being processed by OSPF itself. Type 9 messages are for other routers on the same link.

Type 10 - Area Scope Opaque LSA
Similar to Type 9, but flooded to all routers in an area.

Type 11 - AS Scope Opaque LSA
Similar to Type 9, but flooded to all routers throughout the AS, except for special areas such as stubs.

Area Types

OSPF Areas can be one of several types which alter their behavior in important ways.

Normal
A typical area in which all routers know all possible routes.

Stub Area
An area with no external connections. Since traffic passing out of a stub area must pass through an ABR, it only needs to know about routes to the ABR, not beyond the ABR. Routers in a stub area do not receive Type 5 LSAs.

Totally Stub Area
Similar to a stub area, but routers also do not receive summary LSA messages except for default route information. As such, they do not receive LSA messages of type 3, 4, or 5.

Not-so-Stubby-Area (NSSA)
Similar to a Stub area but it may contain static routes to non-OSPF networks. Routers in an NSSA exchange external routing information in Type 7 LSAs instead of Type 5.

NSSA Totally Stub Area
Similar to both NSSA and a Totally Stub area. As such, they do not receive LSA messages of type 3, 4, or 5.

Metric Types

Type 1 or E1
A Type 1 external metric, also known as E1, uses a similar cost calculation to typical link states, where internal and external costs are added together to find the total cost.

Type 2 or E2
A Type 2 external metric, also known as E2, only considers external costs and ignores internal costs.

OSPF Required Information

Before starting, take the time to gather all the information required to form an OSPF adjacency to a neighbor. At a minimum, FRR will need to know these items:

Local Router ID
Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

Local OSPF Area
A designation for the set of networks to which this router belongs. Can be any number capable of being expressed in dotted quad notation (IPv4 address) or as a 32-bit unsigned integer.

In typical OSPF configurations such as this example, `0.0.0.0` is the backbone area between all routers and each local network has its own area. For simpler deployments, the only area may be `0.0.0.0` on every router and interface. Using a single area disables some features such as route summarization.

OSPF Active Interfaces
The interfaces on this router upon which the OSPF daemon will advertise itself and look for neighbors. These interfaces are connected to network segments with other routers. They may be connected to local networks or remote point-to-point links. These interfaces must be configured with IP addresses.

OSPF Active Interface Cost Values
OSPF calculates the most efficient way to route between networks based on the total cost of a path from source to destination. Less desirable links (e.g. wireless) can be given a higher cost so that paths over faster networks will be used by traffic unless the preferred path is unavailable. For single connections to other networks, this value is not necessary and may be omitted or set to a simple default such as `5` or `10`.

OSPF Passive Interfaces
These interfaces contain networks which should be advertised as reachable through this router, but do not contain other routers.

Summary Routes
A list of networks to advertise instead of using networks from directly attached interfaces. This allows many similar routes to be summarized as one larger route if they can all be contained within a larger subnet.

This can only be done when connecting to an ABR in a configuration with multiple areas.

Note: This is optional, but without this in place, every network which must be advertised to neighbors must be attached to this firewall and the interfaces must be added as passive interfaces in OSPF. Using a summary route is cleaner in that it advertises less (one large subnet vs many small subnets) and that it requires less ongoing configuration.

The example in this section uses the following values:

Table 2: Example OSPF Configuration

Item	Value
Local Router ID	10.2.0.1
OSPF Backbone Area	0.0.0.0
Local OSPF Area	0.0.0.2
Active Interfaces (Cost)	OPT1 (10)
Passive Interfaces	LAN
Summary Routes	10.2.0.0/16

OSPF Example Configuration

This example configuration implements a multi-area OSPF setup using the required information from *Example OSPF Configuration*.

Assumptions

This example makes a few assumptions for brevity and to keep the example simple, including:

- The remote peer is already configured for OSPF with equivalent settings.
- Transit to the peer across a directly attached shared network is already configured, for example over a VPN, shared network segment, or peer-to-peer link.
- The link to the peer is capable of handling multicast traffic.
- Firewall rules pass OSPF traffic, which is protocol `89`. It is not TCP or UDP. Firewall rules must allow multicast traffic destinations for OSPF, and it cannot be restricted to specific sources and destinations in this example.

Example Configuration

The OSPF configuration must be done in the following order initially. Later changes may be made in any order. This is because a valid OSPF configuration requires **Interface** tab entries which must exist before the main OSPF settings can be saved.

OSPF Area Configuration

- Navigate to **Services > FRR OSPF**, **Areas** tab
- Click  **Add** to create a new area
- Set the following options:

  Area
  `0.0.0.0`

  Description
  `Backbone`

- Click **Save**
- Click  **Add** to create a new area
- Set the following options:

  Area
  `0.0.0.23`

  Description
  `Local Area`

  Area Type
  Stub Area (stub)

  Summary Range
  `10.2.0.0/16`

- Click **Save**

OSPF Interface Configuration

- Navigate to **Services > FRR OSPF**, **Interfaces** tab
- Click  **Add** to create a new interface
- Set the following options:

  Interface
  LAN

  Interface is Passive
  Checked

  Area
  `0.0.0.23`

- Click **Save**
- Click  **Add** to create a new interface
- Set the following options:

  Interface
  OPT1

  Ignore MTU
  Checked

  Metric
  `10`

  Area
  `0.0.0.0`

- Click **Save**

OSPF Configuration

- Navigate to **Services > FRR OSPF**, **OSPF** tab
- Set the following options:

  Enable
  Checked

  Router ID
  `10.2.0.1`

  Default Area
  `0.0.0.0`

- Click **Save**

FRR Global Configuration

- Navigate to the **[Global Settings]** tab
- Set the following options:

  Enable
  Checked

  Master Password
  Create a random string to use

- Click **Save**
- Navigate to the **Status** tab
- Confirm that the OSPF neighbor is present, and its routes are in the table

OSPF Configuration

OSPF FRR configuration, as shown in the example, can be fairly straightforward. That said, there are a number of ways to fine-tune the behavior and create complex OSPF routing configurations.

Read through each section before attempting to create a new OSPF configuration.

Note: Though the documentations covers OSPF in *tab* order, OSPF interfaces must be configured before the general OSPF options.

OSPF Tab Configuration

The options on this page configure the general behavior of OSPF in FRR, including various default values used when more specific options are not available.

> **Warning:** When creating an initial OSPF configuration, configure the *interfaces to be used with OSPF* first. The interface configuration must exist before these general OSPF settings can be saved.

General Options

Enable
: Master enable switch for OSPF routing. When checked, FRR will start the OSPF routing daemon and attempt to use the OSPF settings in this section.

Log Adjacency Changes
: When set, instructs the OSPF daemon to log changes in neighbor adjacencies. This is useful for tracking changes to neighbor relationships, especially during initial configuration.

Router ID
: Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

SPF Hold Time
: SPF timers determine when the router will make SPF routing decisions. Lowest time allowed between SPF calculations. Specified in milliseconds from `0-600000`, with a default value of `1000`.

  The maximum time is calculated as 10x this value.

  SPF calculations are adaptive, and if a new event occurs which would otherwise trigger a calculation before the hold timer expires, then the hold is increased by the **SPF Hold Time** value, up to the maximum. This avoids excessive consecutive recalculations.

SPF Delay
: Controls timers that determine when the router will make SPF routing decisions. Minimum time after an event occurs before allowing SPF calculation. Lower values will react faster to changes, but can be less stable. Specified in milliseconds from `0-600000`, with a default value of `200`.

Modules

Enable SNMP AgentX
Enable agentx for accessing FRR OSPF data via SNMP with the NET-SNMP package.

Default Area

Default Area
Default OSPF area for this instance of OSPF. Used when an area is required but not defined elsewhere. See *OSPF Area Configuration* for details.

Default Area Type
Sets the type for the **Default Area**. See *Area Types* and *OSPF Area Configuration* for details.

OSPF Networks

> **Warning:** This section is **deprecated** and will not be covered here. Define *areas* and use areas on *interfaces* instead. Use summary routes instead, or use route-maps and distribute lists to limit distributed networks.

Route Redistribution

This section controls which, if any, IPv4 routes are redistributed to OSPF neighbors from other sources (*Dynamic Routing Protocol Lists*). OSPF can redistribute IPv4 routes from connected networks, kernel routes, BGP, and FRR static routes.

Redistribute
Enables redistribution of routes from the given source.

Metric
Advertise the routes from this source as having the given metric.

Metric Type
The type of metric, either `1` or `2`. See *Metric Types* for details about each type operates.

Route Map
Apply the given route map to the redistributed route advertisements.

Distribute List
Applies the given *access list* to routes redistributed from a given route source.

Default Route Redistribution

Redistribute
Enables origination of a Type 5 AS-External LSA containing default route information into all areas capable of external routing.

Always Redistribute
Always advertise a default route, even when a default route is not present in the local routing table.

Default Metric
Advertise the default route as having the given metric.

Default Metric Type
The type of metric, either `1` or `2`. See *Metric Types* for details about each type operates.

Route Map
Apply the given route map to the outbound default route advertisement.

Advanced

RFC 1583 Compatibility
Enables compatibility with the older OSPF standard from RFC 1583, which has been obsoleted by the newer RFC 2328. The specific change this option enables relates to external path preference calculation and routing loop prevention. See RFC 2328 section **G.2** for specific details.

Opaque LSA
Enables Opaque LSAs capability, as described in RFC 2370.

Reference Bandwidth
A base value, in Mbit/s, which is used when OSPF automatically calculates cost values. The default value is `100` which means that an interface with 100Mbit/s of bandwidth or greater will have a cost of `1`, with lower bandwidth values incurring higher cost values.

All routers in the same area should use the same value, otherwise automatic cost calculations would fail to accurately represent total path costs between routers.

Max Metric

Administratively Enable Max Metric
Sets the administrative distance of routes through this router to infinity, so that other routers will avoid using this router to reach other networks. Networks on this router are still reachable. See RFC 3137 for more information.

Startup Seconds
Conditionally sets the administrative distance of routes through this router to infinity for a period of time after startup.

This allows other routers in the area to avoid using routes through this router until a full convergence is achieved.

Shutdown Seconds
Conditionally sets the administrative distance of routes through this router to infinity for a period of time after shutdown.

This allows other routers in the area to avoid using routes through this router until a full convergence is achieved.

Write Multiplier
Number of interfaces processed per write operation, from `1-100`. Default value is `20`.

ABR Type
Controls the behavior of Area Border Router (ABR) functionality.

cisco|ibm
The default behavior of OSPF in FRR, discussed in RFC 3509. This behavior allows an ABR without a backbone connection to act as an internal router for all connected areas.

shortcut
Discussed in draft-ietf-ospf-shortcut-abr-02, this behavior allows ABRs to consider summary LSAs from all attached areas, rather than being forced to route through a suboptimal path only because it is shorter.

standard
The ABR behavior described in the original OSPF standard. When set, a router attached to multiple areas requires a connection to a backbone. If no backbone is available, traffic attempting to cross areas will be dropped.

OSPF Area Configuration

OSPF Areas are managed at **Services > FRR OSPF** on the **Areas** tab.

The **Areas** tab contains a list of current OSPF areas, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new area.

The remaining sections on this page cover the various options available when creating or editing an area entry.

Area Options

Area
A designation for the set of networks. Can be any number capable of being expressed in dotted quad notation (IPv4 address) or as a 32-bit unsigned integer.

In typical OSPF configurations `0.0.0.0` is the backbone area between all routers and each local network has its own area. For simpler deployments, the only area may be `0.0.0.0` on every router and interface. Using a single area disables some features such as route summarization.

Description
Some text describing this area.

Area Type
Defines how this area behaves. For a list of all types and how they operate, see *Area Types*.

Default Route Cost
Sets the cost applied to default route summary LSA messages sent to stub areas.

ABR Shortcut
For use with **ABR Type** set to *Shortcut* (*Advanced*), this advertises the area as capable of ABR shortcut behavior (draft-ietf-ospf-shortcut-abr-02).

Authentication

This option enables authentication for this area. Communication from peers must contain the expected authentication information to be accepted, and outgoing packets will have authentication information added.

Authentication passwords and keys are configured by the *OSPF interface settings*.

Authentication Type
Sets the type of authentication to use in this area.

None
Do not use authentication in this area.

Message Digest (MD5 Hash)
Enables MD5 HMAC authentication for this area. This is stronger authentication than simple passwords.

Simple Password
Enables simple password authentication for this area.

Route Summarization

Configure summarization of routes inside the given prefixes. Instead of Type 1 (Router) and Type 2 (Network) LSAs, it creates Type 3 Summary LSAs instead.

The options here can be repeated for multiple prefixes. Click  **Add** for each additional prefix.

Summary Prefix
The prefix to summarize.

Do Not Advertise
Disable advertisement for this prefix.

Cost
Apply the specified cost to summarized routes for this prefix.

Substitute Prefix
Instead of advertising the first prefix, advertise this prefix instead.

ABR Summary Route Filtering

Export List
Uses the given ACL to limit Type 3 summary LSA messages for intra-area paths that would otherwise be advertised. This behavior only applies if this router is the ABR for the area in question.

Import List
Similar to `export-list`, but for routes announced by other routers into this area.

Filter List (Out|In)
Similar to **Export List** and **Import List** but uses prefix lists instead of ACLs, and can work in either direction.

OSPF Interface Configuration

OSPF interfaces are managed at **Services > FRR OSPF** on the **Interface** tab.

The **Interfaces** tab contains a list of current OSPF interfaces, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new interface.

OSPF must use one or more interfaces to announce itself to neighbors and to receive announcements from neighbors. At least one interface must be configured and active in order to locate neighbors and form an adjacency.

The remaining sections on this page cover the various options available when creating or editing an interface entry.

Interface Options

Interface
The interface to use with OSPF. Entries should be added for interfaces which connect to other routers (neighbors) as well as interfaces containing networks which should be advertised to OSPF neighbors.

Description
Text describing the purpose of this interface in OSPF.

Network Type
Manually configures a specific type of network used on a given interface, rather than letting OSPF determine the type automatically. This controls how OSPF behaves and how it crafts messages when using an interface.

Tip: Most environments will use either *Broadcast* mode (e.g. Ethernet) or *Point-to-Point* mode (e.g. VPNs).

Broadcast
Broadcast networks, such as typical Ethernet networks, allow multiple routers on a segment and OSPF can use broadcast and multicast to send messages to multiple targets at once. OSPF assumes that all routers on broadcast networks are directly connected and can communicate without passing through other routers.

Non-Broadcast
Non-broadcast networks can contain multiple routers but do not have broadcast or multicast capabilities. Due to this limitation, neighbors must be manually configured using the **Neighbor** tab (*OSPF Neighbor Configuration*).

When using this mode, OSPF simulates a broadcast network using Non-Broadcast Multi-Access (NBMA) mode, but transmits messages to known neighbors directly.

Point-to-Multipoint
Similar to *Non-Broadcast* mode, but connections to manually configured neighbors are treated as a collection of point-to-point links rather than a shared network. Similar to a point-to-point network, OSPF disables DR election.

Point-to-Point
A point-to-point network links a single pair of routers. The interface is still capable of broadcast, and OSPF will dynamically discover neighbors. With this type of network, OSPF disables election of a DR.

Interface is Passive
Configures the specified interface as passive. This prevents the interface from actively participating in OSPF, while still allowing OSPF to operate on networks connected to that interface. This is commonly used for local interfaces without other routers attached. OSPF will announce networks attached to passive interfaces as stub links.

Ignore MTU
When present, OSPF will ignore the MTU advertised by neighbors and can still achieve a full adjacency when peers do not have matching MTU values.

Tip: If two neighbors are stuck in an `ExStart` state, that is typically from an MTU mismatch. If fixing the MTU mismatch is not viable, set this option on both sides.

OSPF Interface Handling

Metric
A manual cost value to apply to this interface, rather than allowing automatic cost calculation to take place.

In situations where multiple paths are possible to the same destination, this allows OSPF to prefer one path over another when all else is equal.

Area
This defines the interface as a member of the given area. If this is left blank, FRR will take the value set in *Default Area*.

Accept Filter
When set, automatically configures lists behind the scenes to prevent this interface subnet from being advertised to, or received from, OSPF.

This can avoid problems seen in Multi-WAN deployments where there are multiple paths to a remote neighbor and a router can end up learning a route to itself across a VPN link, which is problematic.

Authentication

Configures authentication for OSPF neighbors on this interface. All routers connected to this interface must have identical authentication configurations. This can also be enabled in the area settings.

Authentication Type
Sets the type of authentication to use in this area.

None
Do not use authentication in this area. This is the default behavior, but may be explicitly configured to override the authentication configured for this area.

Message Digest (MD5 Hash)
Enables MD5 HMAC authentication for this area. This is stronger authentication than simple passwords.

Simple Password
Enables simple password authentication for this area.

Password
Password to use with *Simple Password* or key to use with *Message Digest* authentication.

This value must match all neighbors reachable through this interface. Simple passwords may be up to 8 characters, Message Digest passwords (keys) may be up to 16 characters.

Advanced

Router Priority
A priority value, from `0-255`, assigned to this router. When determining which router will become the Designated Router (DR), the router with the highest priority is more likely to be elected as the DR.

The default value is `1`. The value `0` is special and will prevent this router from being chosen as DR.

Retransmit Interval
The interval, in seconds from `1-65535`, at which this router will retransmit Link State Request and Database Description messages. This is also known as the `RxmtInterval` timer in OSPF. Default value is `5`.

Hello Interval
The interval, in seconds from `1-65535`, at which this router will send hello messages. This is also known as the `HelloInterval` timer in OSPF. Default value is `10`. This timer should be set to the same value for all routers.

A lower value will result in faster convergence times, but will consume more resources.

Dead Interval
Time, in seconds from `1-65535`, without communication from a neighbor on this interface before considering it dead. This is also known as the `RouterDeadInterval` timer in OSPF. Default value is `40`. This timer should be set to the same value for all routers.

Minimal Hello
When active, the **Dead Interval** is forced to a value of `1` and OSPF will instead send this number of Hello messages each second. This allows for faster convergence, but will consume more resources.

Note: When set, this overrides the values of both **Dead Interval** and **Hello Interval**. Custom values configured in those fields will be ignored by OSPF.

BFD

When set, enables *Bidirectional Forwarding Detection* for OSPF on this interface.

OSPF Neighbor Configuration

OSPF neighbors are managed at **Services > FRR OSPF** on the **Neighbors** tab.

The **Neighbors** tab contains a list of current OSPF neighbors, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new neighbor.

In most cases, OSPF neighbors do not need to be added manually. In certain environments with links which are not capable of using multicast, manual neighbor definitions allow FRR to locate neighbors statically. This can be useful to work around limitations imposed by certain point-to-point links or VPN configurations.

The remaining sections on this page cover the various options available when creating or editing a neighbor entry.

Neighbor Options

OSPF Neighbor IPv4 Address
The IPv4 address of this OSPF neighbor.

Description
Text describing this neighbor.

Neighbor Priority
A priority value applied to neighbors in a down state.

Dead Neighbor Polling Interval
Time, in seconds, between sending OSPF Hello messages to neighbors in a down state.

OSPF Status

The status page for OSPF can be found at **Status > FRR**, **OSPF** tab. It can also be accessed by visiting the **Status** tab while configuring OSPF options.

See also:

For general FRR status information, see *FRR Status*.

OSPF General
General information about the OSPF process, active configuration, and operation.

OSPF Neighbors
A list of OSPF neighbors detected by FRR and their status.

Tip: If a neighbor is stuck in a state such as "ExStart", check that the interface MTUs match. Consider checking the interface option to **Ignore MTU** which will bypass this check.

OSPF Routes
A list of routes being advertised to OSPF and received from OSPF. The list is broken into three parts:

OSPF Network Routing Table
A list of internal networks and their status, how they are reached through neighbors and areas, etc.

OSPF Router Routing Table
A list of other routers through which OSPF networks are reached.

OSPF External Routing Table
A list of external routes from OSPF.

OSPF Database
The OSPF link state database with information about each known router and area.

OSPF Router Database
Detailed information about OSPF routers and how they are connected.

OSPF Interfaces
Detailed information about interfaces on the firewall and how they are used by OSPF.

OSPF CPU Usage
The amount of CPU time used by OSPF for various tasks and threads.

OSPF Memory
The amount of memory (RAM) used by OSPF for various tasks and threads.

30.10.4 Open Shortest Path First v3 (OSPF6)

Open Shortest Path First v3 (OSPF6) is defined by RFC 5340 and is similar to OSPF v2, but operates with IPv6 networks. Thus, it is a link-state routing protocol that automatically locates neighboring IPv6 routers within an autonomous system, typically with multicast, and exchanges IPv6 routing information for networks each neighbor.

OSPF6 is an interior routing protocol (IGP), and facilitates routing between private links or segments of local networks.

Terms used in this section are shared with OSPF, and are covered in *OSPF Terminology*.

OSPF6 Required Information

Before starting, take the time to gather all the information required to form an OSPF6 adjacency to a neighbor. This list is similar to *that of OSPF*. At a minimum, FRR will need to know these items:

Local Router ID
: Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

OSPF6 Area
: A designation for the set of networks to which this router belongs. Can be any number capable of being expressed in dotted quad notation (IPv4 address) or as a 32-bit unsigned integer.

OSPF6 Active Interfaces
: The interfaces on this router upon which the OSPF6 daemon will advertise itself and monitor for neighbors. These interfaces are connected to network segments with other routers. They may be connected to local networks or remote point-to-point links. These interfaces only require an IPv6 link local address.

OSPF6 Active Interface Cost Values
: OSPF6 calculates the most efficient way to route between networks based on the total cost of a path from source to destination. Less desirable links (e.g. wireless) can be given a higher cost so that paths over faster networks will be used by traffic unless the preferred path is unavailable. For single connections to other networks, this value is not necessary and may be omitted or set to a simple default such as `5` or `10`.

OSPF6 Passive Interfaces
: These interfaces contain networks which FRR will advertise as reachable through this router, but do not contain other routers.

The example in this section uses the following values:

Table 3: Example OSPF Configuration

Item	Value
Local Router ID	10.2.0.1
OSPF6 Area	0.0.0.0
Active Interfaces (Cost)	OPT1 (10)
Passive Interfaces	LAN

OSPF6 Example Configuration

This example configuration implements an OSPF6 setup using the required information from *Example OSPF Configuration*.

Assumptions

This example makes a few assumptions for brevity and to keep the example simple, including:

- The remote peer is already configured for OSPF6 with equivalent settings.
- IPv6 transit to the peer across a directly attached shared network is already configured, for example over a VPN, shared network segment, or peer-to-peer link.
- The link to the peer is capable of handling multicast traffic.
- Firewall rules pass OSPF traffic, which is protocol `89`. It is not TCP or UDP. Firewall rules must allow multicast traffic destinations for OSPF, and it cannot be restricted to specific sources and destinations in this example.

Example Configuration

The OSPF6 configuration must be done in the following order initially. Later changes may be made in any order. This is because a valid OSPF6 configuration requires **Interface** tab entries which must exist before the main OSPF6 settings can be saved.

OSPF6 Interface Configuration

- Navigate to **Services > FRR OSPF6**, **Interfaces** tab
- Click  **Add** to create a new interface
- Set the following options:

  Interface
  LAN

  Interface is Passive
  Checked

  Area
  `0.0.0.0`

- Click **Save**
- Click  **Add** to create a new interface
- Set the following options:

  Interface
  OPT1

  Ignore MTU
  Checked

  Metric
  `10`

Area
`0.0.0.0`

- Click **Save**

OSPF6 Configuration

- Navigate to **Services > FRR OSPF6**, **OSPF6** tab
- Set the following options:

Enable
Checked

Router ID
`10.2.0.1`

Default Area
`0.0.0.0`

- Click **Save**

FRR Global Configuration

- Navigate to the **[Global Settings]** tab
- Set the following options:

Enable
Checked

Master Password
Create a random string to use

- Click **Save**
- Navigate to the **Status** tab
- Confirm that the OSPF6 neighbor is present, and its routes are in the table

OSPF6 Configuration

There are a number of ways to fine-tune the behavior and create complex OSPF6 routing configurations. The available configuration parameters are covered throughout this section.

OSPF6 Tab Configuration

The options on this page configure the general behavior of OSPF6 in FRR, including various default values used when more specific options are not available.

Warning: When creating an initial OSPF6 configuration, configure the *interfaces to be used with OSPF6* first. The interface configuration must exist before these general OSPF6 settings can be saved.

General Options

Enable
Master enable switch for OSPF6 routing. When checked, FRR will start the OSPF6 routing daemon and attempt to use the OSPF6 settings in this section.

Log Adjacency Changes
When set, instructs the OSPF6 daemon to log changes in neighbor adjacencies. This is useful for tracking changes to neighbor relationships, especially during initial configuration.

Router ID
Typically the highest numbered local IPv4 address on the firewall. This is also frequently set as the internal or LAN side IPv4 address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

Note: Even though OSPF6 handles IPv6 routing, router IDs are specified using IPv4 addresses in dotted quad notation.

SPF Hold Time
SPF timers determine when the router will make SPF routing decisions. Lowest time allowed between SPF calculations. Specified in milliseconds from `0-600000`, with a default value of `1000`.

The maximum time is calculated as 10x this value.

SPF calculations are adaptive, and if a new event occurs which would otherwise trigger a calculation before the hold timer expires, then the hold is increased by the **SPF Hold Time** value, up to the maximum. This avoids excessive consecutive recalculations.

SPF Delay
Controls timers that determine when the router will make SPF routing decisions. Minimum time after an event occurs before allowing SPF calculation. Lower values will react faster to changes, but can be less stable. Specified in milliseconds from `0-600000`, with a default value of `200`.

Modules

Enable SNMP AgentX
Enable agentx for accessing FRR OSPF6 data via SNMP with the NET-SNMP package.

Default Area

Default Area
Default OSPF area for this instance of OSPF6. Used when an area is required but not defined elsewhere. See *OSPF6 Area Configuration* for details.

Default Area Type
Sets the type for the **Default Area**. See *Area Types* and *OSPF6 Area Configuration* for details.

Route Distribution

The **Distribute Ranges** entries instruct OSPF to associate specific IPv6 prefixes with a given OSPF6 area and to advertise them to neighbors. Additional entries can be created by the  **Add** button.

Note: In most cases these entries are not necessary to add manually, as interface entries will add appropriate statements to distribute interface subnets as needed.

Subnet to Route
The IPv6 prefix to advertise.

Area ID
The area with which to associate the prefix.

Cost
The cost associated with the prefix.

Route Redistribution

This section controls which, if any, IPv6 routes are redistributed to OSPF6 neighbors from other sources (*Dynamic Routing Protocol Lists*). OSPF6 can redistribute IPv6 routes from connected networks, kernel routes, BGP, and FRR static routes.

Redistribute
Enables redistribution of routes from the given source.

Route Map
Apply the given route map to the redistributed route advertisements.

Route Filtering

Export List
Uses the given ACL to limit Type 3 summary LSA messages for intra-area paths that would otherwise be advertised.

Import List
Similar to `export-list`, but for routes announced by other routers into this area.

Filter List (Out|In)
Similar to **Export List** and **Import List** but uses prefix lists instead of ACLs, and can work in either direction.

Advanced

Reference Bandwidth
A base value, in Mbit/s, which is used when OSPF6 automatically calculates cost values. The default value is `100` which means that an interface with 100Mbit/s of bandwidth or greater will have a cost of 1, with lower bandwidth values incurring higher cost values.

All routers in the same area should use the same value, otherwise automatic cost calculations would fail to accurately represent total path costs between routers.

Distance
Sets an administrative distance for routes obtained via OSPF6. This can be configured globally as well as for specific types of OSPF6 routes.

External Distance
Sets the administrative distance for external OSPF6 routes.

Inter-area Distance
Sets the administrative distance for OSPF6 routes between areas.

Intra-area Distance
Sets the administrative distance for OSPF6 routes inside an area.

OSPF6 Interface Configuration

OSPF6 interfaces are managed at **Services > FRR OSPF6** on the **Interface** tab.

The **Interfaces** tab contains a list of current OSPF6 interfaces, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new interface.

OSPF6 must use one or more interfaces to announce itself to neighbors and to receive announcements from neighbors. At least one interface must be configured and active in order to locate neighbors and form an adjacency.

The remaining sections on this page cover the various options available when creating or editing an interface entry.

Interface Options

Interface
The interface to use with OSPF6. Entries should be added for interfaces which connect to other routers (neighbors) as well as interfaces containing networks which should be advertised to OSPF6 neighbors.

Description
Text describing the purpose of this interface in OSPF6.

Network Type
Manually configures a specific type of network used on a given interface, rather than letting OSPF6 determine the type automatically. This controls how OSPF6 behaves and how it crafts messages when using an interface.

Broadcast
Broadcast networks, such as typical Ethernet networks, allow multiple routers on a segment and OSPF6 can use broadcast and multicast to send messages to multiple targets at once. OSPF6 assumes that all routers on broadcast networks are directly connected and can communicate without passing through other routers.

Point-to-Point
A point-to-point network links a single pair of routers. The interface is still capable of broadcast, and OSPF6 will dynamically discover neighbors. With this type of network, OSPF6 disables election of a DR.

Interface is Passive
Configures the specified interface as passive. This prevents the interface from actively participating in OSPF6, while still allowing OSPF6 to operate on networks connected to that interface. This is commonly used for local interfaces without other routers attached. OSPF6 will announce networks attached to passive interfaces as stub links.

Ignore MTU
When present, OSPF6 will ignore the MTU advertised by neighbors and can still achieve a full adjacency when peers do not have matching MTU values.

Tip: If two neighbors are stuck in an `ExStart` state, that is typically from an MTU mismatch. If fixing the MTU mismatch is not viable, set this option on both sides.

OSPF6 Interface Handling

Area
This defines the interface as a member of the given area. If this is left blank, FRR will take the value set in *Default Area*.

Instance ID
An alternate OSPF6 instance identifier for this interface. Typically, omitted or set to `0`.

Metric
A manual cost value to apply to this interface, rather than allowing automatic cost calculation to take place.

In situations where multiple paths are possible to the same destination, this allows OSPF6 to prefer one path over another when all else is equal.

Advanced

Router Priority
A priority value, from `0-255`, assigned to this router. When determining which router will become the Designated Router (DR), the router with the highest priority is more likely to be elected as the DR.

The default value is `1`. The value `0` is special and will prevent this router from being chosen as DR.

Hello Interval
The interval, in seconds from `1-65535`, at which this router will send hello messages. This is also known as the `HelloInterval` timer in OSPF6. Default value is `10`. This timer should be set to the same value for all routers.

A lower value will result in faster convergence times, but will consume more resources.

Dead Interval
Time, in seconds from `1-65535`, without communication from a neighbor on this interface before considering it dead. This is also known as the `RouterDeadInterval` timer in OSPF6. Default value is `40`. This timer should be set to the same value for all routers.

Retransmit Interval
The interval, in seconds from `1-65535`, at which this router will retransmit Link State Request and Database Description messages. This is also known as the `RxmtInterval` timer in OSPF6. Default value is `5`.

BFD

When set, enables *Bidirectional Forwarding Detection* for OSPF6 on this interface.

OSPF6 Area Configuration

OSPF6 Areas are managed at **Services > FRR OSPF6** on the **Areas** tab.

The **Areas** tab contains a list of current OSPF6 areas, if any, and controls to manage the entries (e.g. edit, delete). The  **Add** button creates a new area.

The remaining sections on this page cover the various options available when creating or editing an area entry.

Area Options

Area
A designation for the set of networks to which this router belongs. Can be any number capable of being expressed in dotted quad notation (IPv4 address) or as a 32-bit unsigned integer.

Description
Some text describing this area.

Area Type
Defines how this area behaves. For a list of all types and how they operate, see *Area Types*.

Note: The list of area types compatible with OSPF6 is shorter than OSPF. Consult the GUI drop-down for a current list of available area types.

Range Overrides

The **Ranges** list overrides the more specific ranges automatically determined from interfaces or by manual entry on the main configuration page. For example, this can be used to summarize routes or to prevent certain prefixes from being advertised.

When summarizing routes, instead of Type 1 (Router) and Type 2 (Network) LSAs, OSPF6 creates Type 3 Summary LSAs instead.

The options here can be repeated for multiple prefixes. Click  **Add** for each additional prefix.

Prefix
The prefix to override.

Do Not Advertise
Disable advertisement for this prefix.

Cost
Apply the specified cost to summarized routes for this prefix.

> **Warning:** An entry can either set **Do Not Advertise** or it can set **Cost**, but it **cannot** set both at the same time.

ABR Summary Route Filtering

Export List
Uses the given ACL to limit Type 3 summary LSA messages for intra-area paths that would otherwise be advertised. This behavior only applies if this router is the ABR for the area in question.

Import List
Similar to `export-list`, but for routes announced by other routers into this area.

Filter List (Out|In)
Similar to **Export List** and **Import List** but uses prefix lists instead of ACLs, and can work in either direction.

OSPF6 Status

The status for OSPF6 contains the same areas as the status for OSPF, but for IPv6 instead. See *OSPF Status* for more information.

See also:

For general FRR status information, see *FRR Status*.

30.10.5 Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is used to detect faults between two routers across a link, even if the physical link does not have failure detection capabilities. Even in cases where physical link issues occur and are detected, BFD can coordinate reaction to these failures rather than each component relying on its own failure detection methods.

FRR uses UDP as a transport for BFD between directly connected routers (single hop/next hop) as described in RFC 5880 and RFC 5881.

Each BFD session monitors one link. Multiple BFD sessions are necessary to detect faults on multiple links. BFD sessions must be manually configured between endpoints as there is no method for automated discovery.

When using BFD, both endpoints transmit "Hello" packets back and forth between each other. If these packets are not received within the expected time frame the link is considered down. Links may also be administratively configured as down and will not recover until manually changed.

BFD in FRR is currently compatible with BGP, OSPF, and OSPF6.

BFD is configured at **Services > FRR BFD**.

The **BFD** tab contains only the master **Enable** switch for BFD.

> **Note:** The BFD implementation in FRR does not currently implement authentication.

BFD Peers

A BFD **Peer** entry defines a relationship between this router and a peer so they can exchange BFD information and detect link faults.

Tip: For configurations with multiple peers which have similar settings, create a *BFD profile* with the common settings before creating peer entries.

BFD **Peers** are managed at **Services > FRR BFD** on the **Peers** tab.

The **Peers** tab contains a list of current BFD peers, if any, and controls to manage the entries (e.g. edit, delete). The



Add button creates a new peer.

The remaining sections on this page cover the various options available when creating or editing a peer entry.

Peer Configuration

Peer Address
: The remote BFD peer address. The local and remote peer IP addresses must use the same address family (either IPv4 or IPv6)

Description
: A text description of this BFD peer.

Profile
: A *BFD Profile* from which settings will be used, unless overridden on this entry.

Options

Multihop
: Expect packets with a TTL value less than `254` indicating there is more than one hop between peer addresses. Also listens on the multihop port, `4784`.

  In most cases this option is not necessary, but certain configurations may require it, such as running BGP on loopback interfaces instead of an interface directly connected to a peer.

  Warning: When using multi-hop mode echo-mode will not work, see RFC 5883 section 3.

Shutdown
: Enables or disables the peer administratively.

  When the peer is disabled in this way an "administrative down" message is sent to the remote peer.

Source Address/Interface

Interface
The interface on which to enable BFD.

Local Source Address
The local address used as a source for BFD packets, if it differs from the primary IP address on **Interface**.

Advanced Options

Detect Multiplier
A non-zero value that is, roughly speaking, due to jitter, the number of packets that have to be missed in a row to declare the session to be down. Must be between `2` and `255`.

The remote transmission interval will be multiplied by this value to determine the connection loss detection timer. The default value is `3`.

Receive Interval
The minimum interval, in milliseconds, at which this router can receive control packets from a peer. The default value is `300`.

Transmit Interval
The minimum transmission interval, in milliseconds, for the router to send BFD control packets to peers. The default value is `300`.

Echo Interval
The minimum interval, in milliseconds, at which this router can receive echos from a peer. The default value is `50`.

Echo Mode
Enables or disables echo transmission mode. By default, this mode is disabled.

> **Warning:** Echo mode is not compatible with multi-hop setups, see RFC 5883 section 3.

Tip: FRR documentation recommends that the transmission interval of control packets to be increased after enabling echo mode to reduce bandwidth usage.

For example, increase **Transmit Interval** from `300` to `2000`.

BFD Profiles

BFD profiles hold common sets of BFD settings which can be shared my multiple BFD peers. This can speed up the configuration process when there are numerous similar peers.

BFD **Profiles** are managed at **Services > FRR BFD** on the **Profiles** tab.

The **Profiles** tab contains a list of current BFD profiles, if any, and controls to manage the entries (e.g. edit, delete).

The  **Add** button creates a new profile.

The remaining sections on this page cover the various options available when creating or editing a profile entry.

Note: Most of the options in BFD profiles are the same as those for *BFD Peers*. This document only covers the options which are set on profiles only, not those shared with peers.

Profile Configuration

Name
The name of this profile. This name appears in the drop-down list for **Profile** on the *BFD Peers* configuration.

Passive
Marks sessions using this profile as passive.

A passive session will not attempt to start the connection and will wait for control packets from peer before it begins replying.

Profile Options

Minimum TTL
For multihop sessions only, configure a specific minimum expected TTL value for incoming BFD control packets.

This feature tightens the packet validation requirements to avoid receiving BFD control packets from other sessions.

The default value is `254` which means there is only one hop between this router and the peer.

Using BFD

For BFD to function fully, the BFD session status must be consumed by other interested parties. This can be BGP, OSPF, or OSPF6.

BGP
BFD can be enabled for specific BGP neighbors using the *BFD options for BGP neighbors*.

OSPF/OSPF6
BFD can be enabled for OSPF and OSPF6 by using the appropriate *OSPF interface BFD option* or *OSPF6 interface BFD option*.

30.10.6 Raw FRR Configurations

Rather than using the GUI, administrators can opt to manually manage the FRR configuration. This allows for more complex configurations than are possible in the GUI, but is also more prone to error and less capable of acting dynamically based on certain factors such as changing interface addresses.

Manual TCP MD5 Peers

When using TCP MD5 protection on BGP, this option defines MD5 passwords used to communicate with specific remote peers. This must be defined separately as it requires special handling in the operating system outside FRR.

Raw Configuration Management

The **Update Running** button reads the current FRR configuration into the **Running frr.conf** field on this page.

Raw Configuration Files

Saved frr.conf
: This is the FRR configuration which will be loaded by FRR when starting up.

Running frr.conf
: This is the configuration data currently active in FRR. Editing the running configuration will result in changes to the active FRR configuration, but is not stored long term.

  To save the configuration for later use (e.g. when booting up or restarting FRR), click the **Copy frr.conf Running to Saved** button.

  To refresh this based on the current information from the running instance of FRR, click **Update Running** in the **Raw Configuration Management** section.

30.10.7 Dynamic Routing Protocol Lists

Throughout the FRR package, certain options specify a routing protocol or source of routes. Currently, the following values can be found in these locations, but not every option appears in each area:

connected
: Routes for directly connected networks on up and active interfaces.

kernel
: Routes from the kernel, including static routes defined outside FRR and other non-dynamic routes.

FRR Static
: Static routes defined in the FRR configuration

bgp
: Routes obtained dynamically from BGP neighbors

ospf
: IPv4 routes obtained dynamically from OSPF neighbors

ospf6
: IPv6 routes obtained dynamically from OSPF6 neighbors

30.10.8 Multi-Path Routing

Current versions of pfSense® software include multi-path routing and this behavior is described in *Multi-Path Routing*.

The only way for multi-path routing to work is via FRR. If FRR has a route to the same destination across multiple paths, the traffic can be balanced across all the available paths. In most cases the weight of the balancing is even, resulting in Equal-cost multi-path routing (ECMP).

See also:

See *Multi-Path Routing Behavior* for details on how the OS makes balancing decisions for traffic across multiple paths.

30.11 HAProxy package

HAProxy is a free, very fast and reliable solution offering high availability, load balancing, and proxying for TCP, HTTP and HTTPS-based applications. It is particularly suited for websites struggling under very high loads while needing persistence or Layer7 processing.

Handling tens of thousands of connections is clearly realistic with today's hardware. Its mode of operation makes integrating it into existing architectures very easy and riskless, while still offering the possibility not to expose fragile web servers to the Net.

Refer to the following articles for more information on the listed topics:

See also:

- *Troubleshooting the HAProxy Package*

30.11.1 Package Variants

Two versions of the haproxy packages are available on pfSense® software:

HAProxy
: Tracks a stable version of FreeBSD port.

HAProxy-devel
: Uses haproxy-devel from FreeBSD ports and loosely tracks a HAProxy development branch.

  New features are added to the HAProxy-devel package first then later copied over the HAProxy package.

For info about HAProxy versions, see: https://github.com/PiBa-NL/pfsense-haproxy-package-doc/wiki

30.11.2 Recent Changes

See GitHub log for recent changes:

- https://github.com/pfsense/FreeBSD-ports/commits/devel/net/pfSense-pkg-haproxy
- https://github.com/pfsense/FreeBSD-ports/commits/devel/net/pfSense-pkg-haproxy-devel

30.11.3 Known issues

HAProxy has a list of known bugs by branch and by version.

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.11.4 Differences between this package and HAProxy used directly

HAProxy defines five main sections in its configuration.

Global
: Defines options that process-wide and often OS-specific.

Defaults
: Sets default parameters for all other sections following its declaration.

Frontend
: Describes a set of listening sockets accepting client connections.

Backend
: Describes a set of servers to which the proxy will connect to forward incoming connections.

Listen
: Defines a complete proxy with its frontend and backend parts combined in one section. It is generally useful for TCP-only traffic.

In the pfSense software package, tabs exist to define "frontends" and "servers" but the resulting configuration is actually made up completely of listen sections. This is okay for the most part, but it does prevent advanced usages that need to refer to several backends and the like.

In HAProxy, a single server directive can be made with a blank port, and it will listen on all the ports of the frontend that it is assigned to. The package GUI implies that this will be the case by leaving the port blank.

What actually gets generated instead is a single server directive for each port that the frontend is listening on. This is an important difference when the ports that are being listened on are not interchangeable. Example:

Define a front end for SMTP connections listening on ports 25 and 465. The server is listening on both of those ports, but 25 does not accept SSL/TLS and 465 does. When someone connects to the proxy on port 25, they should get connected to the server on port 25, and when they connect on 465, they get connected to the server on port 465.

In a standard HAProxy configuration where the frontend is set to listen on both ports and a single server directive is made with no port, it will operate the expected way.

In pfSense software, two server directives will be generated; one for each port. HAProxy will not send connections the expected way. It will load balance between them, regardless of whether the frontend and server ports match.

Therefore, in pfSense software a separate frontend must be created for this, as they are essentially different services. Listen on port 25 and 2525, and it doesn't matter whether someone connected on one port gets directed to the other, then they can be combined.

Splitting the servers up by port also means that a separate entry will exist for each one in the stats page, but the port will not be shown. In an HAProxy configuration where a single server directive has no ports and effectively handles multiple (due to inheriting from the frontend) it will only show up in the stats once.

30.12 iperf package

iperf is a tool used for network throughput testing.

30.12.1 Usage

iperf running on pfSense® software is NOT a suitable way of testing firewall throughput, as there is a significant difference between performance of traffic *initiated or terminated on the firewall* and traffic *traversing the firewall*. There are many suitable uses for iperf running on pfSense software, but testing the throughput capabilities of the firewall is not one of them.

Uses for iperf on pfSense software

- Measuring throughput from the internal network to the inside of the firewall
  - Useful in scenarios where portions of the internal network are behind links slower than the firewall network interface. Examples include testing the throughput of a wireless network or private WAN network connected to a router inside the network.
- Testing end to end throughput between two firewalls on the Internet
- Anything else where performance measurement excluding throughput capabilities of the firewall are desirable.

30.12.2 Additional Resources

- iperf Man Page
- iperf - The Easy Tutorial

30.12.3 Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.12.4 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.13 IPsec Export Package

The IPsec Export package generates client configurations for mobile IPsec, making it easier to configure remote access clients. This package is available on pfSense® Plus software.

The IPsec Export package contains an IPsec Profile export page for Apple devices and an IPsec Export page for Windows. Both pages work similarly and give administrators a few extra options to control client behavior.

The package works with most types of mobile IPsec configurations, with some exceptions depending upon settings.

This utility checks configured Mobile Phase 1 and Phase 2 entries and attempts to locate a set of parameters which are compatible with clients. It uses the first match it finds, so order choices in the Phase 1 and Phase 2 list appropriately or manually edit the resulting profile or script as needed.

Note: Certain settings are not compatible with Apple and Windows clients. In these cases, the package will print a notice that it is not possible to export a configuration.

See the Apple Configuration Profile Reference Documentation for details about compatible settings.

For a full list of settings compatible with Windows clients, see the Microsoft Documentation for `Set-VpnConnectionIPsecConfiguration`.

30.13.1 Export Settings

When exporting an IPsec configuration, the following options are available to fine-tune the values put into the generated configuration.

VPN Name
The name of the VPN as seen by the client in their network list. This name is also used when creating the filename of the files exported by the package. It is pre-filled with some basic information, such as the firewall hostname, but it can be customized.

Server Address
Select the server address to be used by the client. This list is generated from the SAN entries on the server certificate.

The hostname used by the client to connect to the server must exist in DNS, and it must be present in the server certificate SAN list for the client to properly validate the certificate.

Set to *Custom Hostname* to fill in a hostname other than one shown in the list.

Custom Hostname
A text field for a custom fully qualified domain name to which the client can connect. As with the **Server Address**, this must exist in DNS and be in the server certificate SAN list.

VPN Client (Apple)
The user for which the package will generate a configuration. Depending on the Mobile IPsec Phase 1 settings, this could either be a user or a TLS certificate.

When using certificates, the list contains certificates which were signed by the CA selected on the mobile IPsec Phase 1 **IPsec Peer Certificate Authority**.

Note: This differs from Windows because Windows can prompt for a username, but Apple requires it to be present in the profile.

TLS User Certificate (Windows)
The TLS user certificate to include in the exported configuration, if needed. This field is only visible when the Mobile IPsec Phase 1 settings require a client certificate (e.g. EAP-TLS).

30.13.2 Export a Client Configuration

The process to export a client for an existing *Mobile IPsec* configuration varies slightly for Apple and Windows.

Apple

- Navigate to **VPN > IPsec Export: Apple**
- Configure the settings as described in *Export Settings*
- Click **View** to display the generated configuration profile
- Review the profile contents and confirm it is acceptable
- Click **Download** to download the configuration profile

Apple Client Configuration

Visit the Apple Configurator Site for details about creating and using profiles. The process varies between iOS and macOS.

Windows

- Navigate to **VPN > IPsec Export: Windows**
- Configure the settings as described in *Export Settings*
- Click **View** to display the generated PowerShell script
- Review the script contents and confirm it is acceptable
- Click **Download** to download a ZIP archive containing the PowerShell script and the required certificates.

If the **Network List** option is active on the **Mobile Clients** tab in IPsec settings, the script will include parameters to set up Split Tunneling on the client as well as commands to configure routes on the VPN for networks configured in the mobile Phase 2 entries.

Windows Client Configuration

On the client system, unzip the configuration archive and run the script. The commands in the PowerShell script will import certificates and set up the VPN on the client workstation.

Running PowerShell scripts on Windows is disabled by default. If scripting is disabled, the commands may be copied and pasted into a PowerShell prompt.

See also:

Local policies may override that behavior. See the PowerShell Execution Policies Documentation for details.

> **Warning:** Some commands may require Administrator access, such as importing the CA certificate. Run these commands at an Administrator-level PowerShell prompt or use an alternate method.

30.14 Using LLDP on pfSense software

The lldpd daemon allows pfSense software to utilize LLDP.

30.14.1 Install

- Navigate to **System > Packages**, **Available Packages** tab
- Search for `lldp` or find it in the list
- Click **Install** and **Confirm**

The package adds a menu entry under **Services > LLDP**.

30.14.2 Setup

The settings are under **Services > LLDP**, **LLDP Settings** tab.

At a minimum, check **Enable**, pick **Interfaces**, and then set other options as desired.

30.14.3 Seeing neighbors

Navigate to **Services > LLDP**, **LLDP Status** tab. The neighbors will be printed in the lower box.

30.15 Netgate Firmware Upgrade Package

The Netgate Firmware Upgrade package provides a mechanism to update firmware on certain Netgate hardware models.

The mechanism used by the package varies by hardware and may be Coreboot, Blinkboot, or other types of firmware.

When installed, the package is located at **System > Netgate Firmware Upgrade**.

30.15.1 Checking for Upgrade

To see if a Netgate device has a firmware update available:

- Navigate to **System > Netgate Firmware Upgrade**
- Compare the **Current Firmware Version** and **Latest Firmware Version**

If the latest version is newer than the current version, the package will offer the option to upgrade.

30.15.2 Upgrading Firmware

When a firmware upgrade is available, upgrade as follows:

- Navigate to **System > Netgate Firmware Upgrade**
- Click **Upgrade and Reboot**
- Click **OK** to confirm the upgrade
- Read the output on the screen for any device-specific or version-specific instructions.

- Wait for the device to complete the upgrade and reboot

30.15.3 Device-Specific Notes

The following sections contain information relevant only to specific models of Netgate hardware.

Netgate 6100

Upgrades to the firmware (Blinkboot) on the Netgate 6100 which also upgrade the microcontroller code require a power cycle after firmware upgrade. This is due to the microcontroller always being on while the device has power. Thus, new microcontroller code will activate until after the device goes through a cold boot (power cycle, not reboot):

- Connect to the device console
- Run the firmware upgrade
- Wait for the firmware upgrade and automatic reboot to complete
- Halt the device gracefully using either the GUI or console menu (*Halting and Powering Off the Firewall*)
- Monitor the console and wait until the device has halted

  Alternately, wait until a few moments after all the lights on the device (front and back) go out.
- Remove power from the device (e.g. unplug or remove power from PDU port)
- Wait a few seconds
- Reapply power to the device (e.g. plug in or apply power to PDU port)
- Monitor the console during the boot sequence for any errors
- Wait for the device to complete its boot sequence

Note: While it is possible to manually remove power after the initial firmware upgrade completes in one pass, it can be problematic to time it properly. Thus, it is safer to follow the procedure as stated to ensure success.

Netgate 7100

Updating the firmware (Coreboot) on the Netgate 7100 causes the device to lose customizations to its default boot order.

In the default boot order if the device has multiple disks (e.g. On-board eMMC plus an add-on SATA device), the device will prefer to boot from the add-on disk rather than the on-board eMMC.

For most users this behavior is correct and expected, but if someone with this configuration has manually opted to boot from eMMC instead of an add-on disk, this boot order preference change must be reconfigured after updating the firmware.

30.16 Nmap package

Nmap is a powerful network scanner that provides port scanning, OS and service identification, and more.

30.16.1 Usage

After installing the package, Nmap will be available at **Diagnostics > nmap** as well as in the shell (SSH or Console).

The Nmap package also installs an OUI database that is used by the pfSense® software GUI to display manufacturer names on pages that list MAC addresses, such as the *ARP Table*, and *DHCP Leases*.

30.16.2 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.17 NUT package

The NUT package provides a way to monitor an Uninterruptable Power Supply (UPS) using Network UPS Tools (NUT) on pfSense® software.

After installation, configure the package at **Services > UPS**.

See also:

Visit the NUT Package forum thread for assistance with this package.

30.17.1 Troubleshooting

If NUT will not start after configuration, it may be due to incorrect settings. The package GUI is unable to validate combinations of settings to ensure they are viable. Check the System log from the GUI from **Status > System Logs** for log entries starting with `nut:`. The culprit is likely explained there, such as selecting a cable for a driver type that does not need (nor permit) the cable selection.

If the system log does not offer adequate information, for example if it prints a generic error such as:

```
nut: Service failed to start: check configuration
```

Log in via SSH and choose option `8` and run the following command:

```
/usr/local/etc/rc.d/nut.sh start
```

Configuration errors will be displayed in the output if any are found.

If the daemon still will not start, users have also reported that rebooting the device after connecting the UPS and configuring the service may help resolve the problem.

See also:

The pfSense-packages Redmine project contains a list of known issues with this package.

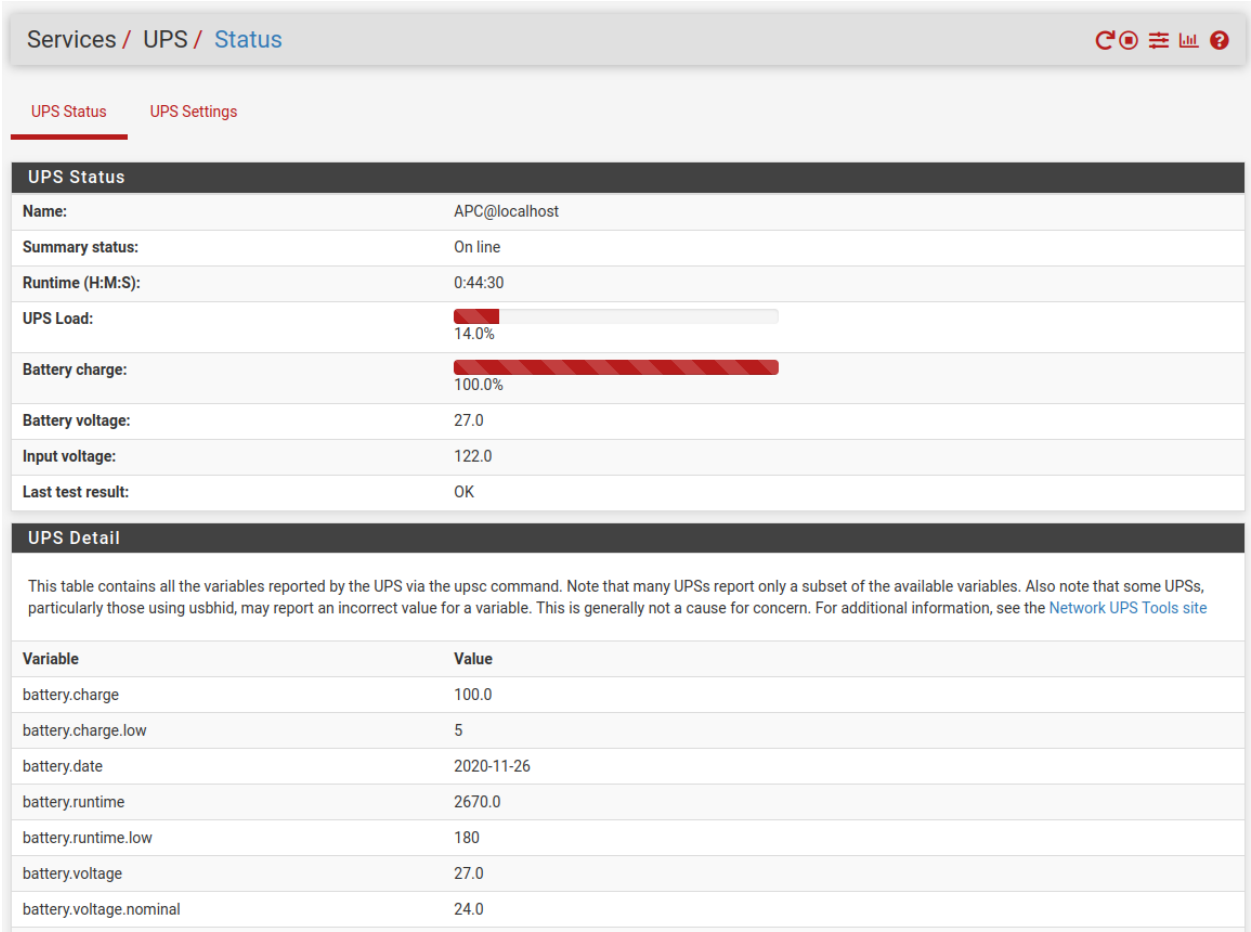


Fig. 4: NUT status screen

30.18 Open VM Tools package

This package installs VMware Tools for pfSense® software, using the Open VM Tools available from VMware. It is built using the open-vm-tools-nox11 FreeBSD port.

30.18.1 Usage

There is no GUI, the services are automatically started at boot time.

30.18.2 Verifying functionality

There are two portions of this package, the **vmware-guestd** process and the kernel modules.

Verifying vmtoolsd

Navigate to **Diagnostics > Command**, and in **Execute shell command**, enter the following:

```
ps uxawww | grep vmtoolsd
```

Output similar to the following will be shown:

```
$ ps uxawww | grep vmtoolsd
1026  ??  Ss     0:20.84 /usr/local/bin/vmtoolsd -c /usr/local/share/vmware-tools/tools.
↪conf -p /usr/local/lib/open-vm-tools/plugins/vmsvc
19157  ??  SN     0:00.00 sh -c ps ax|grep vmware
19159  ??  RN     0:00.00 grep vmware (sh)
```

As long as `vmtoolsd` is shown in the output, it is working.

30.18.3 Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.18.4 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.19 OpenVPN Client Export Package

The easiest way to configure an OpenVPN client on most platforms is to use the OpenVPN Client Export Package on pfSense® software.

30.19.1 Installing the Export Package

Install the OpenVPN Client Export Utility package as follows:

- Navigate to **System > Packages**, **Available Packages** tab
- Locate the **OpenVPN Client Export** package in the list
- Click  **Install** next to that package listing to install
- Click  **Confirm** to confirm the installation

30.19.2 Using the Export Package

Once installed, the package is located at **VPN > OpenVPN**, on the **Client Export** tab. That page presents several options which control the behavior of exported clients. The firewall can optionally save selections on this page as new defaults for future use.

Settings

The options for the package include:

OpenVPN Server

Remote Access Server
The OpenVPN server instance for which the package will export a client.

The list includes only Remote Access mode OpenVPN servers. If there is only one OpenVPN remote access server there will only be one choice in the list. The list will be empty if the firewall has no OpenVPN servers set to a Remote Access mode.

Client Connection Behavior

Host Name Resolution
Controls the format of `remote` directive entries the package uses in client configuration files.

Interface IP Address
Uses the OpenVPN server interface IP address directly.

This is typically the best choice for installations with a static IP address on WAN.

Automagic Multi-WAN IPs
Searches all port forwards and automatically creates `remote` directives for all port forwards that target the OpenVPN server interface address and port. It uses the destination IP address on the port forward in the `remote` directives for the client configuration, so the clients will use the correct external address(es).

This option is useful when redirecting connections using port forwards for deployments that utilize, multi-WAN, multiple ports on the same WAN, or servers which bind to *Localhost*.

Automagic Multi-WAN DDNS Hostnames
Similar to the previous option, but instead of using IP addresses it uses the first Dynamic DNS entry it finds which matches the chosen destination.

This is useful if one or more of the WANs involved in OpenVPN has a dynamic IP address.

Installation Hostname
Uses hostname assigned to the firewall at **System > General Setup** for the `remote` directive in the client configuration.

Note: The hostname must exist in public DNS, so clients can resolve it by name.

Dynamic DNS Hostname Entries
The GUI includes entries in the selection list for each Dynamic DNS hostname in the firewall configuration.

These are typically the best choice for a server on a single WAN with a dynamic IP address.

Other
Displays a **Host Name** field for a custom hostname or IP address.

This is useful when the firewall is behind one or more layers of NAT, which may make the external address difficult or impossible for the package to discover. It is also useful for situations where the firewall does not manage the DNS entry (static or dynamic) for a hostname.

Verify Server CN
Controls how the client verifies the identity of the server certificate.

The package places the CN of the server certificate in the client configuration, so that if another valid certificate pretends to be the server with a different CN, it will not match and the client will refuse to connect.

Automatic - Use verify-x509-name where possible
Uses the `verify-x509-name` directive in OpenVPN to set a specific string the client will expect to match the common name on the server certificate.

Do not verify the server CN
Disables client verification of the server certificate common name. This is not a secure, as the client will accept any server certificate signed by the CA.

Block Outside DNS
Makes Windows 10 clients block access to DNS server except across OpenVPN while connected, forcing clients to use only VPN DNS servers.

This is only relevant on Windows 10 clients using OpenVPN version 2.3.9 and later as they are the only clients prone to leak DNS requests in this way. The option has no effect on other platforms and they will ignore the directive.

Legacy Client
Controls whether the package uses OpenVPN 2.5.x and later directives in the configuration, or directives for older legacy clients.

When set, the package forms configurations using older options accepted by outdated clients such as OpenVPN 2.4.x and avoids newer directives only understood by current versions of OpenVPN.

Silent Installer
When set, the package adds flags to the Windows installer which allows for silent and unattended client deployment.

Note: The installer generated by the package is not signed. As such, deploying the installer executable may require special software.

Use Random Local Port
Controls whether the client binds to a specific local port.

When checked, which is the best practice, the client will use a random source port. When unchecked, two OpenVPN connections cannot be run simultaneously on the client device as they would attempt to use the same port.

Certificate Export Options

PKCS#11 Certificate Storage
Configures the client to use PKCS#11 storage device (e.g. cryptographic token, HSM, smart card) instead of local files.

PKCS#11 Providers
The client-side path to PKCS#11 provider(s) files (e.g. DLL, module). Separate multiple entries with a space.

PKCS#11 ID
The object ID on the PKCS#11 device.

Use Microsoft Certificate Storage
Creates the installer in such a way that it places the CA and user certificate in Microsoft certificate storage rather than using inline entries or files directly.

Use a password to protect the PKCS#12 file contents or key in Viscosity bundle
Controls whether the package protects the certificates and keys supplied to the client with a password.

When checked, the GUI displays fields to enter and confirm a password.

Note: If the OpenVPN server requires user authentication this will cause users to see two different password prompts when loading the client: One to decrypt the keys and certificates and another for OpenVPN user authentication upon connecting.

Proxy Options

Use Proxy
If the client is located behind a proxy, check this option. The GUI displays several fields to configure the proxy. Supply a Proxy **Type**, **IP Address**, **Port**, and **Proxy Authentication** with credentials (if the proxy server requires authentication).

Advanced

Additional configuration options
Any extra *custom OpenVPN directives* for the package to include in the client configuration.

This is roughly equivalent to the **Advanced options** box on the OpenVPN configuration screens, but from the perspective of the client.

Save as default

The package does not save the settings on this page by default, so they must be set each time the package is used. The **Save as default** button stores the current package settings as new default values for future use.

OpenVPN Clients List

OpenVPN Clients contains a list of potential clients that the package can export. The contents of the list depend on the server configuration and which users and/or certificates are present on the firewall.

The following list describes how the server configuration style affects the list in the package:

Remote Access (SSL/TLS)
The list contains entries for user certificates signed by the same CA as the OpenVPN server

Remote Access (SSL/TLS + User Auth – Local Users)
The list contains entries for local users which also have an associated certificate signed by the same CA as the OpenVPN server.

Tip: This is the type of server configured by the OpenVPN wizard.

Remote Access (SSL/TLS + User Auth – Remote Authentication)
The list contains user certificates signed by the same CA as the OpenVPN server. Because the users are remote, the package cannot determine the usernames, so the package assumes that the username is identical to the common name of the certificate.

Remote Access (User Auth – Local Users or Remote Authentication)
The list contains a single configuration entry for all users since there are no per-user certificates or settings. Each client uses the same configuration.

The **Search** box above the client list filters the list to only entries which match a given string. This feature is particularly useful on firewalls with many clients, allowing administrators to quickly locate a specific entry.

Note: If the list contains no entries, or if a specific entry is missing from the list, then either the user does not exist or the user does not have an appropriate certificate.

See *Local Database* for the correct procedure to create a user and certificate.

Client Install Package Types

Each client entry contains numerous options which export the configuration and associated files in different ways. Each choice accommodates a different potential type of client.

Inline Configurations

These choices export a single configuration file with the certificates and keys inline. This format is ideal for use on all platforms, especially Android and iOS clients, or for manually copying a configuration to a device which already has a client installed.

This option will work for any client type based on OpenVPN version 2.1 or newer.

Most Clients
Exports a configuration which is usable by any standard OpenVPN client on platforms such as Windows, macOS, or BSD/Linux.

Tip: This format works well with Tunnelblick on macOS. Download the inline configuration file and drag it into the configuration folder for Tunnelblick.

Android
Exports a configuration for the Android OpenVPN client mentioned in *Installing the OpenVPN Client on Android*, which uses a different set of directives and omits certain features.

OpenVPN Connect (iOS/Android)
Exports a configuration for the OpenVPN Connect client on iOS or Android described in *Installing the OpenVPN Client on iOS*.

Bundled Configurations

Archive
Exports a ZIP archive containing the configuration file, the server TLS key (if it has one), and a PKCS#12 file which contains the CA certificate, client key, and client certificate. This option is usable with Linux clients, Tunnelblick, Windows, and many others when configuring the files manually.

File Only
Exports only the basic configuration file, no certificates or keys. This would mainly be used to update the configuration file for systems which used the ZIP archive or executable installer before, without downloading the other associated files. For example, this is useful if the user authentication did not change but some aspect of the server configuration or package preferences is different.

Windows Installers

The Windows installer options export a simple-to-use executable installer file which contains the OpenVPN client software plus the configuration data. The installer runs like the normal Windows OpenVPN client installer, but it also copies all the settings and certificates the clients needs when it connects to the VPN.

See also:

See *Installing the OpenVPN Client on Windows* for notes on how to install and run the Windows client.

Currently, there are four options available:

Current Windows Installer - 64-bit/32-bit
Export an installation bundle for the latest available version of the Windows OpenVPN client, currently version 2.5.x. The installer is available in 64-bit and 32-bit varieties. This is the best version to use where possible, and it works on Windows 11, Windows 10, and more.

Legacy Windows Installer - 10/2016/2019
Export an installation bundle for the legacy version of the Windows OpenVPN client, currently version 2.4.x. This version is specific to Windows 10 and similar vintage versions of Windows.

Legacy Windows Installer - 7/8/8.1/2012r2
Export an installation bundle for the legacy version of the Windows OpenVPN client, currently version 2.4.x. This version is specific to Windows 7/8.x and similar vintage versions of Windows.

Note: Users must click next/finish **all the way through** the installation process. Do not click cancel or X out the installer at any step, as this can leave the client system with the client installed but without an imported configuration.

Viscosity Bundle

This exports the configuration similar to the archive method, but in a format used by the Viscosity OpenVPN client for macOS and Windows. Install the Viscosity client separately, then export this bundle and click it to import the configuration.

SIP Phone archives

The package can export configurations formatted for several popular VoIP phones which natively include OpenVPN. Notable examples are the Yealink T28 and T38G, and SNOM phones.

The process to install and configure OpenVPN on the phone varies by model, check the manufacturer documentation for more information.

The package only displays these options for OpenVPN servers set to SSL/TLS only *without* authentication.

Note: The phone must have a proper clock setup and/or NTP server otherwise the certificates will fail to validate, and the VPN will not connect.

Warning: Typically, these handsets are only compatible with the SHA1 certificate hash. Ensure the CA, server certificate, and client certificates all use *SHA1*, or they may fail. They may also only be compatible with a limited set of encryption algorithms such as AES-128-CBC. Consult the phone documentation for details, and use the strongest possible combination of options.

30.20 OpenVPN Client Import Package

Note: This package is only available on Netgate pfSense® Plus software.

The OpenVPN client import package can take a unified OpenVPN client configuration file as exported by an OpenVPN server and automatically turn it into an OpenVPN client instance on pfSense Plus software. The unified OpenVPN configuration file format includes all the certificates and keys required for the connection, allowing the client instance to be created with minimal effort.

In many cases the newly imported client instance starts and passes traffic on completion of the import, but in some cases adjustments must be made to the imported client configuration by editing the resulting OpenVPN client instance.

The package can be installed using the *Package Manager* on pfSense Plus software. Once the package is installed, it can be accessed at **VPN > OpenVPN** on the **Import** tab.

30.20.1 How it Works

The import process attempts to read the configuration file and map directives from the file to their equivalent settings in pfSense Plus software. Unknown directives are placed into the **Custom options** area in the resulting client instance.

If the configuration being imported contains certificates, the import package will create appropriate CA and certificate entries if they do not already exist.

Note: If the configuration requires certificates, but they are not present in the imported configuration file, they can be manually imported in the certificate manager and then manually selected in the OpenVPN client instance after it has been imported.

Once the import process is complete, the new client is stored and, if it is enabled and has a complete configuration, the client is immediately started.

30.20.2 Imported OpenVPN Client Configuration

When importing a configuration there are several options specific to pfSense Plus software which cannot be automatically determined from the imported configuration. These must be filled in manually before the import process can be completed.

These options are equivalent to their counterparts in the *OpenVPN Configuration Options*. Consult that document for additional details on these settings.

Config File
: The OpenVPN configuration file (e.g. `<name>.ovpn`) to import.

  The OpenVPN client configuration file can be from another instance of pfSense software, a VPN provider, or other OpenVPN compatible server so long as it uses the standard OpenVPN configuration format.

Disabled
: When set, the client will be marked as disabled on import so it will not start automatically.

Server Mode
: Chooses between whether this client is connecting to an SSL/TLS server with certificates, or to a shared key server.

Name
A descriptive name for this client instance.

Interface
The firewall interface to be used by this client instance for outbound connections. In most cases this will be **WAN** but may also be another interface, or a virtual IP address.

Username
The username to use if the OpenVPN server requires a username and password. May be left blank if the server does not require user authentication.

Password
The password to use if the OpenVPN server requires a username and password. May be left blank if the server does not require user authentication.

30.20.3 Client Import Example

The process to import a client generally follows this format:

- Obtain an OpenVPN configuration file in inline format from the OpenVPN server (e.g. `username.ovpn`)

Note: If the server is also running pfSense software, use the *OpenVPN Client Export Package* and download the inline configuration using the **Most Clients** button.

- Navigate to **VPN > OpenVPN**, **Import** tab on the client firewall
- Click **Browse** in the **.ovpn config file** field and select the configuration file obtained from the server (e.g. `username.ovpn`)
- Fill in the other options as described in *Imported OpenVPN Client Configuration*
- Click **Import**

At that point the client instance will be created and started automatically. If the configuration was incomplete or needs other changes, then do so as follows:

- Navigate to **VPN > OpenVPN**, **Clients** tab
- Find the newly imported client in the list and click  on its row
- Make final adjustments needed
- Click **Save**

See also:

See also: *OpenVPN Configuration Options*

30.21 pfBlocker-NG Package

pfBlocker-NG introduces an enhanced alias table feature to pfSense® software.

This package enables users to:

- Assign many IP address URL lists from sites like I-blocklist to a single alias and then choose a rule action.
- Block countries and IP address ranges.
- Use native functions of pfSense software instead of file hacks and table manipulation.

Features include:

- Geographical/Country Blocking
- IP block lists
- Dashboard widget
- XMLRPC Sync
- Frequently updated lists
- Many options to control what to block and how to block
- Network lists can be used in custom rules

30.21.1 General Setup

Set the interfaces to be monitored by pfBlocker-NG (both inbound and outbound), where the inbound is the Internet connection.

To prevent devices or users from accessing sites in the selected countries/IP addresses, select local interfaces under **outbound**.

30.21.2 Setting up Lists

This is the IPBlocklist feature, enter IP addresses here to specifically block. It must be in the file format or CIDR. Create a list for each type of action to be taken by pfBlocker.

Options are:

Deny Both
Will deny access on Both directions.

Deny Inbound
Will deny access from selected lists to the local network.

Deny Outbound
Will deny access from local users to IP address lists selected to block.

Permit Inbound
Will allow access from selected lists to the local network.

Permit Outbound
Will allow access from local users to IP address lists selected to block.

Disabled
Will just keep selection and do nothing to selected Lists.

Alias Only
Will create an alias with selected Lists to help custom rule assignments.

The rest of the tabs (except sync) specify the other lists included with the package. They are separated by continent except for the spammer list which contains countries from around the globe that are known to harbor spammers.

Sync tab configures pfBlocker to sync its configuration to other pfSense devices.

30.21.3 Available lists

Spamhaus
DROP and EDROP.

- http://www.spamhaus.org/drop/drop.txt
- http://www.spamhaus.org/drop/edrop.txt

DShield
Most Active Attacking IPs.

- http://feeds.dshield.org/top10-2.txt

iblocklist.com
A number of lists are available.

- http://www.iblocklist.com/lists.php

30.21.4 FAQ

I'm getting memory errors while applying pfBlocker lists, how to fix this?
Increase table size to avoid memory errors in Advanced settings.

I can't see any pfBlocker rules applied, what's wrong?
pfBlocker requires at least one firewall entry (any interface) for it to be active. One way to verify is to check the front page widget.

pfBlocker always moves its rules to the top, how can I stop this?
Change rule action to Alias only and then apply custom rules using pfBlocker aliases with an arbitrary sequence.

How can I apply pfBlocker lists in floating rules?
Aliases are used for customized filter entries and float rules.

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.22 Siproxd package

Warning: For many modern SIP configurations this package is unnecessary. Please only install and use siproxd if it is absolutely required.

Siproxd is a proxy/masquerading daemon for SIP. It handles registrations of SIP clients on a private IP network and performs rewriting of the SIP message bodies to make SIP connections possible via a masquerading firewall. It allows SIP phones and soft phones (like kphone, linphone) to work behind an IP masquerading firewall or router.

The most useful thing siproxd does is allow multiple phones to use a static source port of 5060 when registering to the same remote PBX. **If the PBX is local, this package should not be installed or used**. Most remote PBX systems are OK with the phones having random source ports, but that was not the case many years ago. Unless the remote PBX is absolutely strict about the 5060 source port requirement for each phone, this package is not needed.

30.22.1 Install Siproxd

In the pfSense® software GUI, navigate to **System > Packages**:

- Find the **siproxd** package.
- Click  to install, and confirm installation.
- Wait for it finish.

30.22.2 Configure siproxd

Under **Services > siproxd**:

- **Inbound interface** will generally be `LAN`.
- **Outbound interface** will generally be `WAN`.
- Fill in any other values (where appropriate).
- **RTP port range (lower)** should be an even number.
- Click **Save**.

30.22.3 Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.22.4 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.23 IDS / IPS

pfSense® software can act in an Intrusion Detection System (IDS) / Intrusion Prevention System (IPS) role with add-on packages like Snort and Suricata.

Note: The Snort and Suricata packages share many design similarities, so in most cases the instructions for Snort carry over to Suricata with only minor adjustments.

30.23.1 Snort

Configuring the Snort Package

Snort is an intrusion detection and prevention system. It can be configured to simply log detected network events to both log and block them. Thanks to OpenAppID detectors and rules, Snort package enables application detection and filtering. The package is available to install in the pfSense® software GUI from **System > Package Manager**. Snort operates using detection signatures called rules. Snort rules can be custom created by the user, or any of several pre-packaged rule sets can be enabled and downloaded.

The Snort package currently offers these pre-packaged rules:

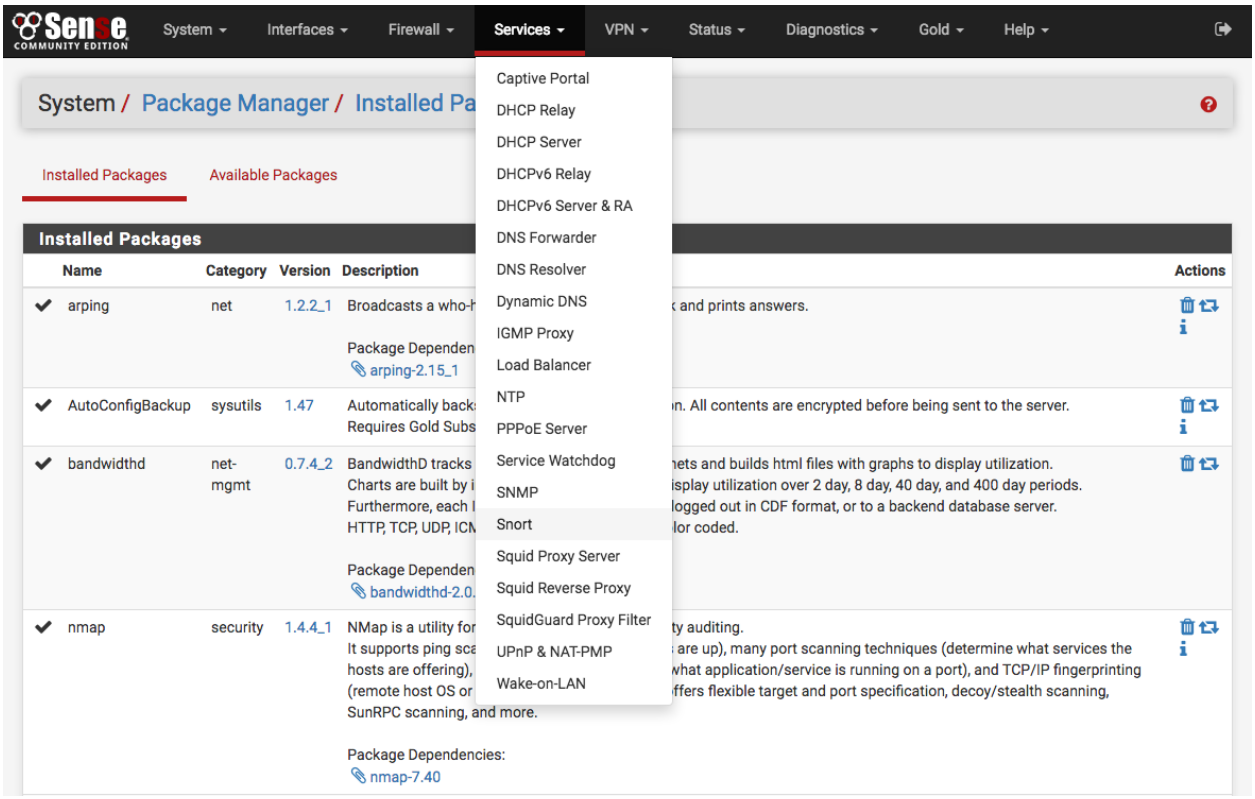
- Snort VRT (Vulnerability Research Team) rules
- Snort GPLv2 Community Rules
- Emerging Threats Open Rules
- Emerging Threats Pro Rules
- OpenAppID Open detectors and rules for application detection

The Snort GPLv2 Community Rules and the Emerging Threats Open Rules are both available for free with no registration required. The Snort VRT rules are offered in two forms. One is a registered-user version which is free, but requires registration at http://www.snort.org. The registered-user free version only provides access to rules that are 30-days old or more in age. A Snort VRT paid subscription can be purchased, and it offers twice-weekly (and sometimes more frequent) updates to the rules. The Emerging Threats Pro rules are offered to paid subscribers only and offer almost daily updates to address fast-changing threats.

The best practice is to obtain a paid subscription from Snort or Emerging Threats in order to download the most current rules. This is highly recommended for commercial applications.

Launching Snort configuration GUI

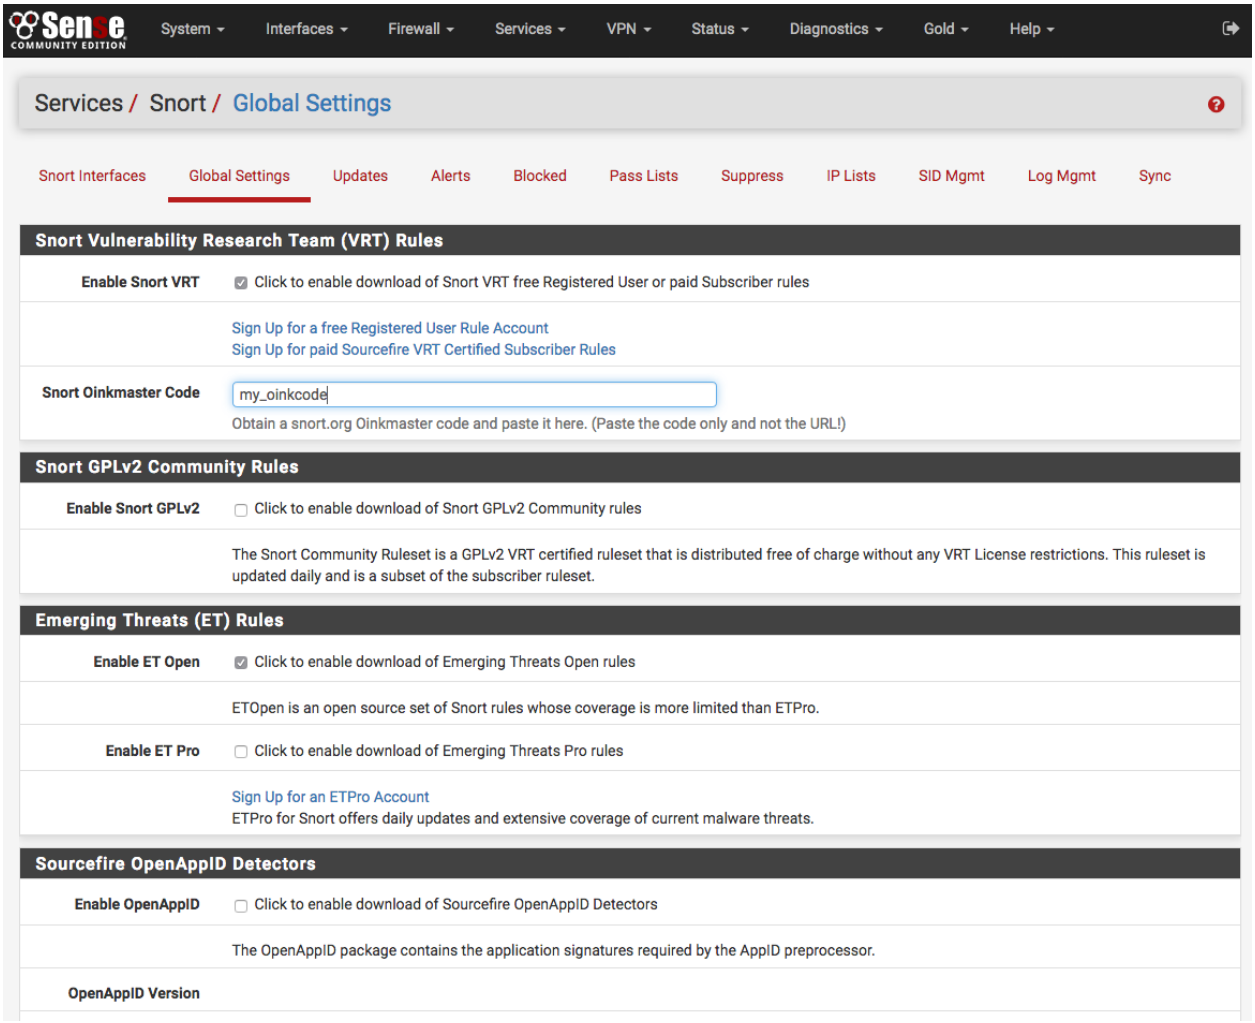
To launch the Snort configuration application, navigate to **Services > Snort** from the menu in the GUI.



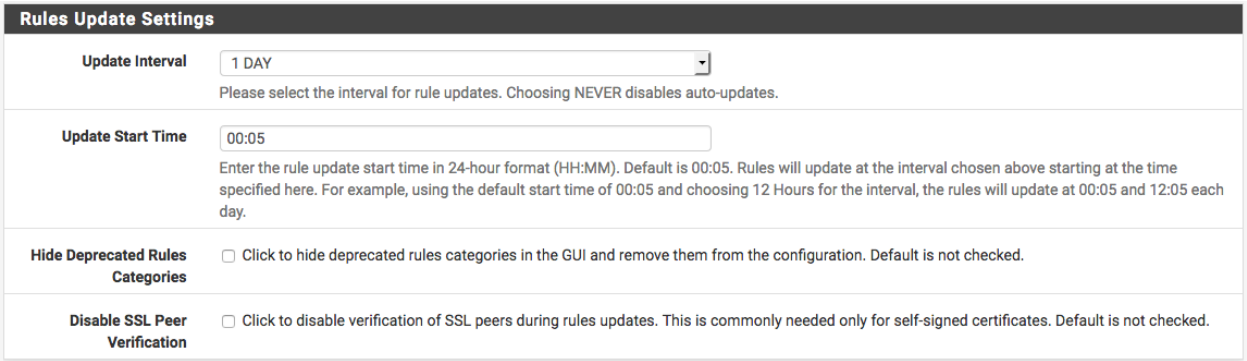
Setting up Snort package for the first time

Click the **Global Settings** tab and enable the rule set downloads to use. If either the Snort VRT or the Emerging Threats Pro rules are checked, a text box will be displayed to enter the unique subscriber code obtained with the subscription or registration.

More than one rule set may be enabled for download, but note the following caveats. If a paid subscription is available for the Snort VRT rules, then all the Snort GPLv2 Community rules are automatically included within the file downloaded with the Snort VRT rules; therefore, do not enable the GPLv2 Community rules if a paid-subscriber account is used for the Snort VRT rules. All the Emerging Threats Open rules are included within the paid subscription for the Emerging Threats Pro rules. If the Emerging Threats Pro rules are enabled, the Emerging Threats Open rules are automatically disabled.



Once the desired rule sets are enabled, next set the interval for Snort to check for updates to the enabled rule packages. Use the **Update Interval** drop-down selector to choose a rule update interval. In most cases every *12* hours is a good choice. The update start time may be customized if desired. Enter the time as hours and minutes in 24-hour time format. The default start time is *3* minutes past midnight local time. So with a *12*-hour update interval selected, Snort will check the Snort VRT or Emerging Threats websites at 3 minutes past midnight and 3 minutes past noon each day for any posted rule package updates.

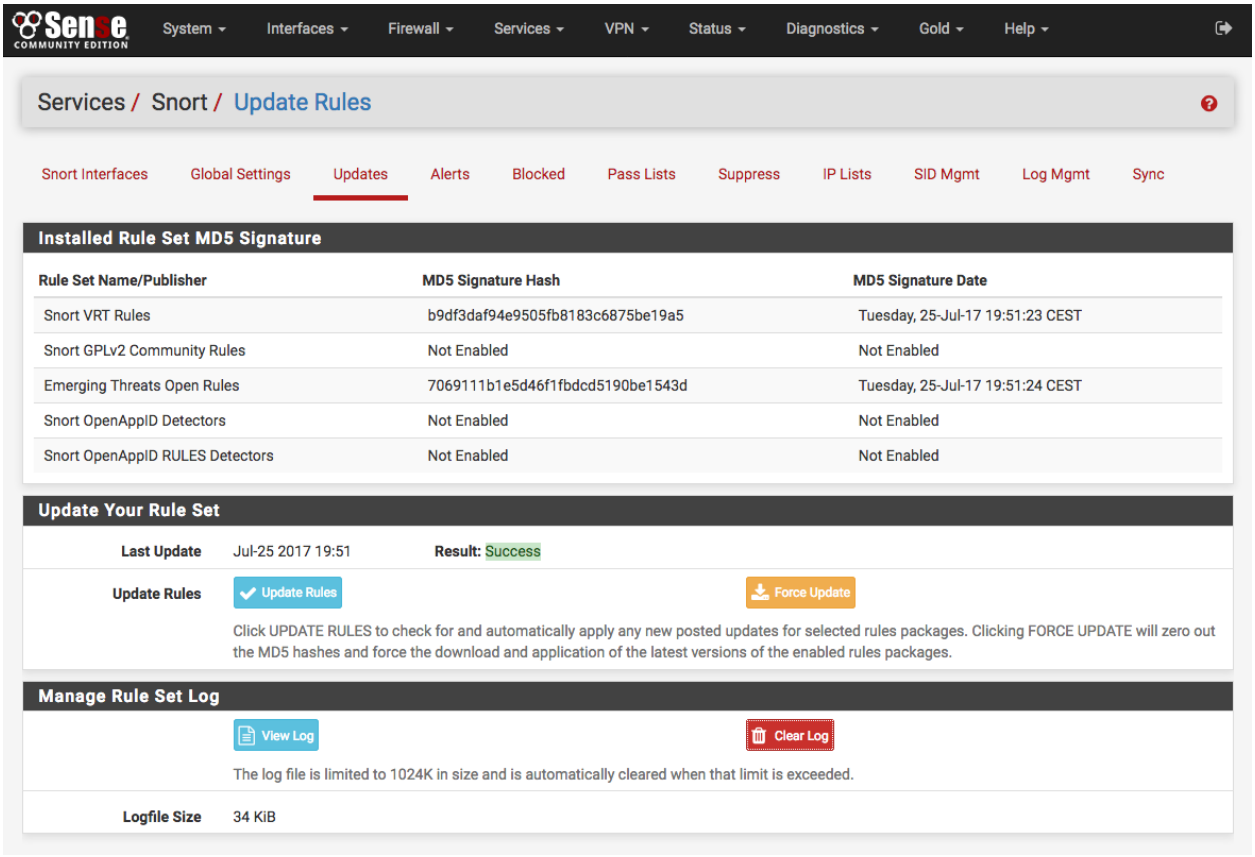


Update the rules

The **Updates** tab is used to check the status of downloaded rules packages and to download new updates. The table shows the available rule packages and their current status (not enabled, not downloaded, or a valid MD5 checksum and date).

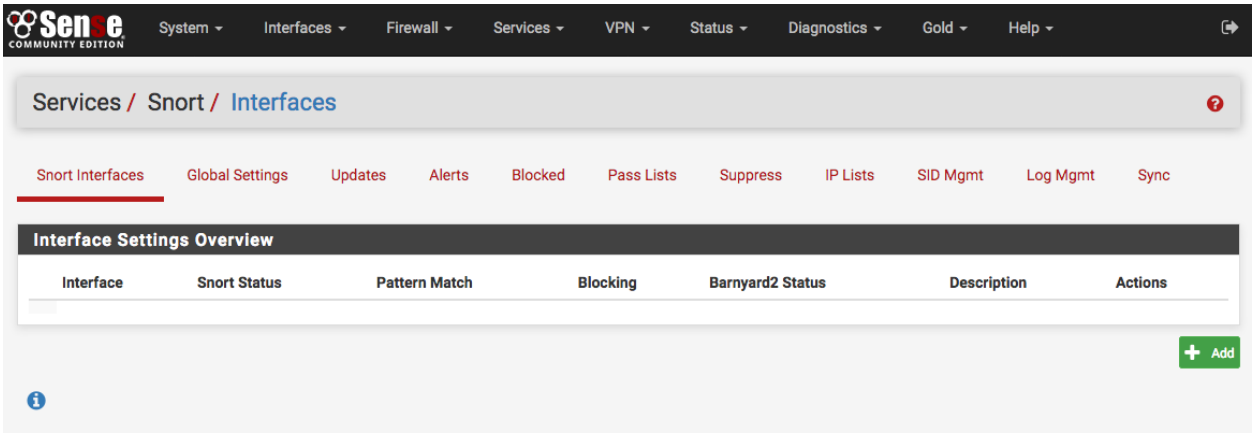
Click on the **Update Rules** button to download the latest rule package updates. If there is a newer set of packaged rules on the vendor website, it will be downloaded and installed. The determination is made by comparing the MD5 of the local file with that of the remote file on the vendor website. If there is a mismatch, a new file is downloaded. The **FORCE** button can be used to force download of the rule packages from the vendor website no matter how the MD5 hash tests out.

In the screenshot below, the Snort VRT and Emerging Threats Open rule packages have been successfully downloaded. The calculated MD5 hash and the file download date and time are shown. Also note the last update time and result are shown in the center of the page.

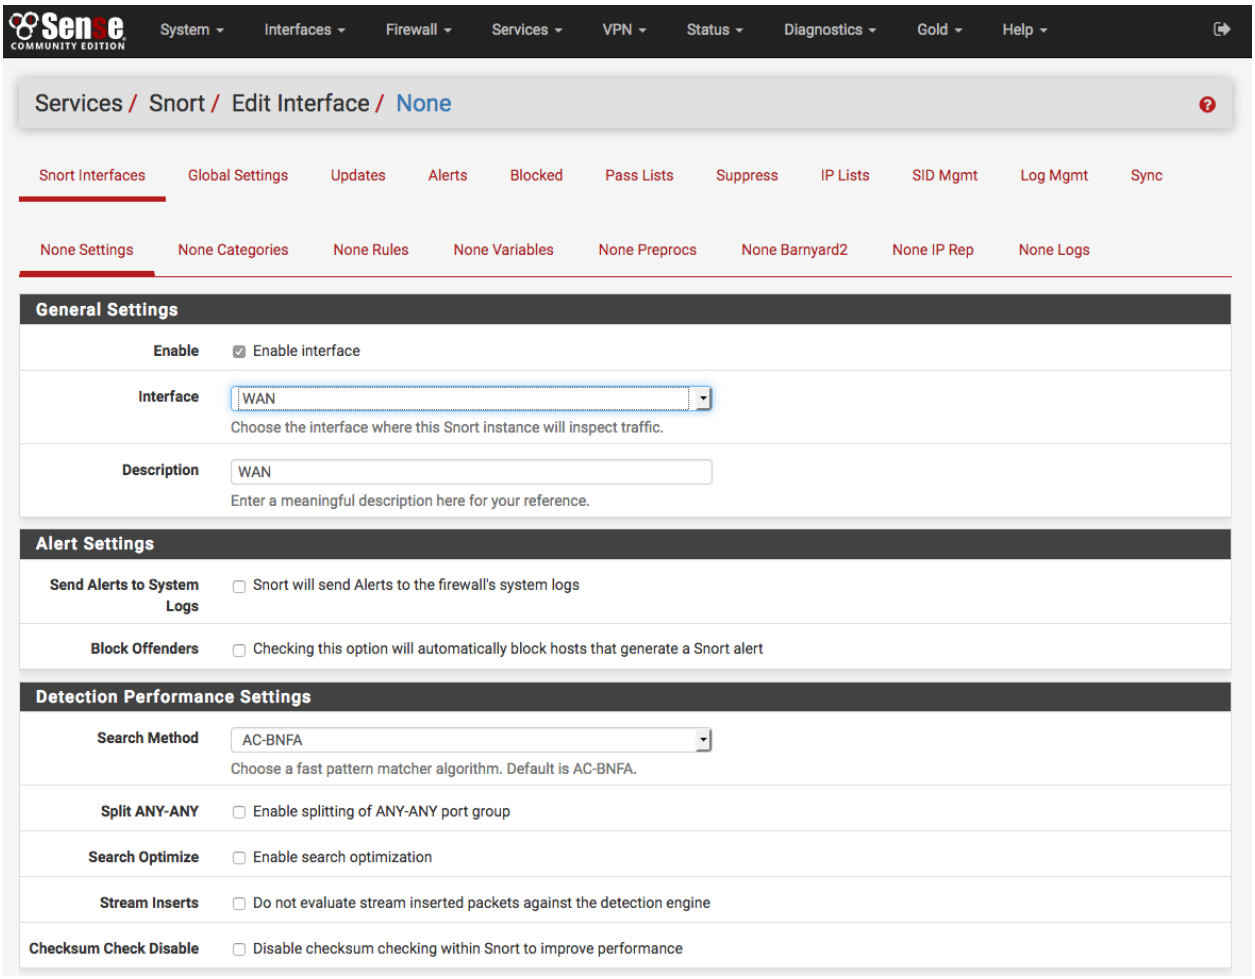


Add Snort to an interface

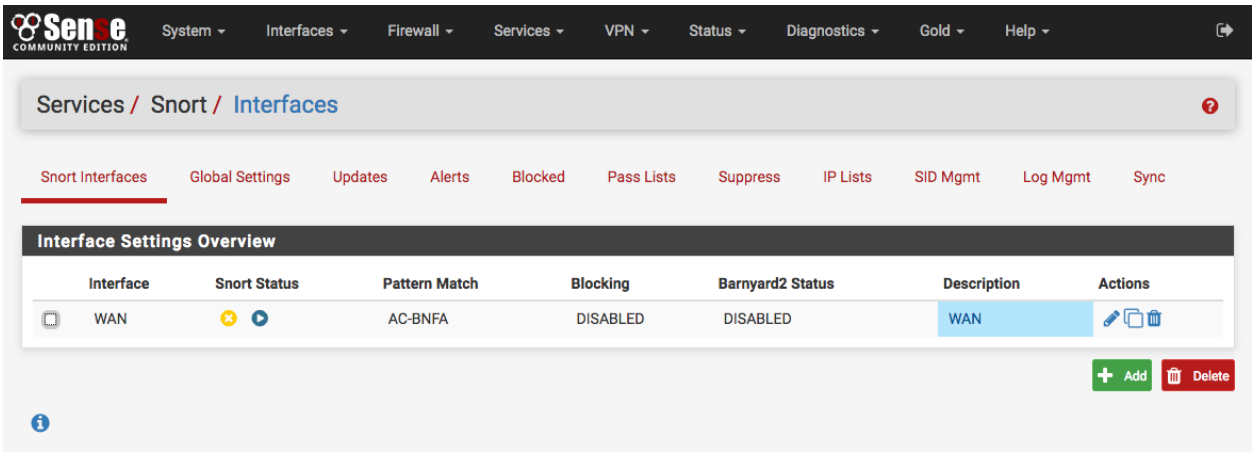
Click the **Snort Interfaces** tab and then the + icon to add a new Snort interface.



A new Interface Settings tab will open with the next available interface automatically selected. The interface selection may be changed using the **Interface** drop-down if desired. A descriptive name may also be provided for the interface. Other interface parameters may also be set on this page. Be sure to click the **SAVE** button down at the bottom of the page when finished.



After saving, the browser will be returned to the **Snort Interfaces** tab. Note the warning icons in the image below showing no rules have been selected for the new Snort interface. Those rules will be configured next. Click the  icon (shown highlighted with a red box in the image below) to edit the new Snort interface again.



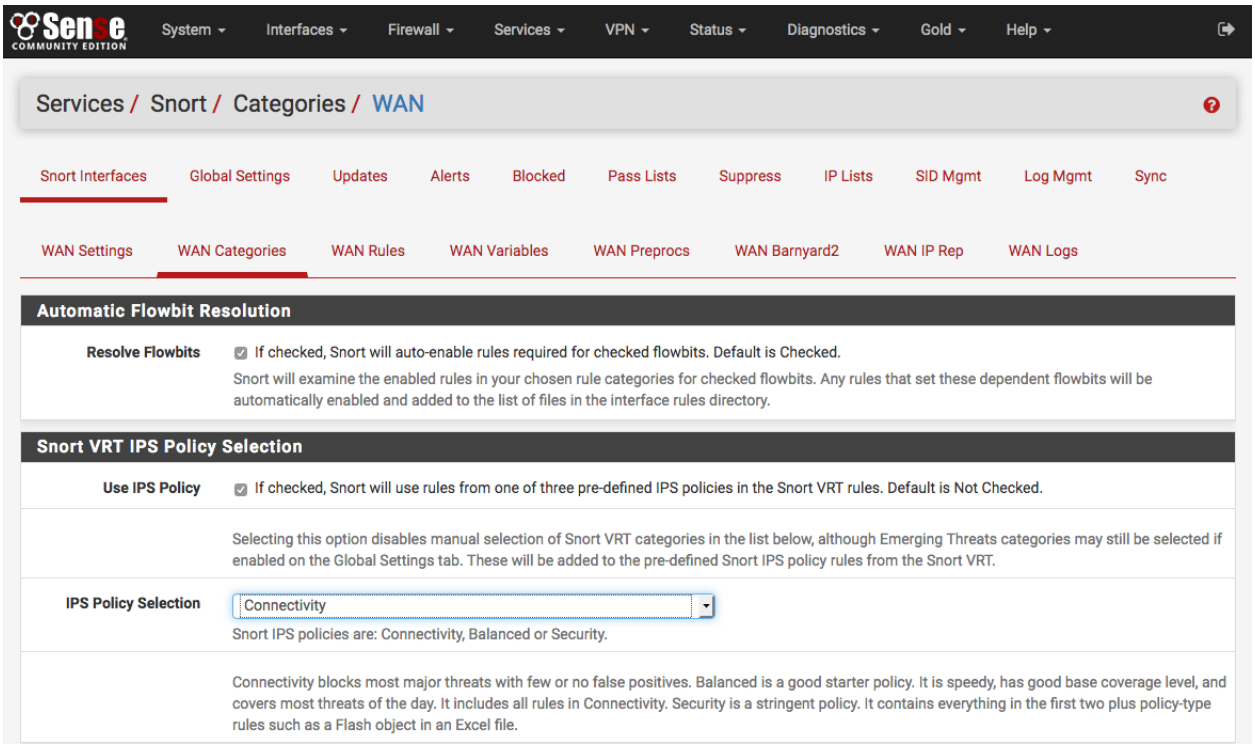
Select which types of rules will protect the network

Click the **Categories** tab for the new interface.

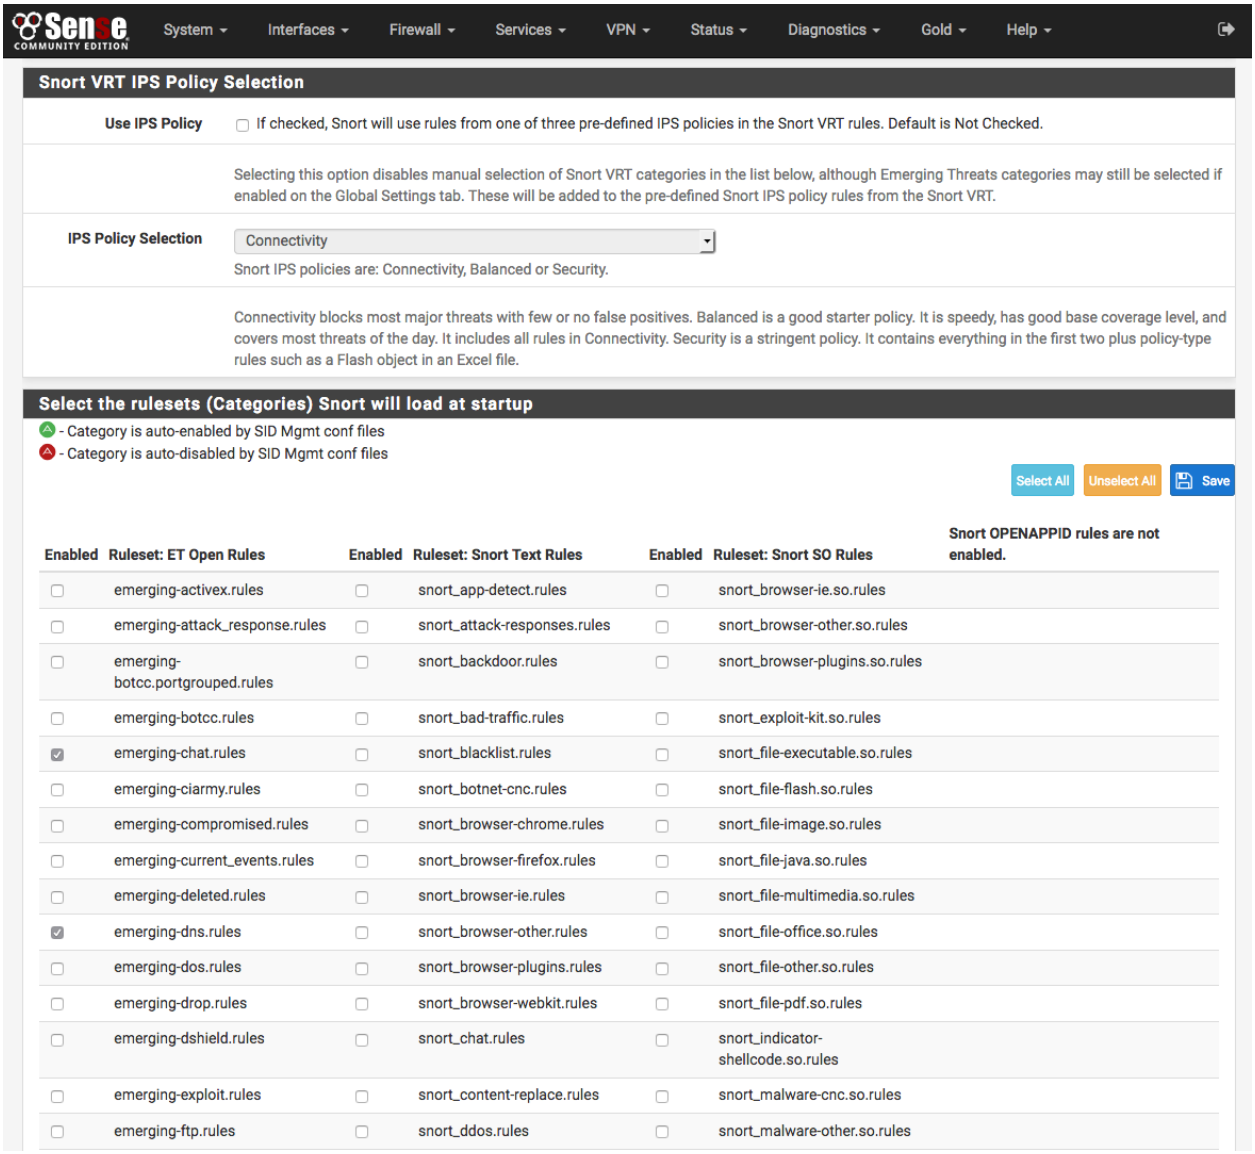
If a Snort VRT Oinkmaster code was obtained (either free registered user or the paid subscription), enabled the Snort VRT rules, and entered the Oinkmaster code on the Global Settings tab then the option of choosing from among three pre-configured IPS policies is available. These greatly simplify the process of choosing enforcing rules for Snort to use when inspecting traffic. The IPS policies are only available when the Snort VRT rules are enabled.

The three Snort VRT IPS Policies are: (1) Connectivity, (2) Balanced and (3) Security. These are listed in order of increasing security. However, resist the temptation to immediately jump to the most secure *Security* policy if Snort is unfamiliar. False positives can frequently occur with the more secure policies, and careful tuning by an experienced administrator may be required.

Tip: If Snort is unfamiliar, then using the less restrictive *Connectivity* policy in non-blocking mode (the default setting) is recommended as a starting point to identify and whitelist false positives. Once experience with Snort has been gained in this network environment, blocking mode may be enabled (via the **Block Offenders** option in the **Snort Interface Settings** tab) and a more restrictive IPS policy may be chosen.



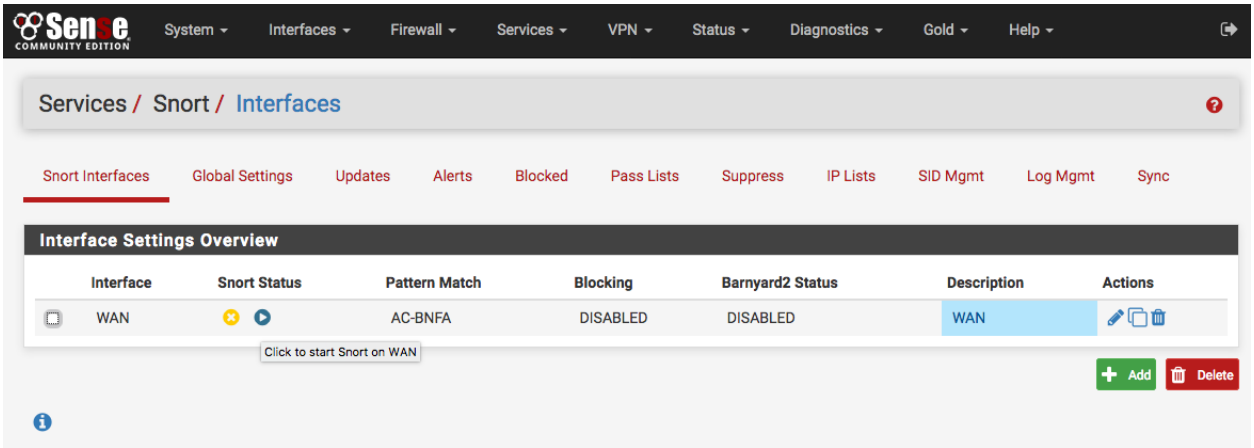
If the Snort VRT rules were not enabled, or if any of the other rule packages are to be used, then make the rule category selections by checking the checkboxes beside the rule categories to use.



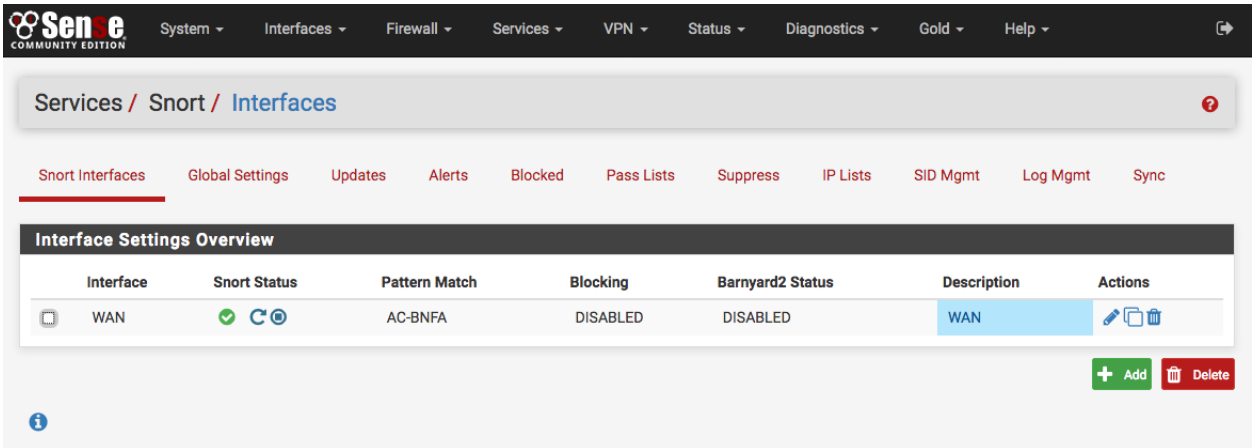
Click  **Save** to store the current selection and build the rules file for Snort to use.

Starting Snort on an interface

Click the **Snort Interfaces** tab to display the configured Snort interfaces. Click the  icon (shown highlighted with a red box in the image below) to start Snort on an interface.



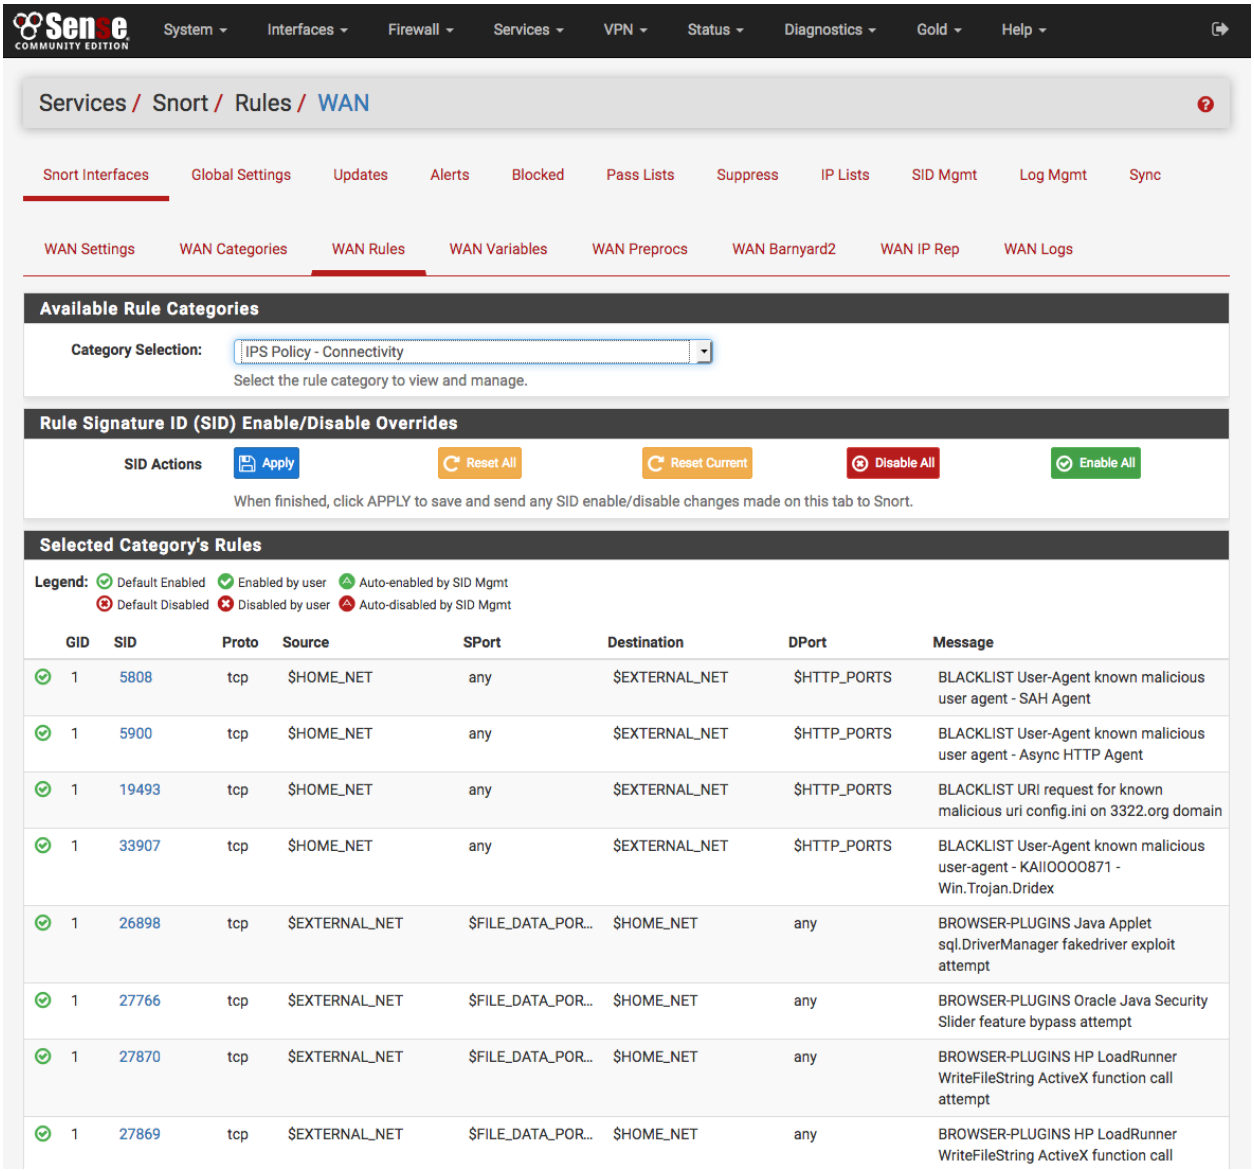
It will take several seconds for Snort to start. Once it has started, the icon will change to as shown below. To stop a running Snort instance on an interface, click the icon.

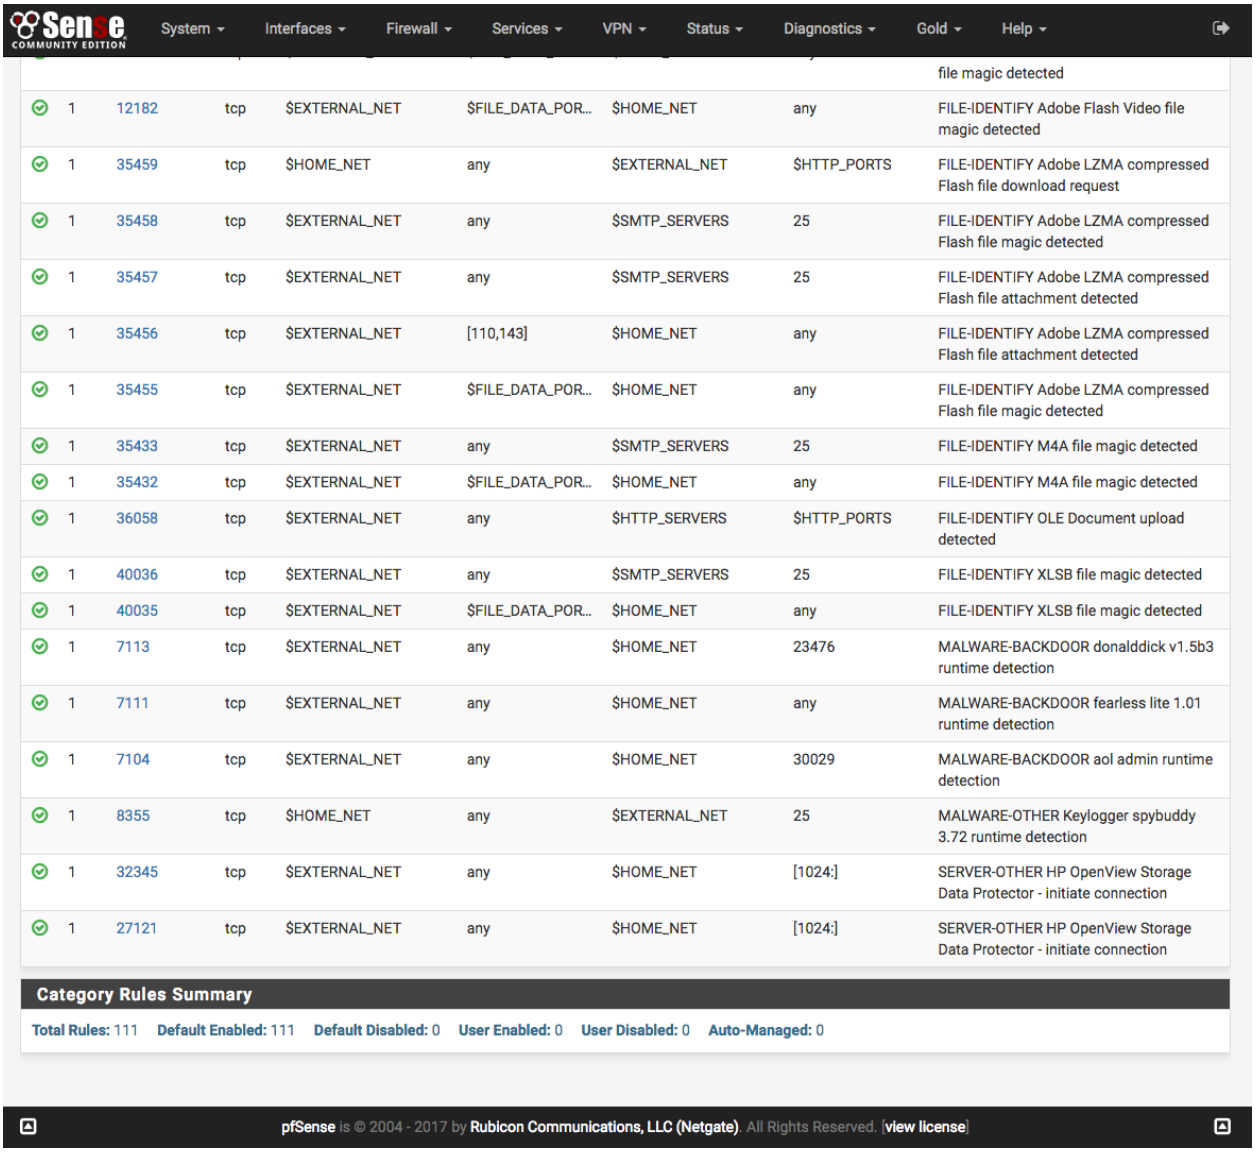


Select which types of signatures will protect the network

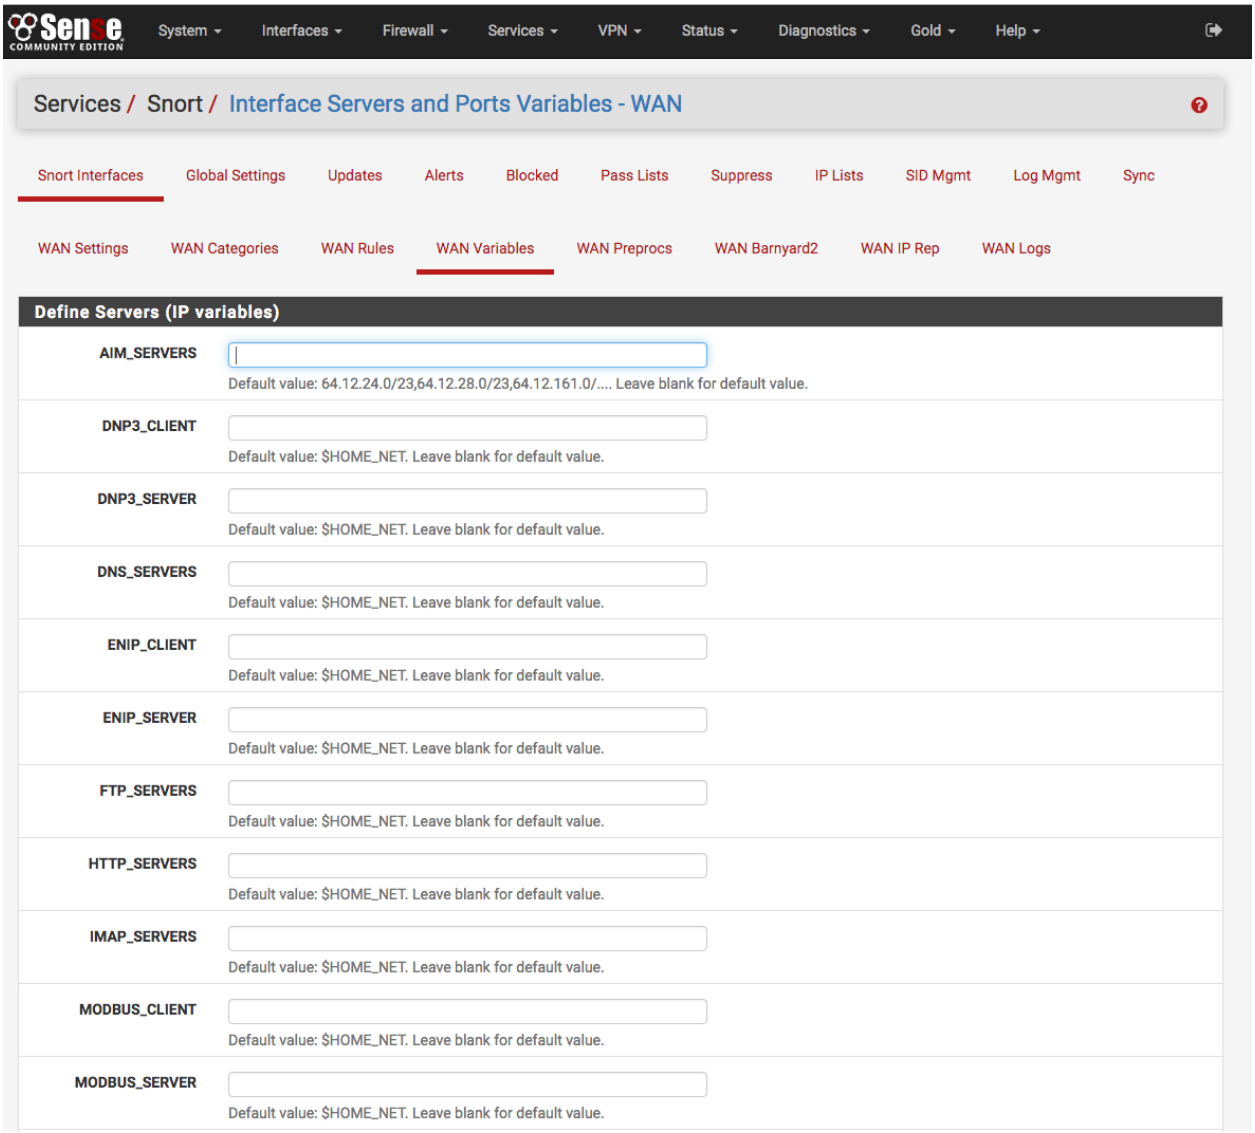
Click the **Rules** tab for the interface to configure individual rules in the enabled categories. Generally this page is only used to disable particular rules that may be generating too many false positives in a particular network environment. Be sure they are in fact truly false positives before taking the step of disabling a Snort rule!

Select a rule category from the **Category** drop-down to view all the assigned rules. Click the or icon at the far-left of a row to toggle the rule's state from enabled to disabled, or click or to toggle from disabled to enabled. The icon will change to indicate the state of the rule. At the top of the rule list is a legend showing the icons used to indicate the current state of a rule.


System
Interfaces
Firewall
Services
VPN
Status
Diagnostics
Gold
Help
Services / Snort / Rules / WAN
Snort Interfaces
Global Settings
Updates
Alerts
Blocked
Pass Lists
Suppress
IP Lists
SID Mgmt
Log Mgmt
Sync
WAN Settings
WAN Categories
WAN Rules
WAN Variables
WAN Preprocs
WAN Barnyard2
WAN IP Rep
WAN Logs
Available Rule Categories
Category Selection:
IPS Policy - Connectivity
Select the rule category to view and manage.
Rule Signature ID (SID) Enable/Disable Overrides
SID Actions
Apply
Reset All
Reset Current
Disable All
Enable All
When finished, click APPLY to save and send any SID enable/disable changes made on this tab to Snort.
Selected Category's Rules
Legend:
Default Enabled
Enabled by user
Auto-enabled by SID Mgmt
Default Disabled
Disabled by user
Auto-disabled by SID Mgmt
GID SID Proto Source SPort Destination DPort Message
1 5808 tcp $HOME_NET any $EXTERNAL_NET $HTTP_PORTS BLACKLIST User-Agent known malicious user agent - SAH Agent
1 5900 tcp $HOME_NET any $EXTERNAL_NET $HTTP_PORTS BLACKLIST User-Agent known malicious user agent - Async HTTP Agent
1 19493 tcp $HOME_NET any $EXTERNAL_NET $HTTP_PORTS BLACKLIST URI request for known malicious uri config.ini on 3322.org domain
1 33907 tcp $HOME_NET any $EXTERNAL_NET $HTTP_PORTS BLACKLIST User-Agent known malicious user-agent - KAIIOOOO871 - Win.Trojan.Dridex
1 26898 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any BROWSER-PLUGINS Java Applet sql.DriverManager fakedriver exploit attempt
1 27766 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any BROWSER-PLUGINS Oracle Java Security Slider feature bypass attempt
1 27870 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any BROWSER-PLUGINS HP LoadRunner WriteFileString ActiveX function call attempt
1 27869 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any BROWSER-PLUGINS HP LoadRunner WriteFileString ActiveX function call

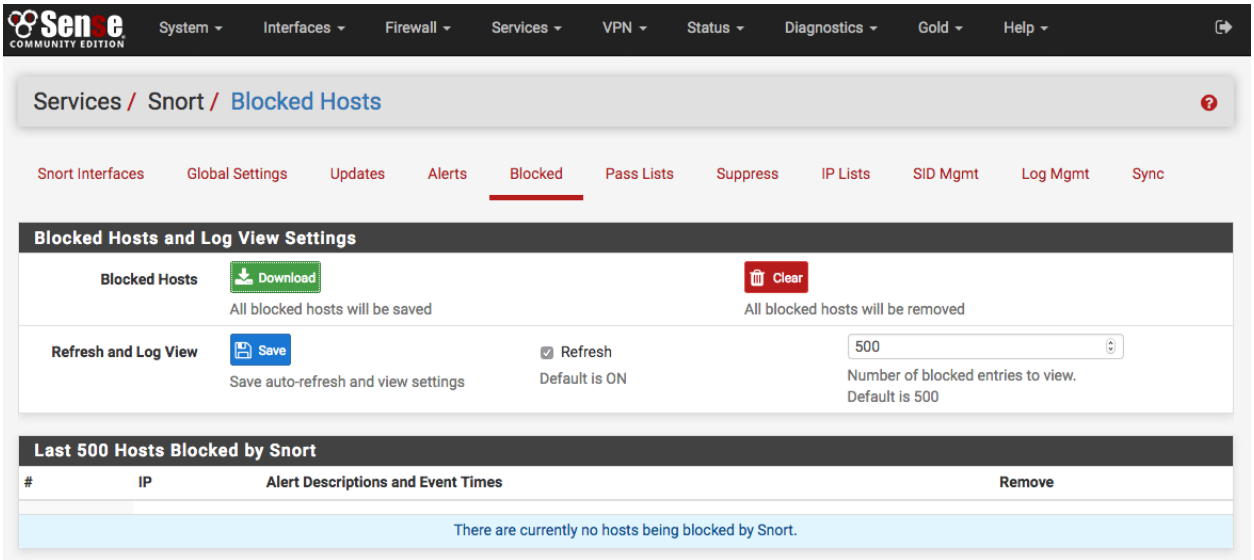

pfSense COMMUNITY EDITION
System
Interfaces
Firewall
Services
VPN
Status
Diagnostics
Gold
Help
file magic detected
1 12182 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any FILE-IDENTIFY Adobe Flash Video file magic detected
1 35459 tcp $HOME_NET any $EXTERNAL_NET $HTTP_PORTS FILE-IDENTIFY Adobe LZMA compressed Flash file download request
1 35458 tcp $EXTERNAL_NET any $SMTP_SERVERS 25 FILE-IDENTIFY Adobe LZMA compressed Flash file magic detected
1 35457 tcp $EXTERNAL_NET any $SMTP_SERVERS 25 FILE-IDENTIFY Adobe LZMA compressed Flash file attachment detected
1 35456 tcp $EXTERNAL_NET [110,143] $HOME_NET any FILE-IDENTIFY Adobe LZMA compressed Flash file attachment detected
1 35455 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any FILE-IDENTIFY Adobe LZMA compressed Flash file magic detected
1 35433 tcp $EXTERNAL_NET any $SMTP_SERVERS 25 FILE-IDENTIFY M4A file magic detected
1 35432 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any FILE-IDENTIFY M4A file magic detected
1 36058 tcp $EXTERNAL_NET any $HTTP_SERVERS $HTTP_PORTS FILE-IDENTIFY OLE Document upload detected
1 40036 tcp $EXTERNAL_NET any $SMTP_SERVERS 25 FILE-IDENTIFY XLSB file magic detected
1 40035 tcp $EXTERNAL_NET $FILE_DATA_POR... $HOME_NET any FILE-IDENTIFY XLSB file magic detected
1 7113 tcp $EXTERNAL_NET any $HOME_NET 23476 MALWARE-BACKDOOR donalddick v1.5b3 runtime detection
1 7111 tcp $EXTERNAL_NET any $HOME_NET any MALWARE-BACKDOOR fearless lite 1.01 runtime detection
1 7104 tcp $EXTERNAL_NET any $HOME_NET 30029 MALWARE-BACKDOOR aol admin runtime detection
1 8355 tcp $HOME_NET any $EXTERNAL_NET 25 MALWARE-OTHER Keylogger spybuddy 3.72 runtime detection
1 32345 tcp $EXTERNAL_NET any $HOME_NET [1024:] SERVER-OTHER HP OpenView Storage Data Protector - initiate connection
1 27121 tcp $EXTERNAL_NET any $HOME_NET [1024:] SERVER-OTHER HP OpenView Storage Data Protector - initiate connection
Category Rules Summary
Total Rules: 111 Default Enabled: 111 Default Disabled: 0 User Enabled: 0 User Disabled: 0 Auto-Managed: 0
pfSense is © 2004 - 2017 by Rubicon Communications, LLC (Netgate). All Rights Reserved. [view license]

Define servers to protect and improve performance



Managing blocked hosts

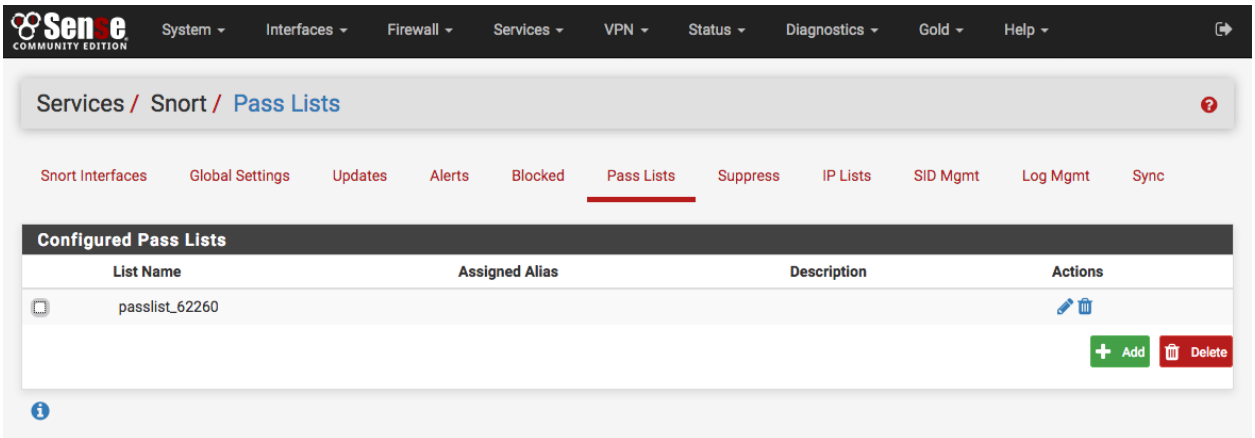
The **Blocked** tab shows what hosts are currently being blocked by Snort (when the block offenders option is selected on the **Interface Settings** tab). Blocked hosts can be automatically cleared by Snort at one of several pre-defined intervals. The blocking options for an interface are configured on the Snort **Interface Settings** tab for the interface.



Managing Pass lists

Pass Lists are lists of IP addresses that Snort should never block. These may be created and managed on the **Pass Lists** tab. When an IP address is listed on a Pass List, Snort will never insert a block on that address even when malicious traffic is detected.

To create a new Pass List, click . To edit an existing Pass List, click the . To delete a Pass List, click . Note that a Pass List may not be deleted if it is currently assigned to one or more Snort interfaces.

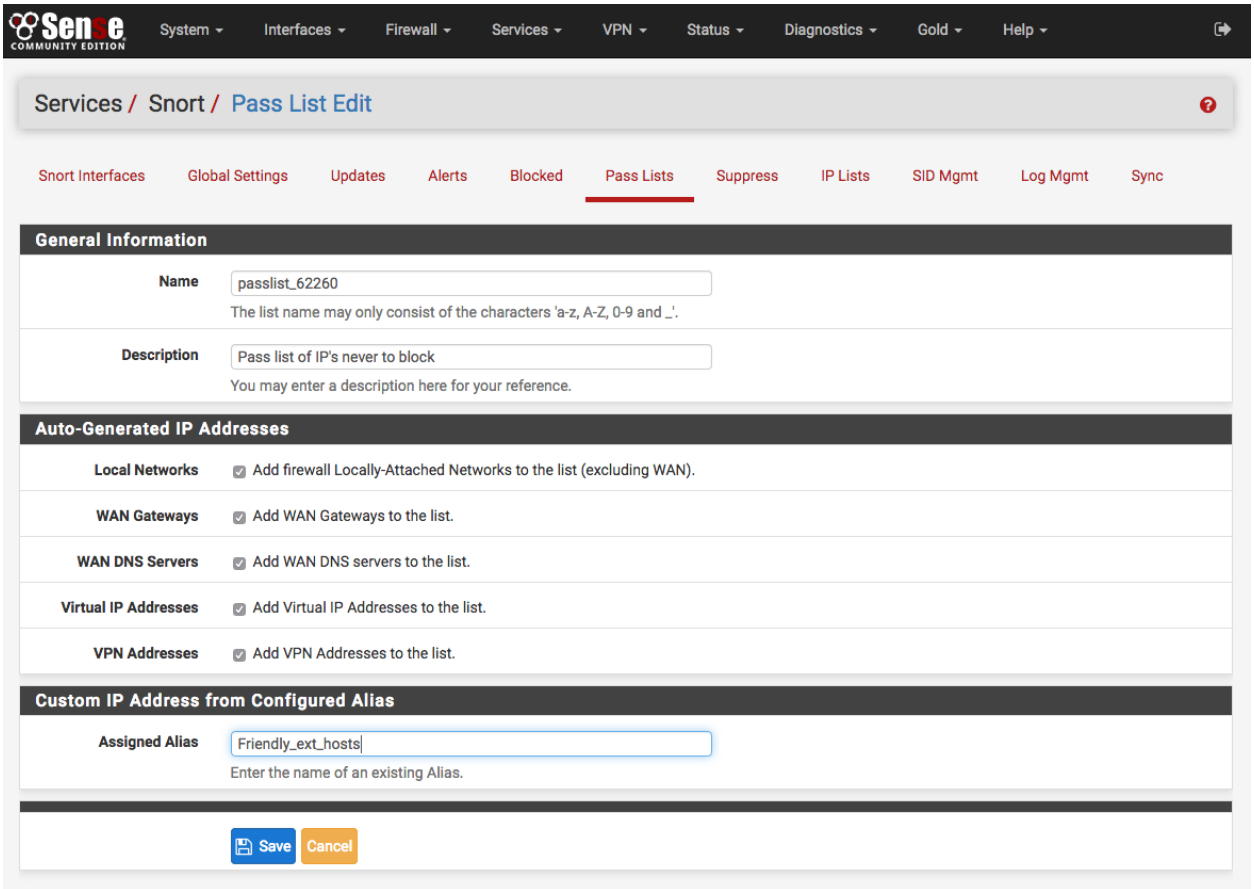


A default Pass List is automatically generated by Snort for every interface, and this default list is used when no other list is specified. Pass Lists are assigned to an interface on the **Interface Settings** tab.

Customized Pass List may be created and assigned to an interface. This might be done when trusted external hosts exist that are not located on networks directly connected to the firewall. To add external hosts in this manner, first create an Alias under **Firewall > Aliases** and then assign that alias to the **Assigned Aliases** field. In the example shown below, the alias "*Friendly_ext_hosts*" has been assigned. This alias would contain the IP addresses of the trusted external hosts.

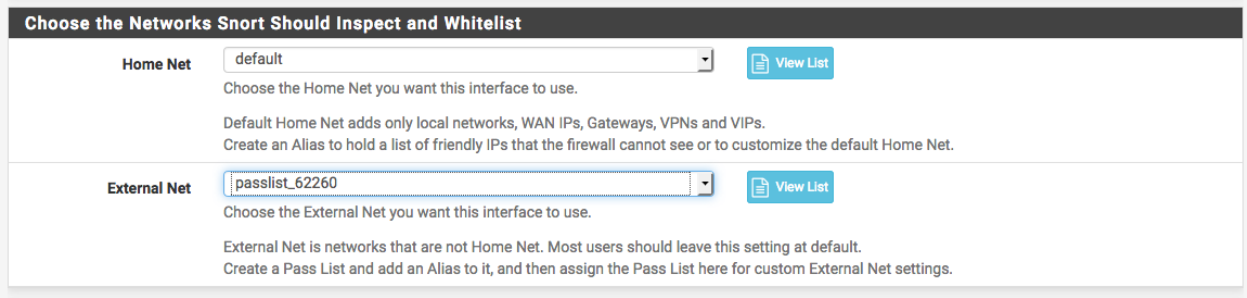
When creating a custom Pass List, leave all the auto-generated IP addresses checked in the **Add auto-generated IP addresses** section. Not selecting the checkboxes in this section can lead to blocking of critical addresses including the

firewall interfaces themselves. This could result in being locked out of the firewall over the network! Only uncheck boxes in this section when absolutely necessary.



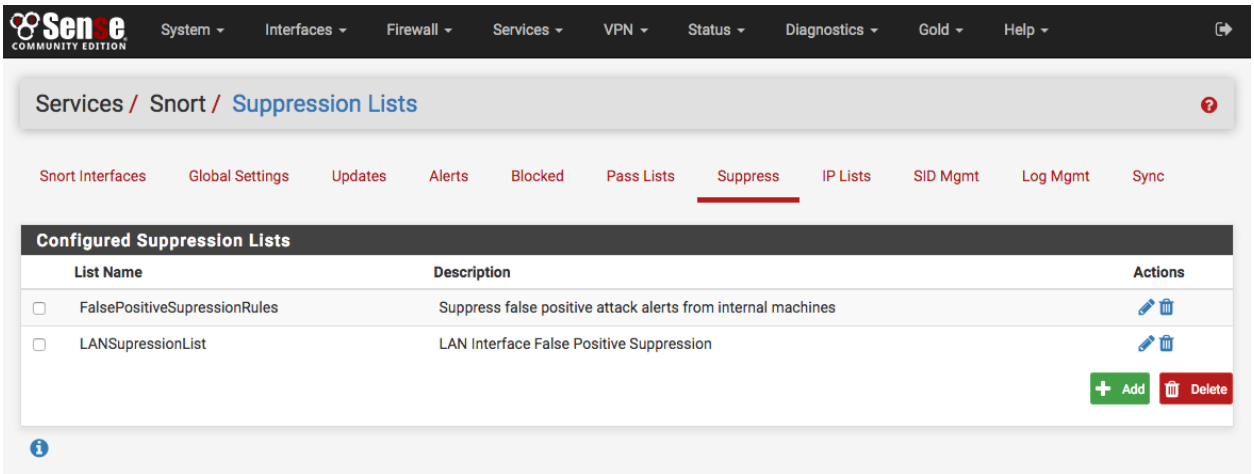
Click the **ALIASES** button to open a window showing previously defined aliases for selection. Remember to click **SAVE** to save changes.

Note: Remember that simply creating a Pass List is only the first step! It must be selected by going to the **Interface Settings** tab for the Snort interface and assigning the newly created Pass List as shown below. After assigning and saving the new Pass List, restart Snort on the affected interface to pick up the change.

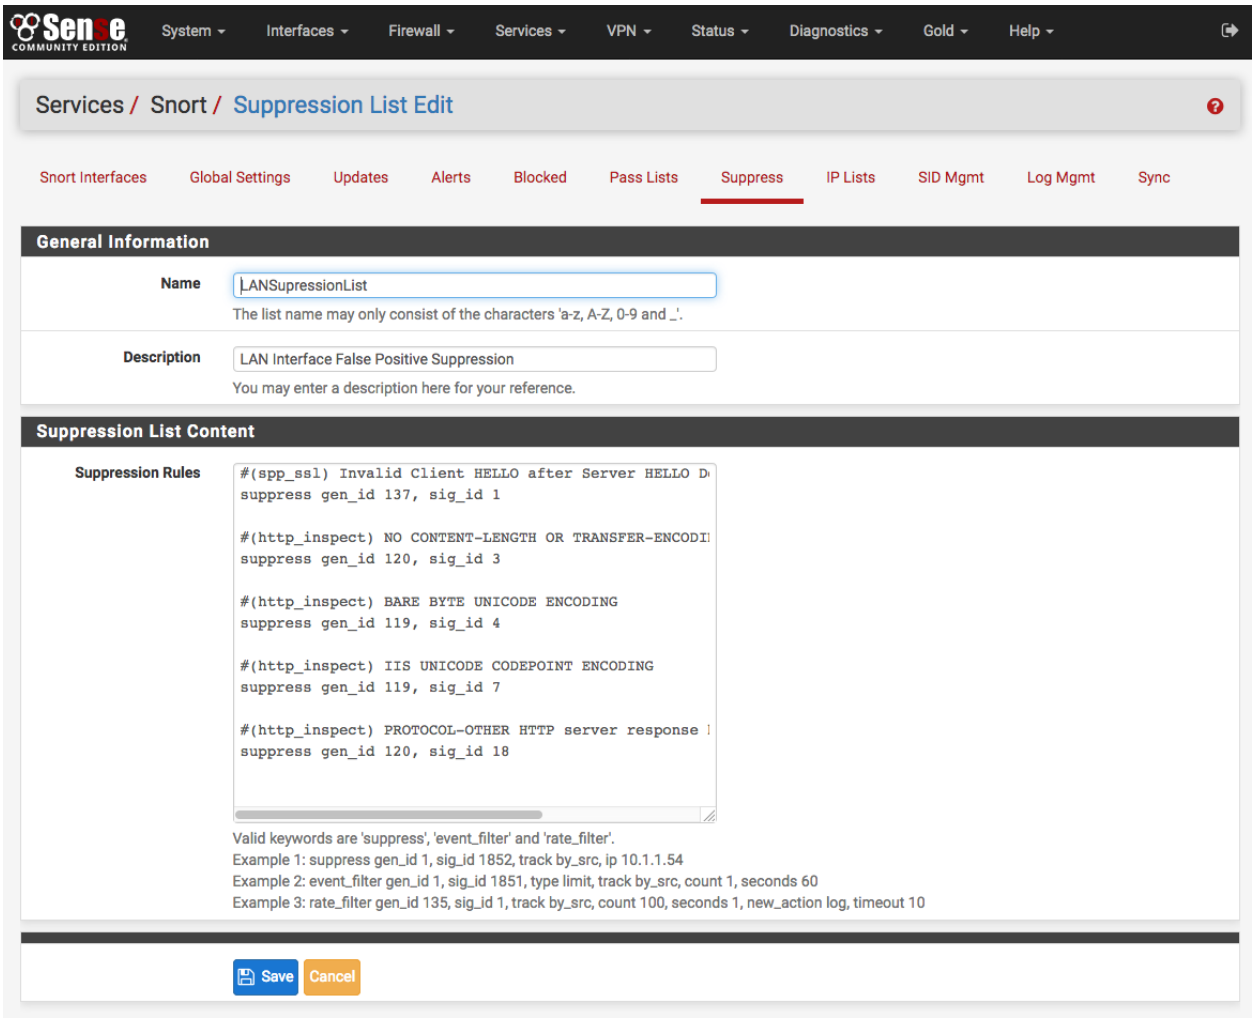


Alert Thresholding and Suppression

Suppression Lists allow control over the alerts generated by Snort rules. When an alert is suppressed, then Snort no longer logs an alert entry (or blocks the IP address if block offenders is enabled) when a particular rule fires. Snort still inspects all network traffic against the rule, but even when traffic matches the rule signature, no alert will be generated. This is different from disabling a rule. When a rule is disabled, Snort no longer tries to match it to any network traffic. Suppressing a rule might be done in lieu of disabling the rule when alerts should only be stopped based on either the source or destination IP. For example, to suppress the alert when traffic from a particular trusted IP address is the source. If any other IP is the source or destination of the traffic, the rule would still fire. To eliminate all alerts from the rule, then it is more efficient to simply disable the rule rather than to suppress it. Disabling the rule will remove it from Snort's list of match rules and therefore makes for less work Snort has to do.



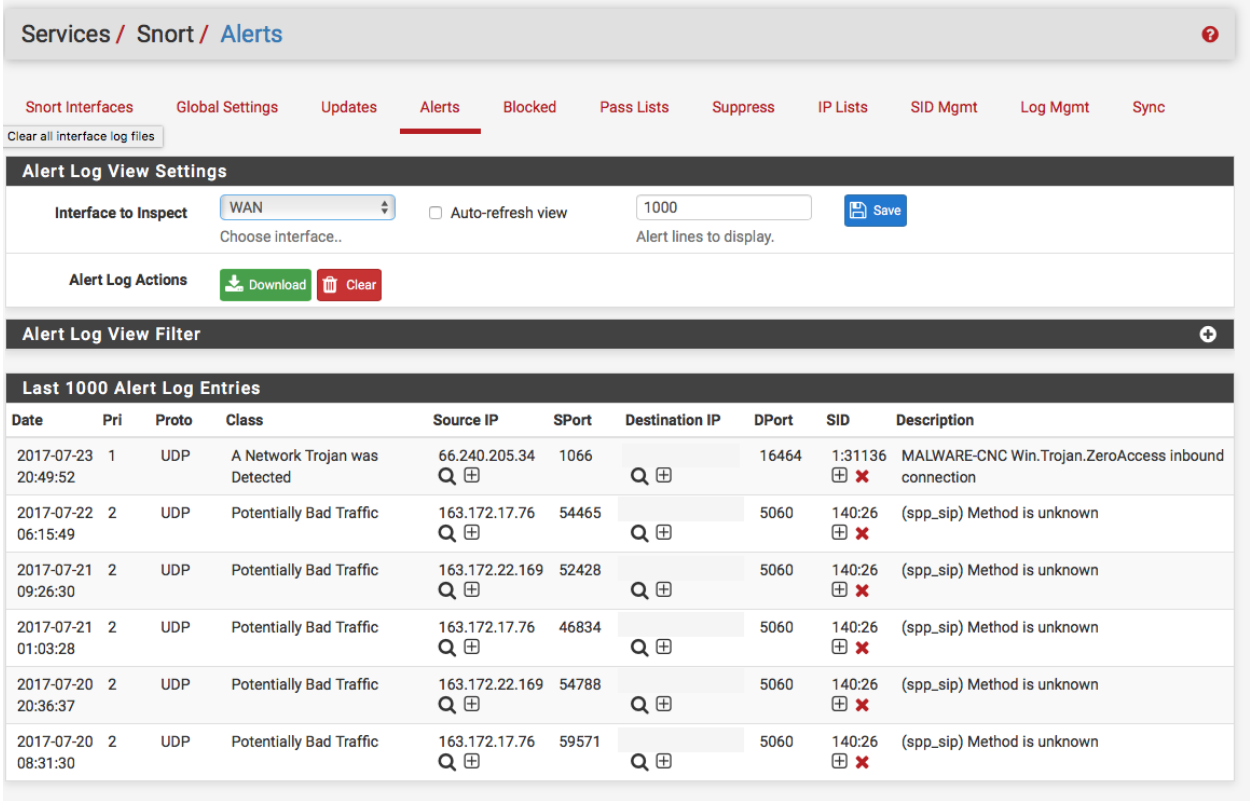
On the Suppress List Edit page, a new suppress list entry may be manually added or edited. It is usually easier and faster to add suppress list entries by clicking  shown with the alert entries on the **Alerts** tab. Remember to click the **SAVE** button to save changes when manually editing Suppress List entries.



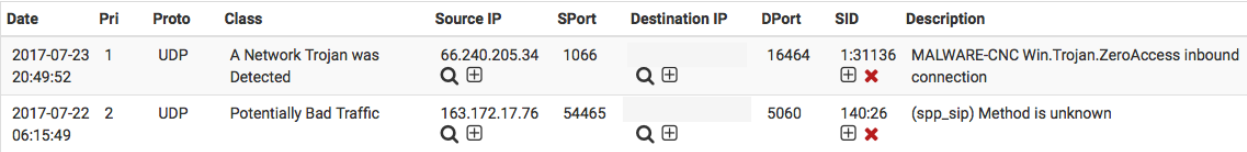
Getting to know the alerts

The **Alerts** tab is where alerts generated by Snort are viewed. If Snort is running on more than one interface, choose the interface whose alerts should be viewed in the drop-down selector.

Use the **DOWNLOAD** button to download an archive containing all the logged alerts to a local machine. The **CLEAR** button is used to erase the current alerts log. Destination IP addresses have been redacted from the screenshot.



Alert Details



The **Date** column shows the date and time the alert was generated. The remaining columns show data from the rule that generated the alert.

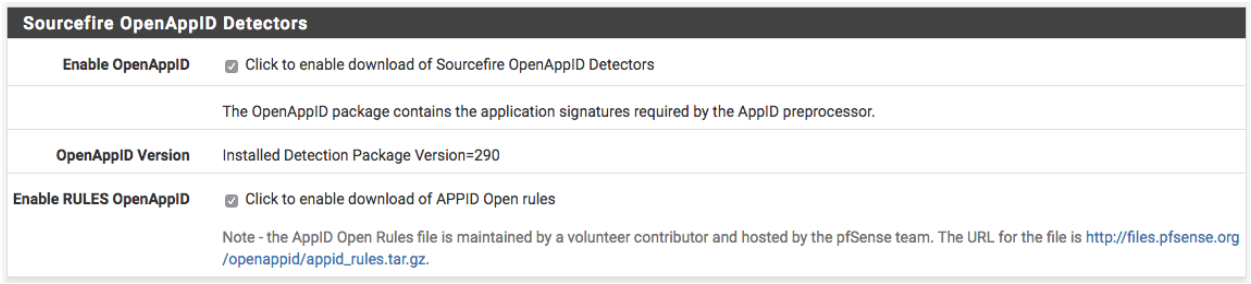
In the **Source**, **Destination** columns are  icons for performing reverse DNS lookups on the IP addresses as well as a  icon used to add an automatic *Suppress List* entry for the alert using the IP address and SID (signature ID). This will prevent future alerts from being generated by the rule for that specific IP address only. If either of the Source or Destination addresses are currently being blocked by Snort, then a  icon will also be shown. Clicking that icon will remove the block for the IP address.

The SID column contains two icons. The  icon will automatically add that SID to the *Suppress List* for the interface and suppress future alerts from the signature for all IP addresses. The  icon in the SID column will disable the rule and remove it from the enforcing rule set. When a rule is manually disabled, the icon in the SID column changes to  .

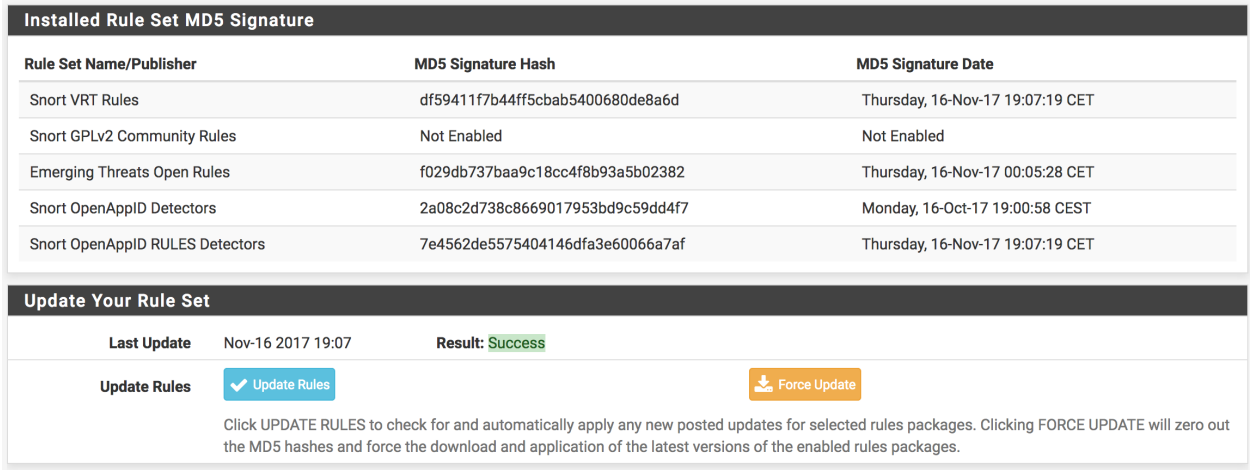
Application ID detection with OpenAppID

OpenAppID is an application-layer network security plugin for the open source intrusion detection system Snort. Learn more about it here.

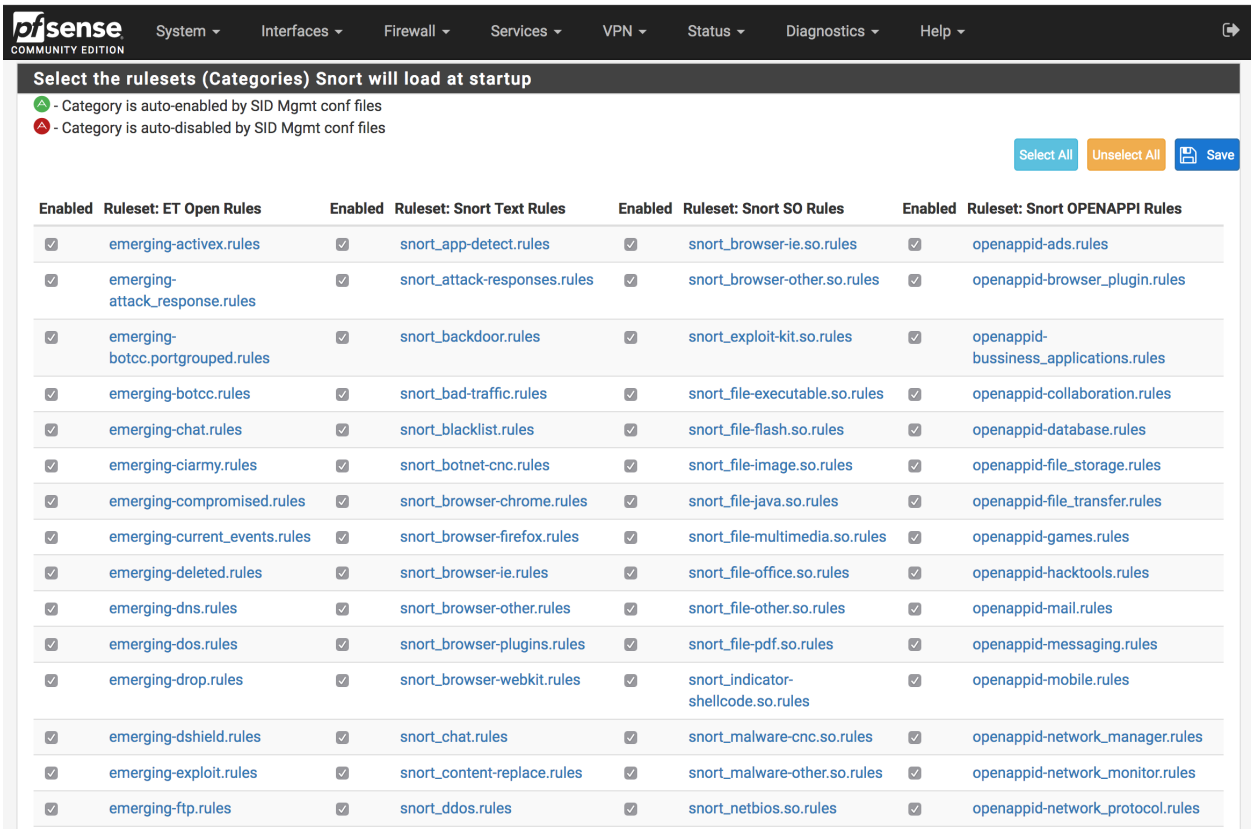
Enabling OpenAppID and its rules is done from Snort **Global Settings**. Select both checkboxes to enable detectors and rules download. Save the page.



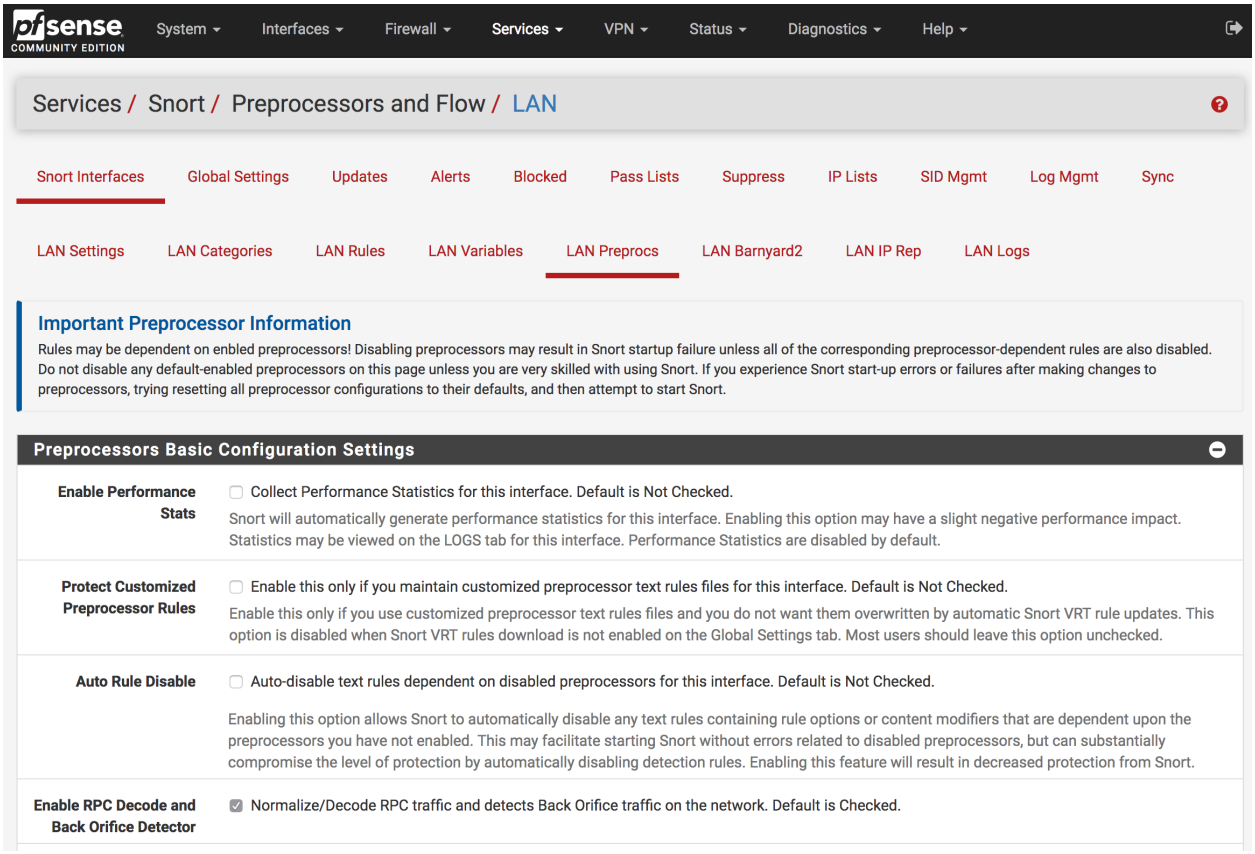
After enabling the detectors and rules go to Snort Updates tab and click on **Update Rules**. Wait for all the rules to update. Once done, the page will show OpenAppID detectors and rules have been updated.



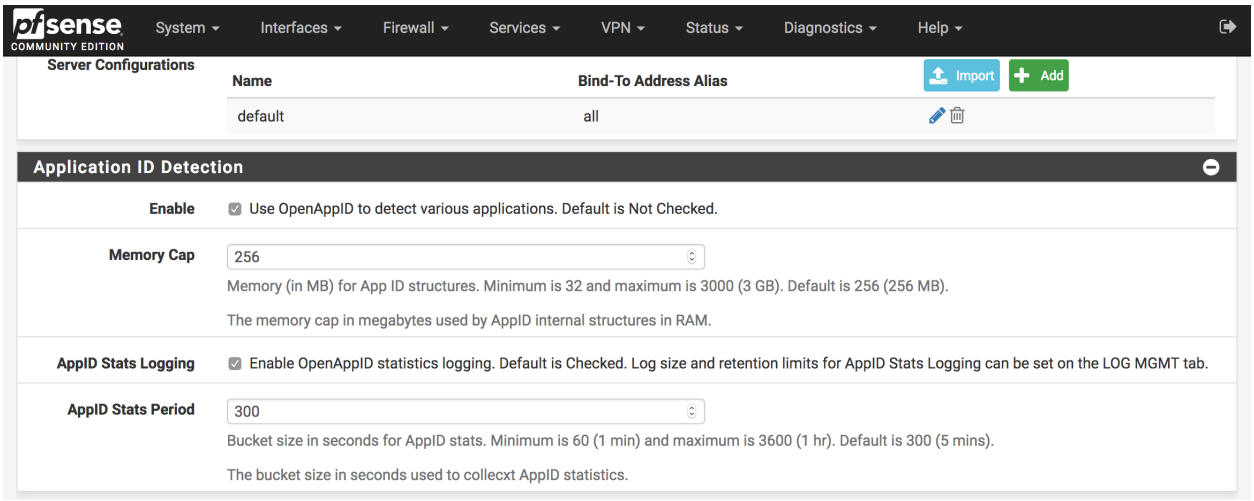
The following steps assume the firewall already has a Snort interface for LAN. Edit the LAN interface and navigate to LAN categories tab. When there, make sure the **Snort OPENAPPID Rules** from the right column are all selected and click **Save**.



Lastly, while still editing Snort interface, navigate to **LAN Preprocessor** tab.

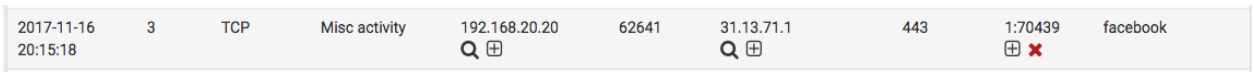


Scroll down to **Application ID Detection** section and select both **Enable** and **AppID Stats Logging** checkboxes. Save the page and OpenAppID will be activated on the Snort interface.

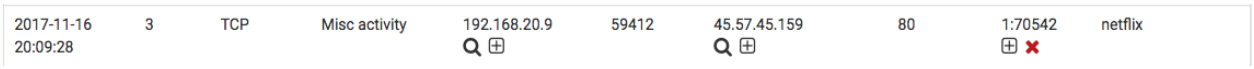


Viewing detected applications can be done from **Alerts** tab. The following screenshots are examples of identified services and applications:

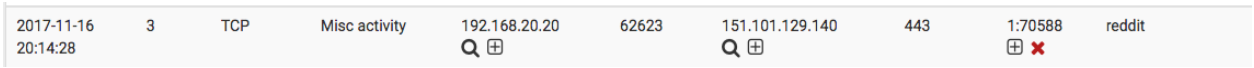
Facebook



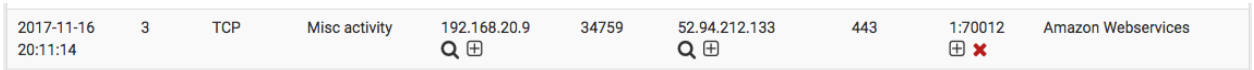
Netflix



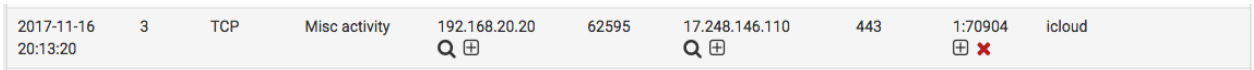
Reddit



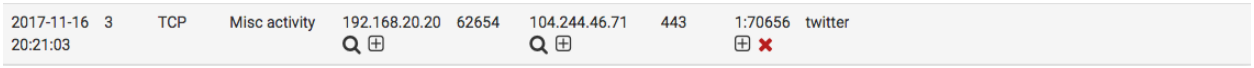
Amazon Web Services



iCloud



Twitter



Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

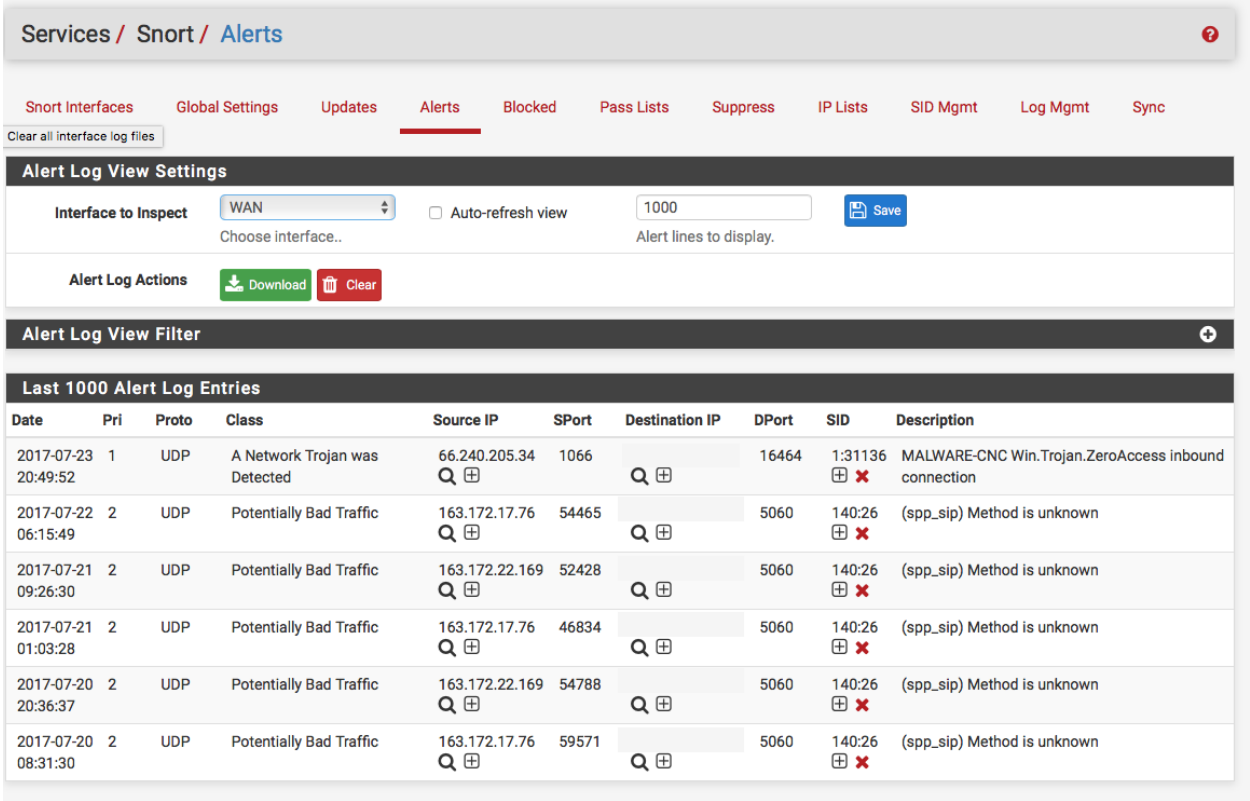
Package Assistance

Netgate TAC can only assist with the installation of this package. Netgate Professional Services can assist with custom configurations.

Snort Alerts

The **Alerts** tab is where alerts generated by Snort may be viewed. If Snort is running on more than one interface, choose the interface to view alerts for in the drop-down selector.

Use the **DOWNLOAD** button to download an archive containing all the logged alerts to a local machine. The **CLEAR** button is used to erase the current alerts log.



Alert Details

Date	Pri	Proto	Class	Source IP	SPort	Destination IP	DPort	SID	Description
2017-07-23 20:49:52	1	UDP	A Network Trojan was Detected	66.240.205.34	1066		16464	1:31136	MALWARE-CNC Win.Trojan.ZeroAccess inbound connection
2017-07-22 06:15:49	2	UDP	Potentially Bad Traffic	163.172.17.76	54465		5060	140:26	(spp_sip) Method is unknown

The **Date** column shows the date and time the alert was generated. The remaining columns show data from the rule that generated the alert.

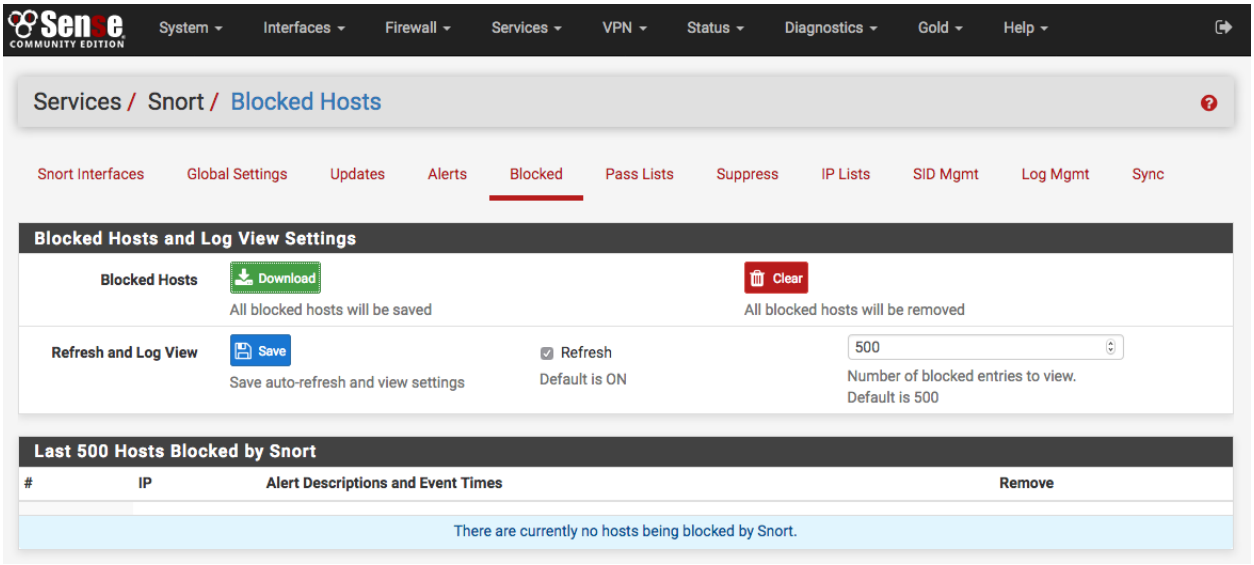
In the **Source**, **Destination** columns are  icons for performing reverse DNS lookups on the IP addresses as well as a  icon used to add an automatic *Suppress List* entry for the alert using the IP address and SID (signature ID). This will prevent future alerts from being generated by the rule for that specific IP address only. If either of the Source or Destination addresses are currently being blocked by Snort, then a  icon will also be shown. Clicking that icon will remove the block for the IP address.

The SID column contains two icons. The  icon will automatically add that SID to the *Suppress List* for the interface and suppress future alerts from the signature for all IP addresses. The  icon in the SID column will disable the rule and remove it from the enforcing rule set. When a rule is manually disabled, the icon in the SID column changes to  .

Snort Blocked Hosts

The **Blocked** tab shows what hosts are currently being blocked by Snort (when the **block offenders** option is selected on the **Interface Settings** tab). Blocked hosts can be automatically cleared by Snort at one of several pre-defined intervals. The blocking options for an interface are configured on the Snort **Interface Settings** tab for the interface. To manually remove a blocked host, click the  icon in the right-hand column.

The  icon performs a reverse DNS lookup on the blocked host IP address when clicked.



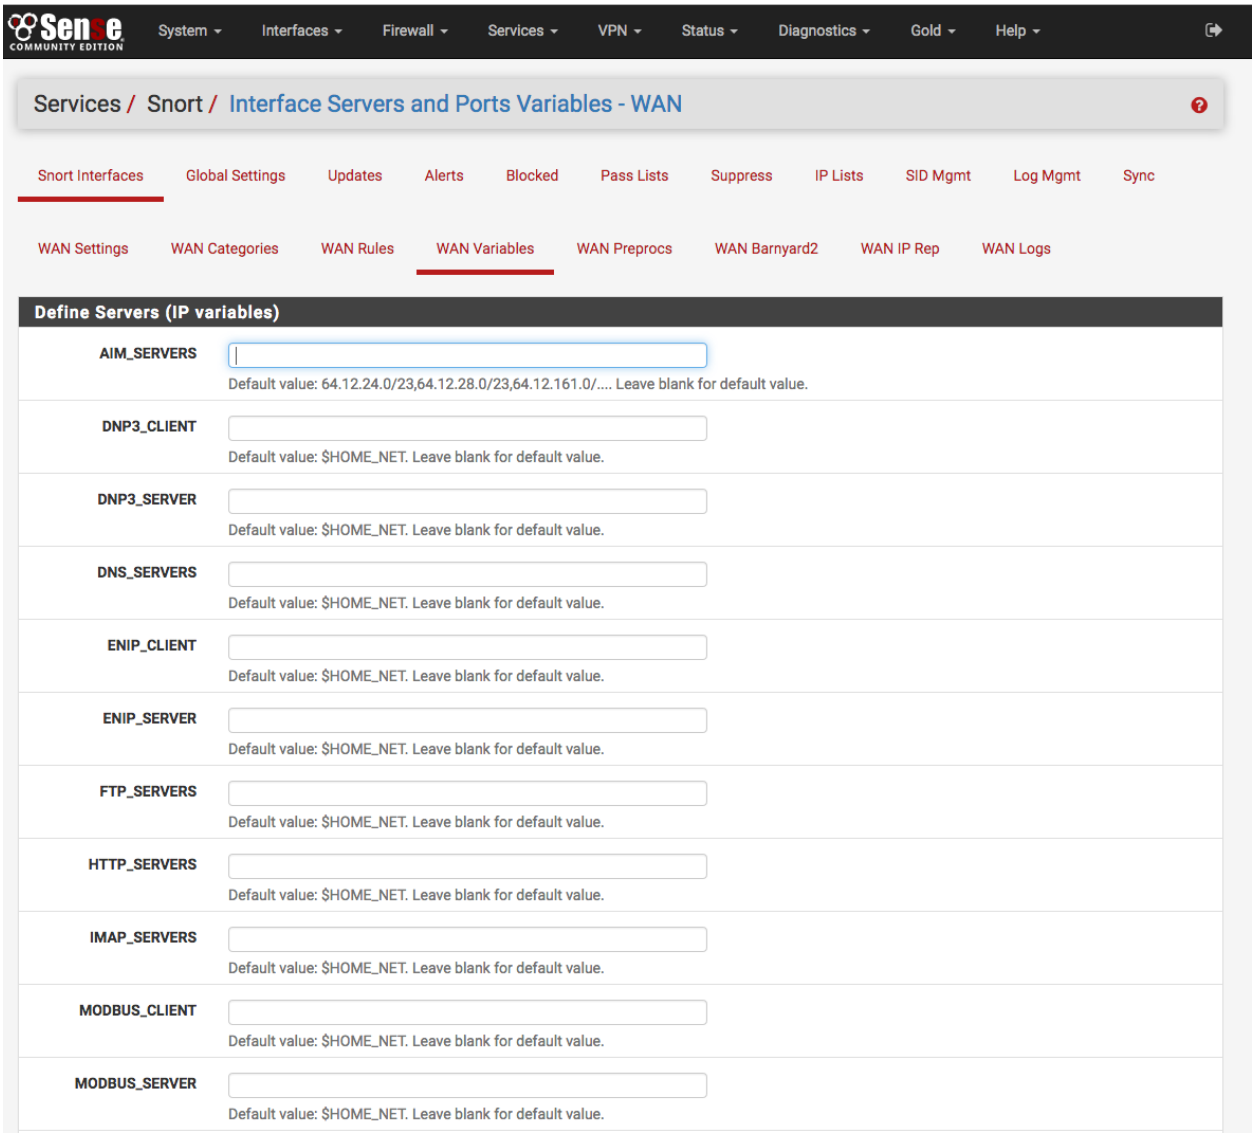
Snort Server Definitions

Define servers to protect and improve performance

The **Variables** tab is where the specific types of hosts on the network are configured. For example, the specific IP addresses or network ranges containing web servers to protect may be defined. This can make Snort more efficient because it won't waste time scanning for web server threats on IP addresses where web servers do not exist. Similarly, Snort performance can be optimized by instructing it which addresses contain other critical servers such as SMTP, POP, DNS, etc.

The exact ports or port ranges used for certain services on the network may also be specified.

Each value entered on this page can only be an existing Alias. Start typing the name of the Alias into a text box and a drop-down selection of matching entries will appear for selection. Aliases are created under **Firewall > Aliases** from the menu.



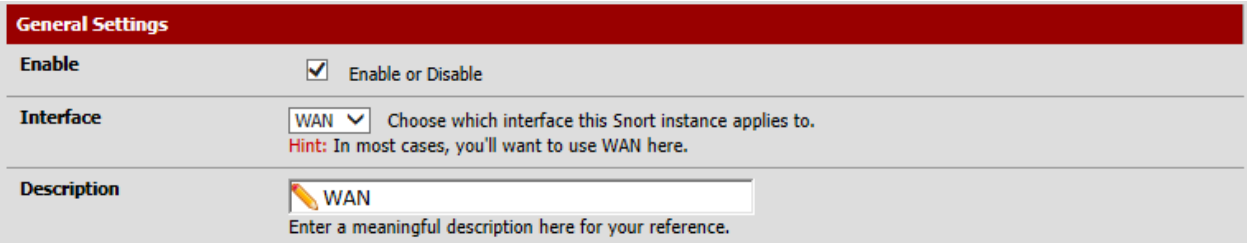
Snort interface Settings

General Settings

Enable
Used to enable or disable Snort on the selected interface. Snort is enabled on the interface when this box is checked.

Interface
Used to choose which physical firewall interface this Snort instance protects.

Description
Used to provide an optional friendly name for the interface.



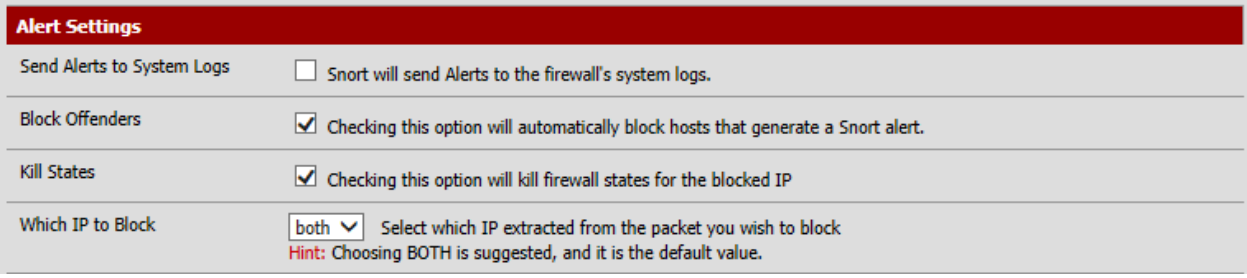
Alert Settings

Send Alerts to System Logs
When checked, all Snort alerts will be copied to the system log on the firewall.

Block Offenders
When checked, Snort will automatically insert a firewall block of the host generating an alert.

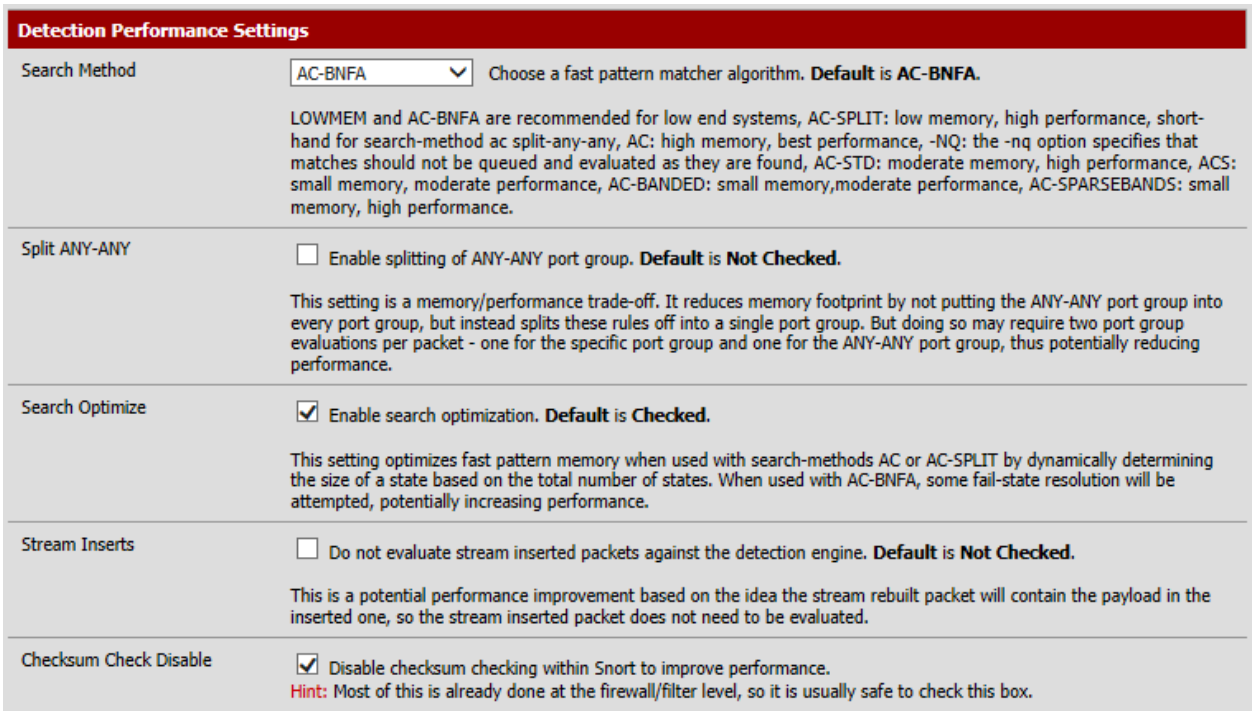
Kill States
When checked, Snort will kill all existing state table entries for the IP address it blocks. This should generally be enabled (box checked).

Which IP to Block
This determines which IP address extracted from the packet that generated an alert will be blocked. The choices are SOURCE, DESTINATION or BOTH. BOTH is the recommended default.



Detection Performance Settings

Search Method
Used to select the pattern matcher algorithm used by Snort in the signature detection engine.

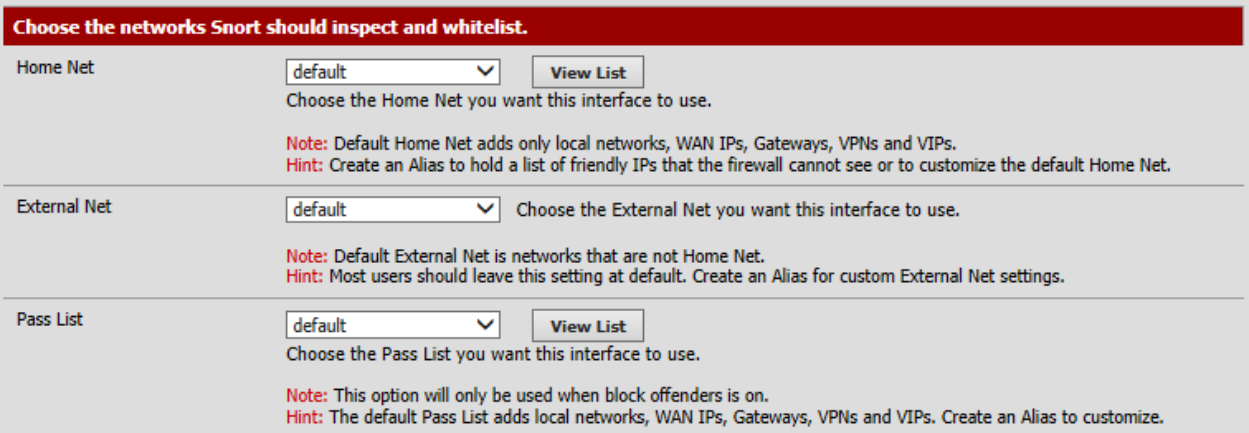


Choose the networks Snort should inspect and whitelist

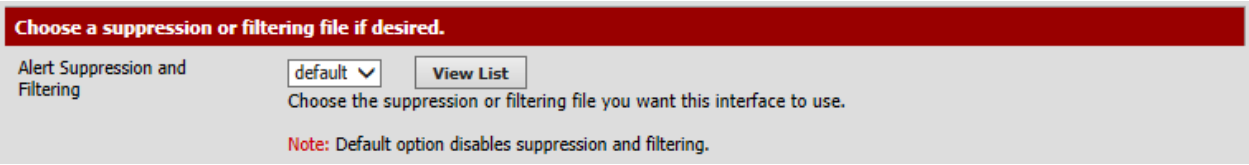
Home Net
: Selects the network Snort will use as the HOME_NET variable. Default is the recommended choice and contains the firewall WAN IP address and WAN gateway, all networks locally-attached to a firewall interface, the configured DNS servers, VPN addresses and Virtual IP addresses. Additional HOME_NET networks may be created on the IP LISTS tab, and then return to this tab to assign them to the Snort interface by selecting the appropriate list in the drop-down selector. View the contents of the selected list by clicking the **View List** button.

External Net
: Selects the network will use as the EXTERNAL_NET variable. Default is the recommended choice and contains all networks not included in HOME_NET. Create additional EXTERNAL_NET networks on the IP LISTS tab, and then return to this tab to assign them to the Snort interface by selecting the appropriate list in the drop-down selector.

Pass List
: Selects the networks and IP addresses that Snort will never block. These represent "trusted hosts". Even if a trusted host generates a Snort alert, it will not be blocked if the IP address is on a Pass List. The default Pass List contains the same addresses as HOME_NET. Create additional pass lists on the IP LISTS tab, and then return to this tab to assign them to the Snort interface by selecting the appropriate list in the drop-down selector. Snort must be restarted on the interface when making changes to the Pass List. View the contents of the selected list by clicking the **View List** button.



Choose a suppression or filtering file if desired

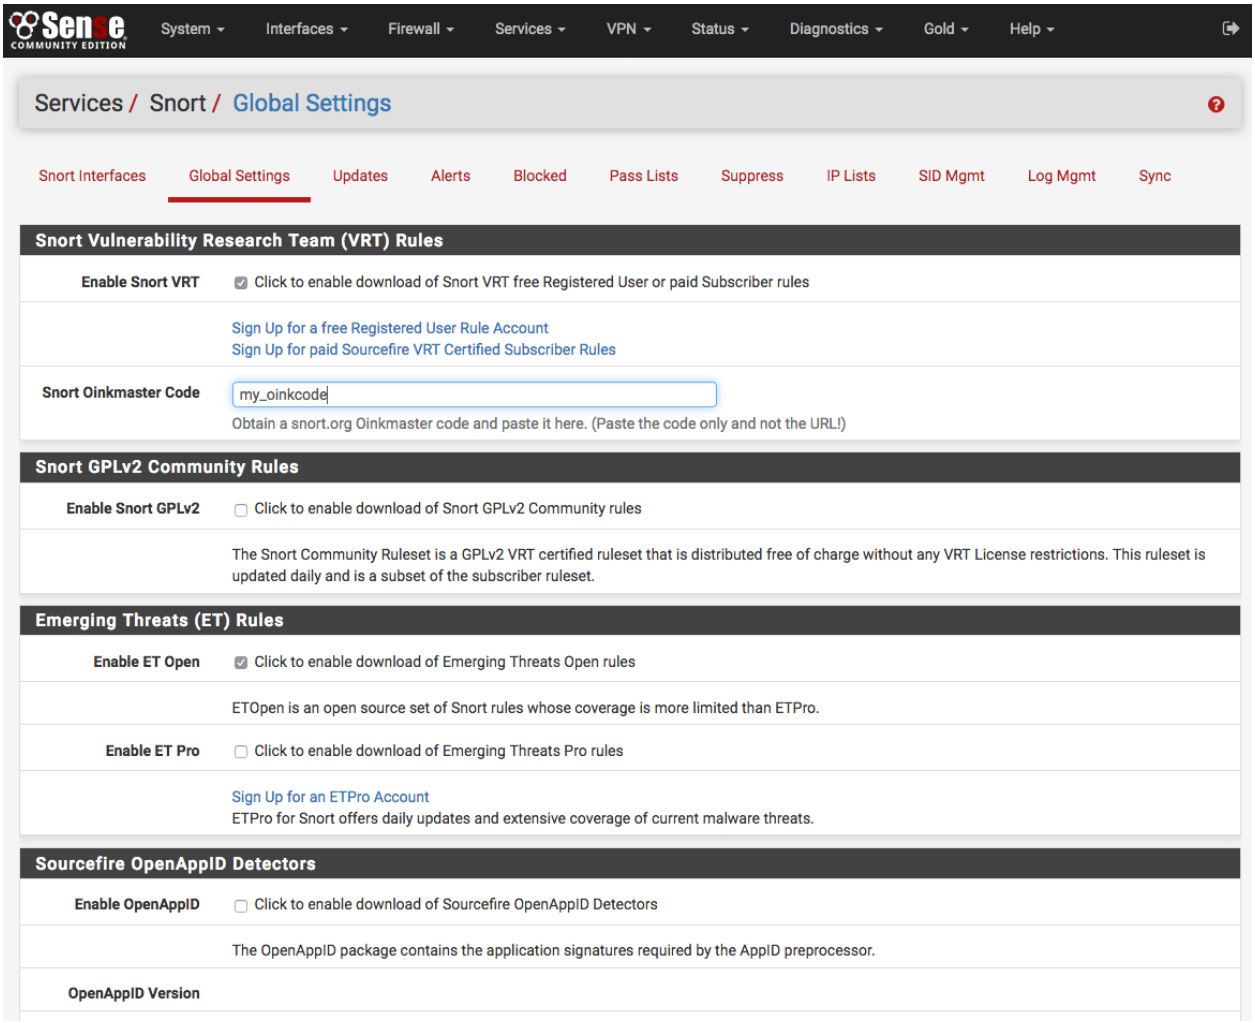


Snort interface Global Settings

This tab is used to enable rule set packages for download, configure the rules package update interval and start time, configure Snort logging directory size limits and determine whether Snort settings are saved when the package is removed from the system.

Please Choose The Type Of Rules To Download

More than one rule set may be enabled for download, but note the following caveats. If there is a paid subscription for the Snort VRT rules, then all the Snort GPLv2 Community rules are automatically included within the file downloaded with the Snort VRT rules; therefore do not enable the GPLv2 Community rules if there is a paid-subscriber account for the Snort VRT rules. All the Emerging Threats Open rules are included within the paid subscription for the Emerging Threats Pro rules. If the Emerging Threats Pro rules are enabled, the Emerging Threats Open rules are automatically disabled.



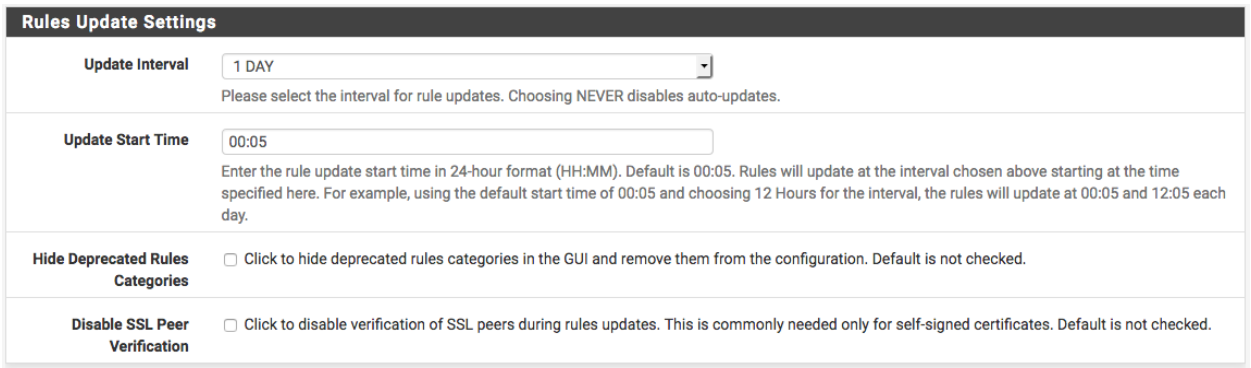
To use the Snort VRT rules package, check the **Install Snort VRT rules** checkbox and then enter the Oinkmaster code in the text box that appears.

To use the ETPro rules package, check the box next to **ETPro** and then enter the ETPro subscription code in the text box that appears.

Rules Update Settings

Use the **Update Interval:** drop-down selector to choose the periodicity for checking for updates to the enabled rules packages. When any value other than NEVER is selected, the **Update Start Time** text box is available for entering a start time in 24-hour format using hours and minutes only.

In most cases every 12 hours is a good choice. The update start time can be customized if desired. Enter the time as hours and minutes in 24-hour time format. The default start time is 3 minutes past midnight local time. So with a 12-hour update interval selected, Snort will check the Snort VRT or Emerging Threats websites at 3 minutes past midnight and 3 minutes past noon each day for any posted rule package updates.



General Settings

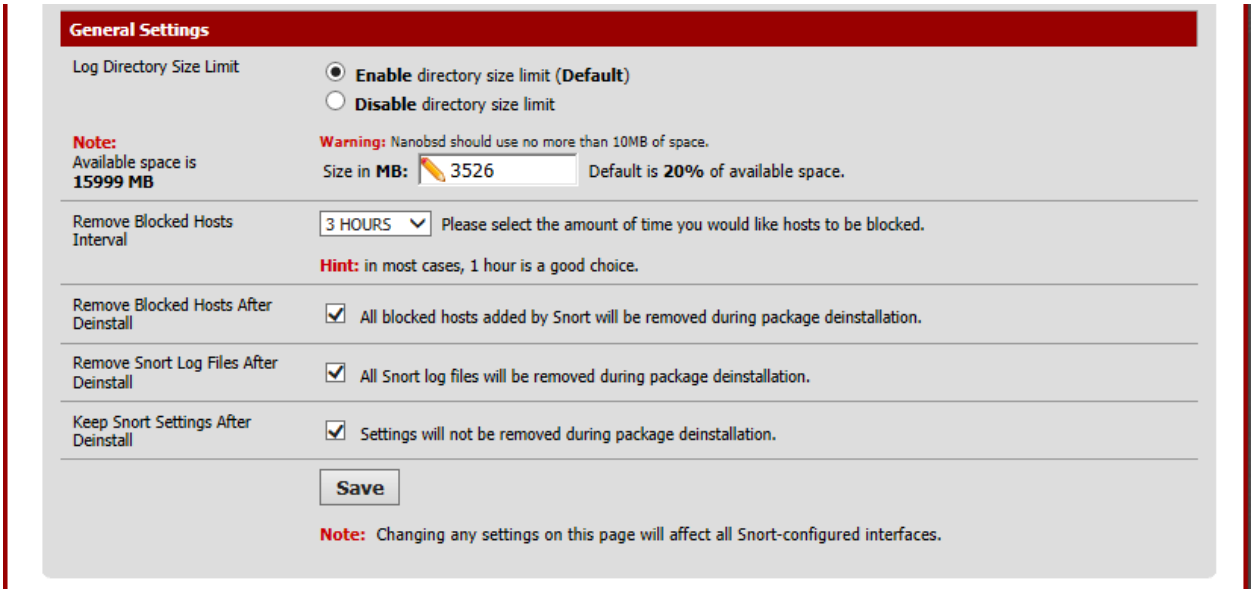
Log Directory Size Limit
When enabled, sets an absolute hard upper limit on the total size of the Snort logging subdirectory in `/var/log/snort`. This can prevent Snort from filling up the `/var` volume on the firewall. When the Snort logging directory size (the total size of all files within the Snort log directory tree) exceed the value set, all files are automatically pruned (deleted) and the Snort process is signaled to soft-restart and resynchronize logging. The default size limit is 20% of the available space on the volume. This may be overridden by setting a value in megabytes (MB) in the text box provided.

Remove Blocked Hosts Interval
Controls how long Snort-blocked IP addresses must be inactive before being cleared. Once per interval specified, Snort executes a cron job that tests all the IP addresses it has inserted into the firewall's block table for activity. IP addresses that have had no further network activity within the time specified are removed from the block table.

Remove Blocked Hosts After Deinstall
Determines whether Snort-blocked IP addresses are automatically removed when the Snort package is uninstalled.

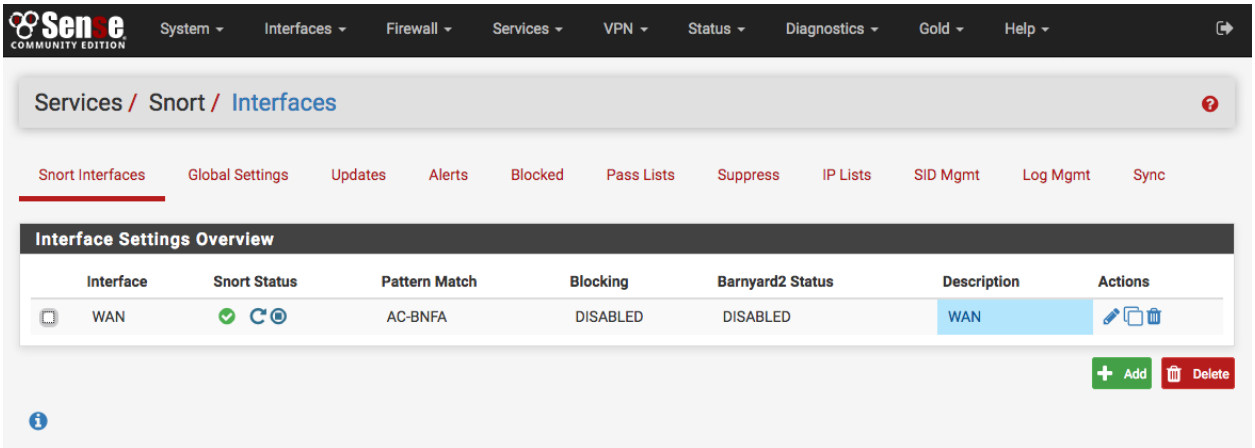
Remove Snort Log Files After Deinstall
Determines whether log files generated by Snort are retained or removed when the Snort package is removed.

Keep Snort Settings After Deinstall
Controls whether the Snort configuration is retained when the Snort package is removed.

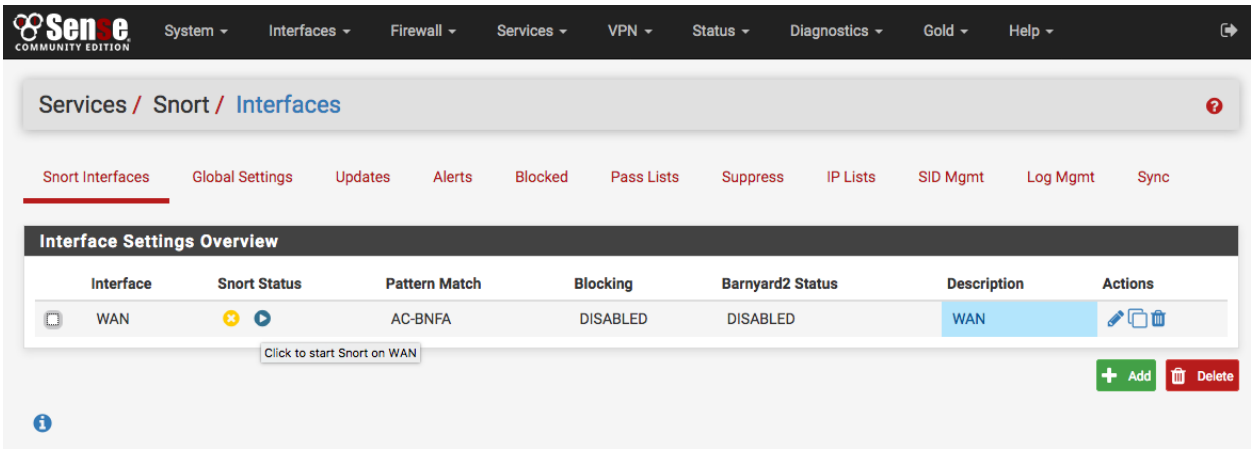


Snort Interfaces

The **Snort Interfaces** tab is where one can add, edit or delete a Snort instance from a physical network interface. A snort instance can also manually be started and stopped. If *Barnyard2* is configured on an interface, it can also be started or stopped.



The green icon indicates a running Snort process for the interface. To stop a running process, click the  icon and wait for it to change to  as shown below.



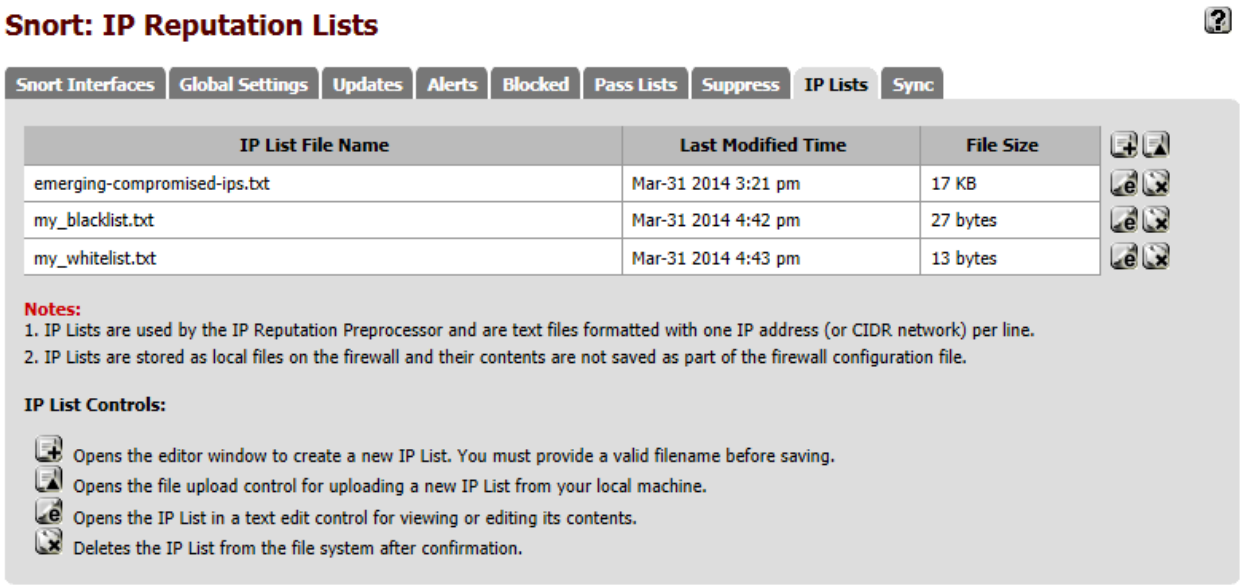
To add a new Snort configuration for an interface, click **Add**.

To edit an existing Snort configuration, click edit icon.

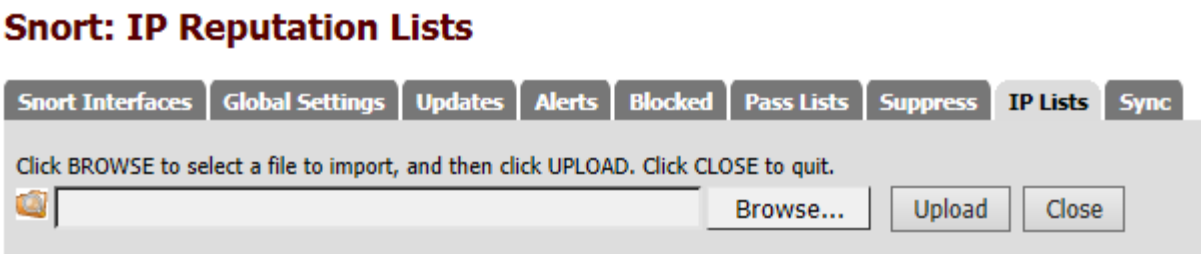
To delete an existing Snort configuration, click the checkbox on the left of the interface to select it, then click **Delete**. A prompt for confirmation will appear before deleting the chosen interface. Multiple interfaces may be selected and deleted at once.

Managing Snort IP Address Lists

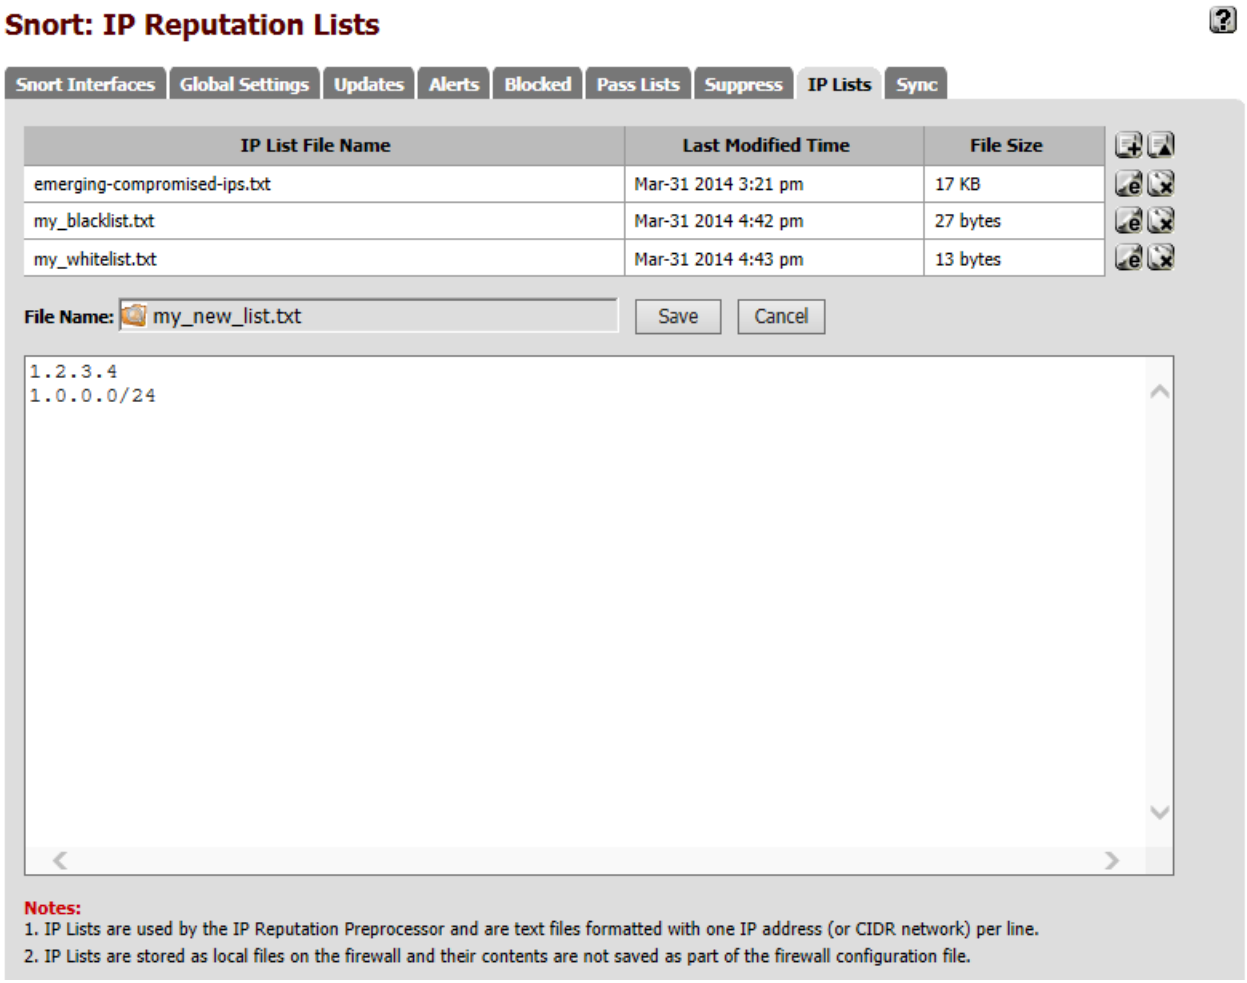
Use this tab to manage the IP lists files for the IP Reputation preprocessor. IP lists are text-format files containing one IP address or network (expressed in CIDR notation) per line.



To upload an IP list file to the firewall, click the  icon to open the file upload dialog as shown below. Browse to the file on the local machine using the **BROWSE** button, then click the **UPLOAD** button to upload the file to the firewall for use by the IP Reputation preprocessor in Snort.



- Click the  icon to create a new IP list.
- Click the  icon beside the list to edit an existing IP list.
- Click the  icon to save changes to the list when finished.
- Click **Cancel** to abandon any changes.



Snort IP Address Reputation Preprocessor

This tab allows configuration of the parameters specific to the IP Reputation preprocessor on the interface. It also allows the assignment of blacklist and whitelist files of IP addresses to the interface.

The available fields are:

Enable
When checked, the IP Reputation preprocessor is active on this Snort instance.

Memory Cap
Sets the amount of system memory in megabytes (MB) to reserve for storage of the IP lists associated with this preprocessor. The default is 500 MB and should be sufficient for most installations.

Scan Local
When checked, Snort will include RFC 1918 IP addresses ranges when comparing IP addresses to the blacklists and whitelists. If an RFC 1918 IP addresses is in the whitelist files, or some are blacklist files, then this option should be enabled. The default is disabled.

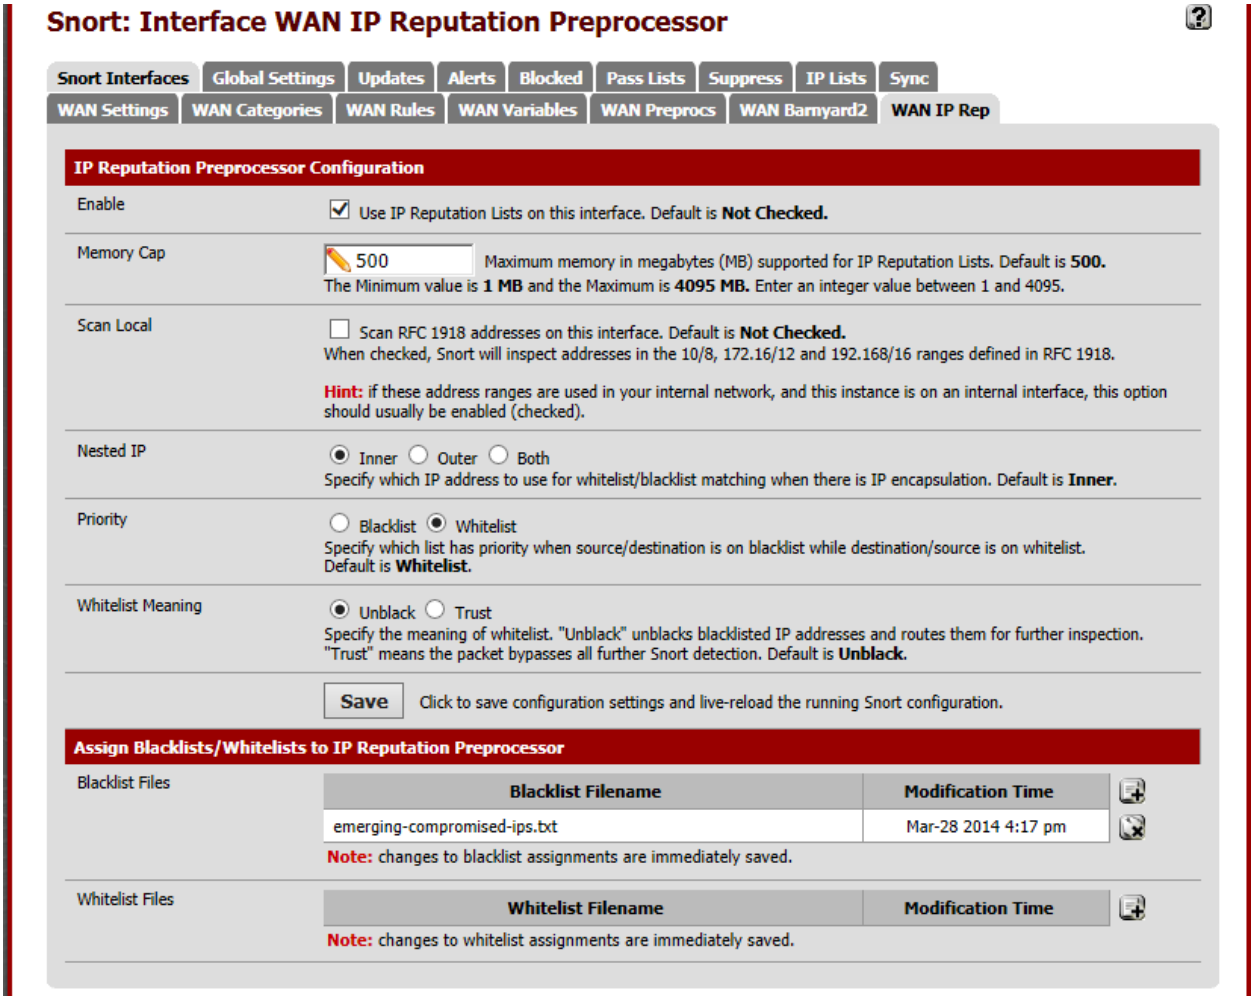
Nested IP
This tells Snort which IP address to compare to the IP lists in the whitelist and blacklist files when there is IP encapsulation. The default is **Inner**.

Priority
Instructs Snort which IP list has priority when the source and destination IP addresses of a packet are each on separate IP lists. For example, if the source IP address is on a blacklist while the destination IP address is on a whitelist, this option tells Snort whether to block the traffic if blacklist has priority, or pass the traffic if whitelist has priority.

Whitelist Meaning
This tells Snort what action to take with whitelisted IP addresses. The two options are **Un-black** and **Trust**.

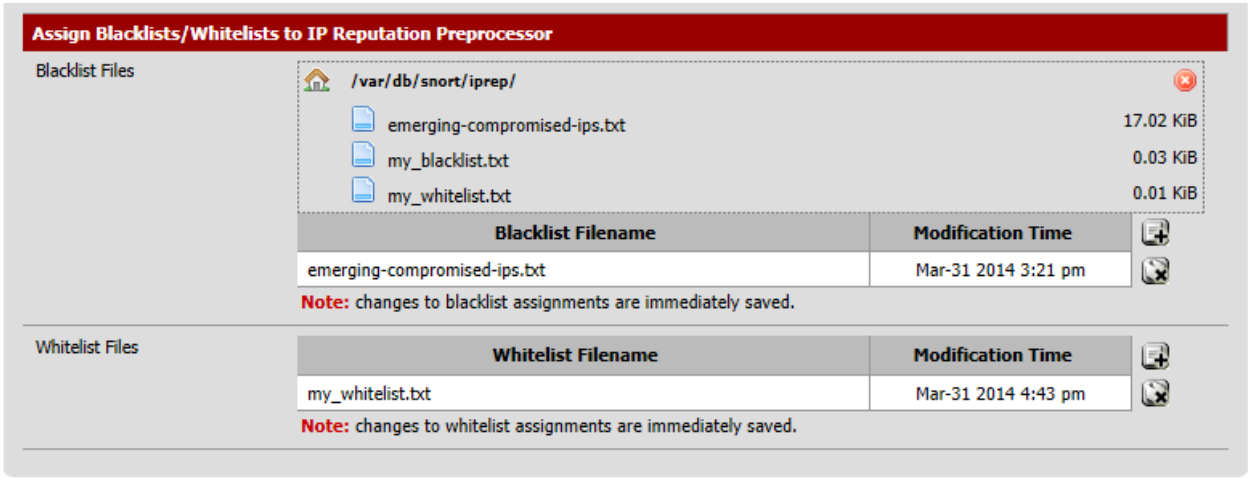
Un-black
A blacklisted IP which is listed in the whitelist is not immediately blocked. Instead, it is routed through the Snort detection engine for normal inspection. If it generates no alerts, the traffic is allowed. If the inspection results in a Snort alert for the traffic, it will be blocked.

Trust
Any IP address on the whitelist (including any that may also be on a blacklist) is immediately allowed to pass with no further inspection. Caution should be exercised when using the Trust mode of operation to insure the IP addresses on the whitelist are in fact trustworthy.



The  and  icons at the bottom of the page are used to assign or remove blacklist and whitelist files to or from the interface.

Click the  icon to open a file selection dialog. Choose an IP list file from the list by clicking on the name.



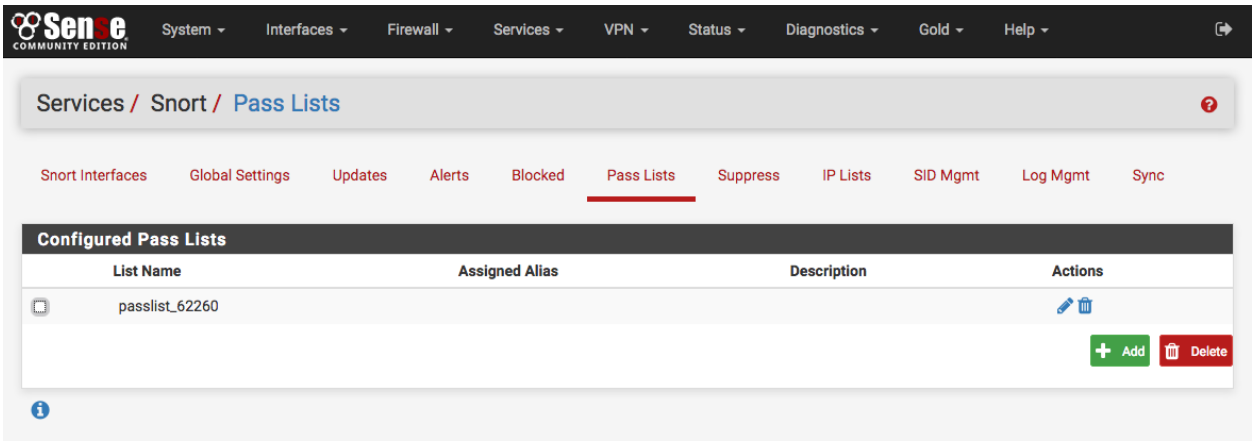
Snort IP Address Reputation

Coming soon…

Snort Pass Lists

Pass Lists are lists of IP addresses that Snort should never block. Pass lists can be created and managed on the **Pass Lists** tab. When an IP address is listed on a Pass List, Snort will never insert a block on that address even when malicious traffic is detected.

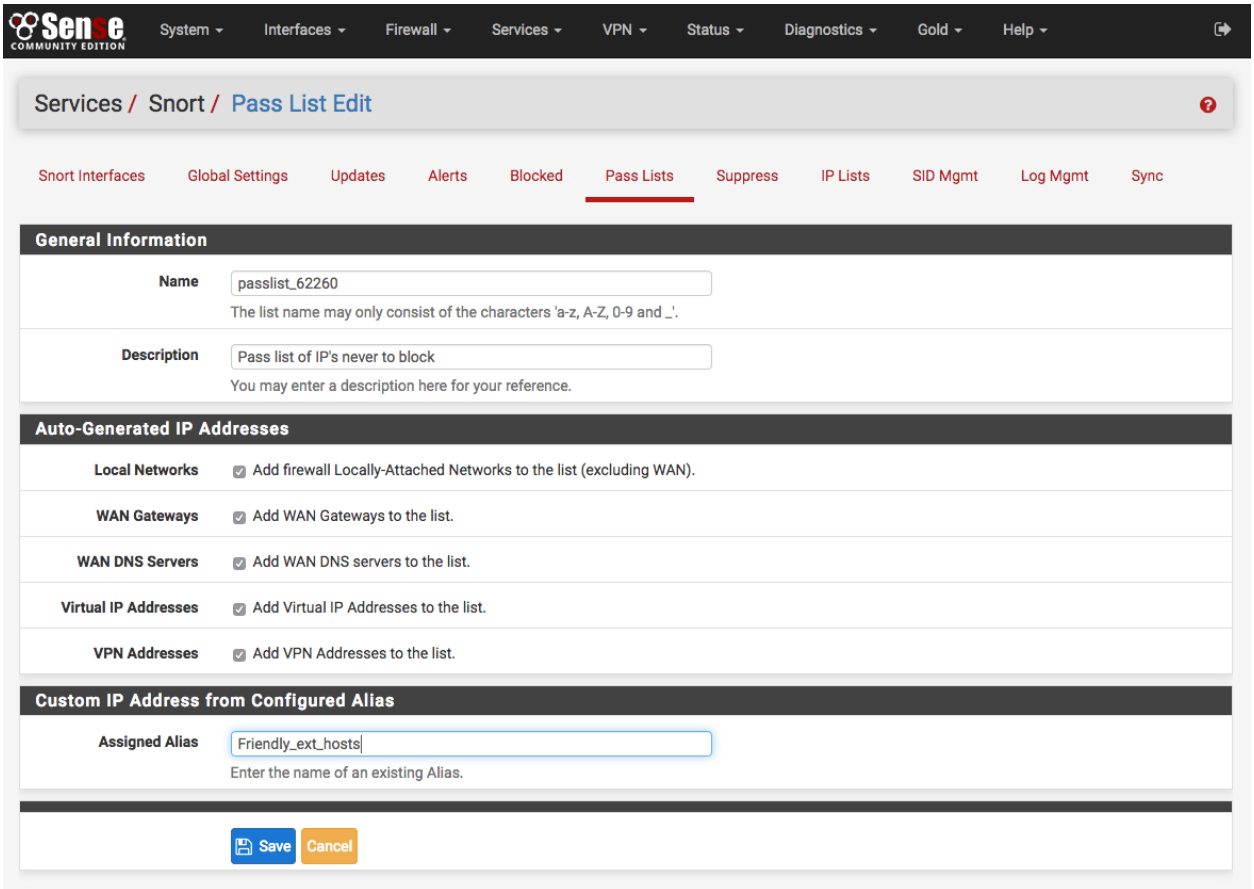
To create a new Pass List, click the  icon. To edit an existing Pass List, click the  icon. To delete a Pass List, click the  icon. Note that a Pass List cannot be deleted if it is currently assigned to one or more Snort interfaces.



A default Pass List is automatically generated by Snort for every interface, and this default list is used when no other list is specified. Assign Pass Lists to an interface on the **Interface Settings** tab.

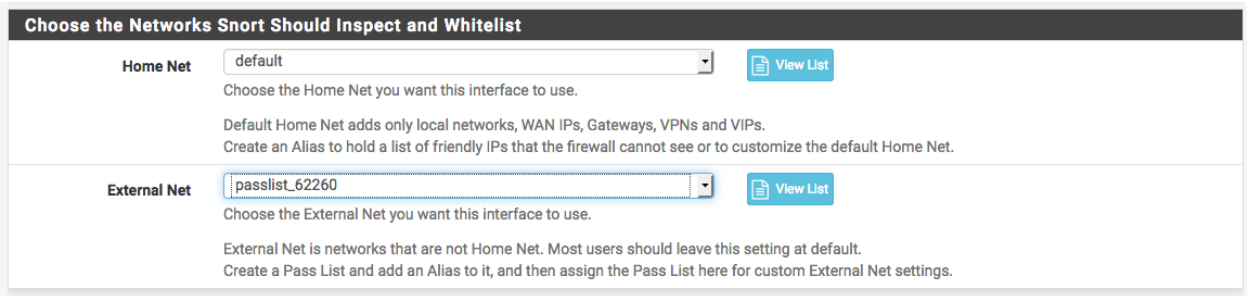
Customized Pass Lists can be created and then assigned to an interface. This might be needed when trusted external hosts are needed that are not located on networks directly connected to the firewall. To add external hosts in this manner, first create an Alias under **Firewall > Aliases** and then assign that alias to the **Assigned Aliases:** field. In the example shown below, the alias "*Friendly_ext_hosts*" has been assigned. This alias would contain the IP addresses of the trusted external hosts.

When creating a custom Pass List, leave all the auto-generated IP addresses checked in the **Add auto-generated IP addresses** section. Not selecting the checkboxes in this section can lead to blocking of critical addresses including the firewall interfaces themselves. This could result in being locked out of the firewall over the network! Only uncheck boxes in this section when a valid need is present.



Click the **ALIASES** button to open a window showing previously defined aliases for selection. Remember to click **SAVE** to save changes.

Note: Remember that simply creating a Pass List is only the first step! Go to the **Interface Settings** tab for the Snort interface and assign the newly created Pass List as shown below. After assigning and saving the new Pass List, restart Snort on the affected interface to pick up the change.



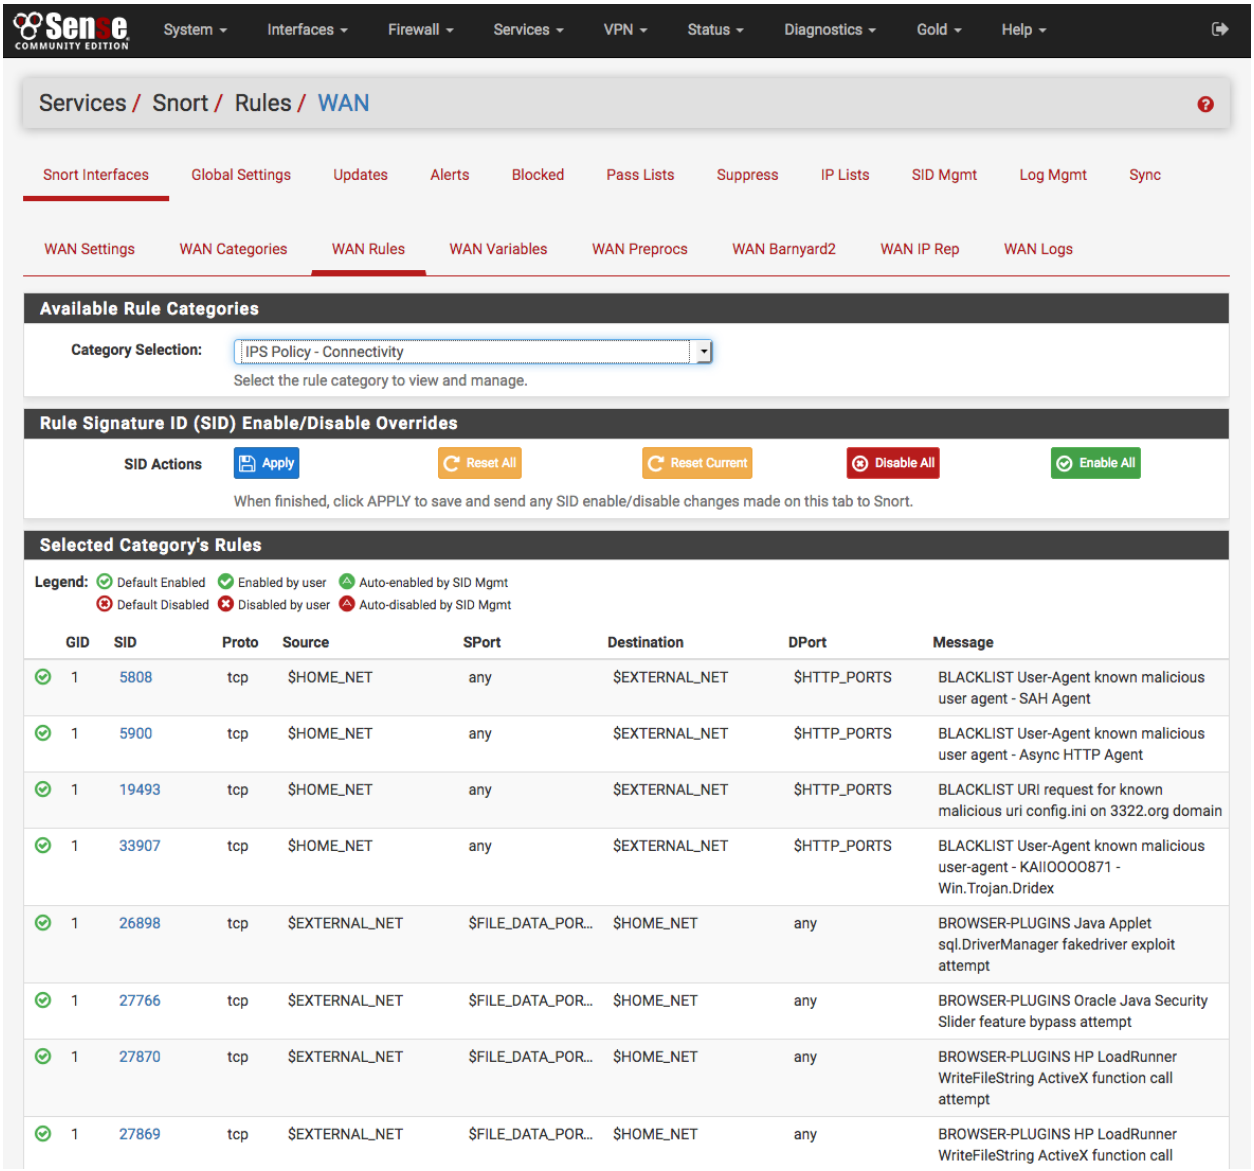
Snort Rules

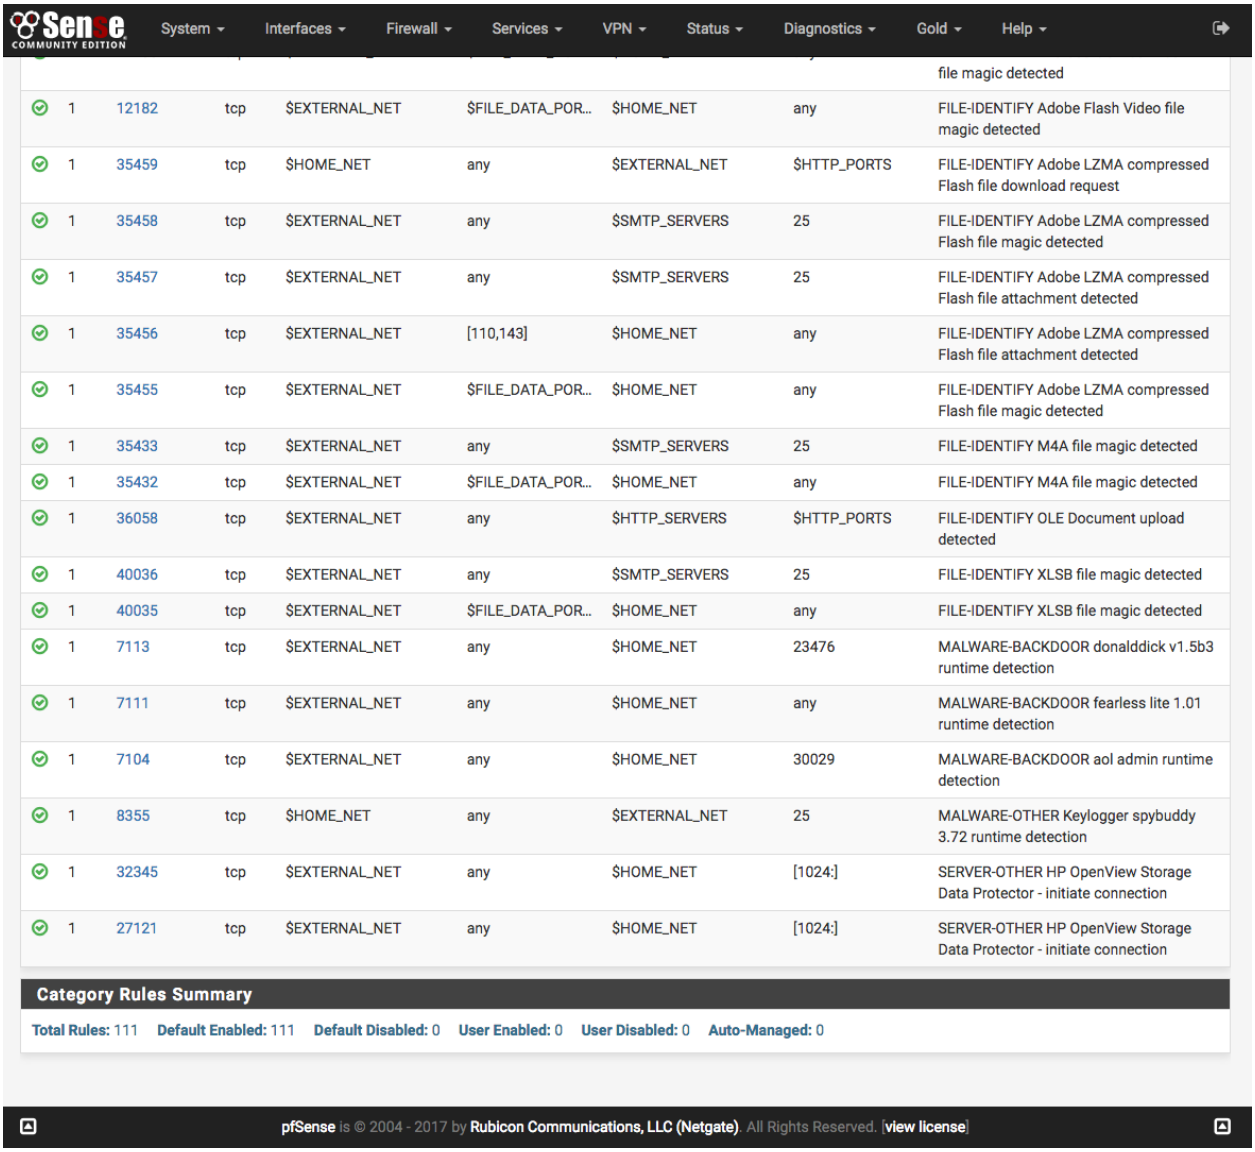
Rules

Use the **Rules** tab for the interface to configure individual rules in the enabled categories. Generally this page is only used to disable particular rules that may be generating too many false positives in a network environment. Be sure they are in fact truly false positives before taking the step of disabling a Snort rule!

- Select a rule category from the **Category:** drop-down to view all the assigned rules.
- Click the or icon at the far-left of a row to toggle the rule's state from enabled to disabled.
- Click or to toggle from disabled to enabled.

  The icon will change to indicate the state of the rule. There is a legend at the top of the rule list which describes the rule state icons.



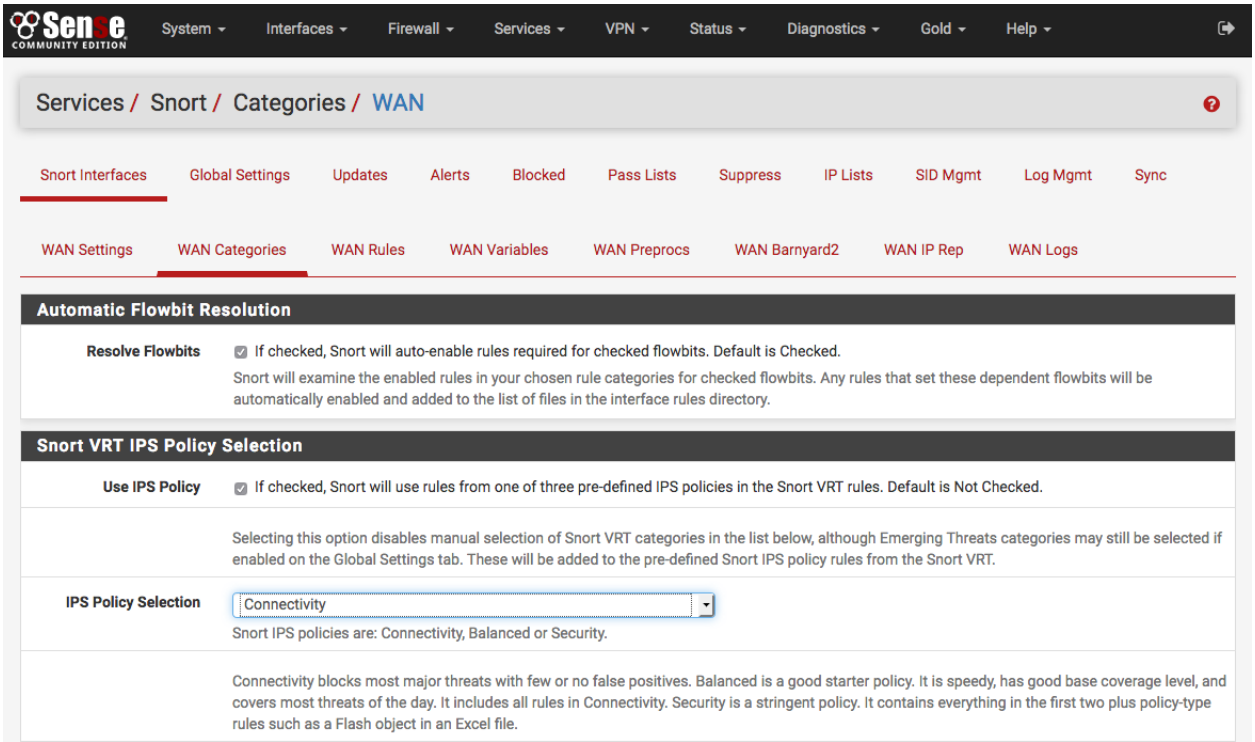


Snort Rulesets

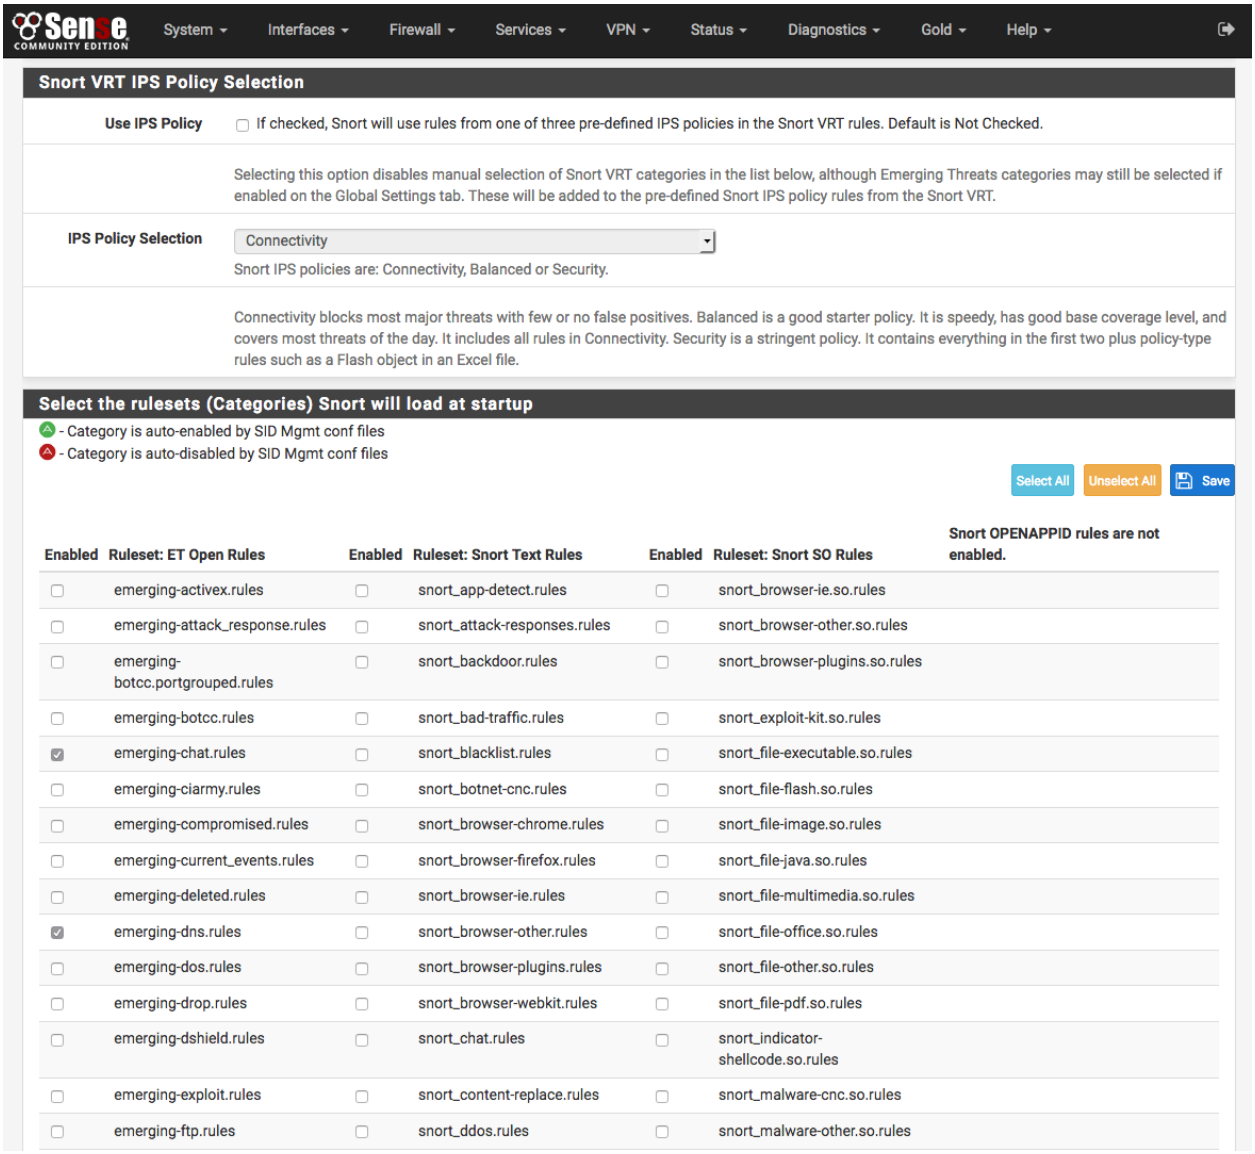
Categories

If a Snort VRT Oinkmaster code has been obtained (either free registered user or the paid subscription), and the Snort VRT rules have been enabled, and the Oinkmaster code has been entered on the Global Settings tab then the option of choosing from among three pre-configured IPS policies is available. These greatly simplify the process of choosing enforcing rules for Snort to use when inspecting traffic. The IPS policies are only available when the Snort VRT rules are enabled.

The three Snort VRT IPS Policies are: (1) Connectivity, (2) Balanced and (3) Security. These are listed in order of increasing security. However, resist the temptation to immediately jump to the most secure “Security” policy if new to using Snort. False positives can frequently occur with the more secure policies, and careful tuning by an experienced administrator may be required. So if new to Snort, then using the less restrictive “Connectivity” policy in non-blocking mode is recommended as a starting point. Once experience with Snort has been gained in the network environment, blocking mode can be enabled and then move up to more restrictive IPS policies.



If the Snort VRT rules are not enabled, or to use any of the other rule packages, then make the rule category selections by checking the checkboxes beside the rule categories to use.

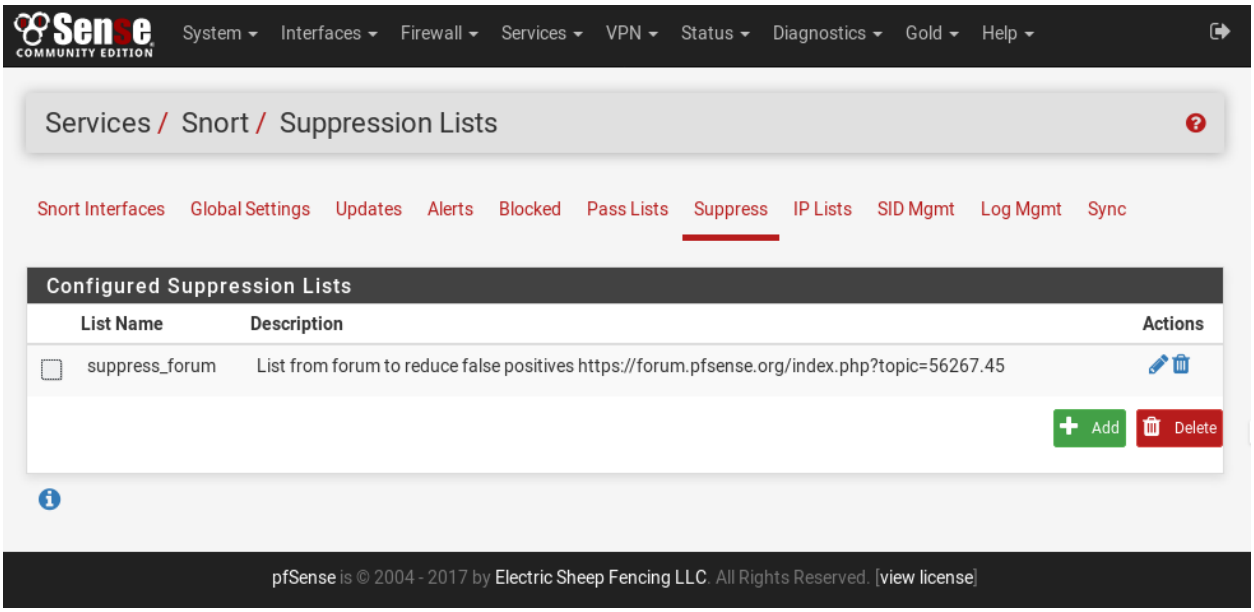


Click  **Save** to store the current selection and build the rules file for Snort to use.

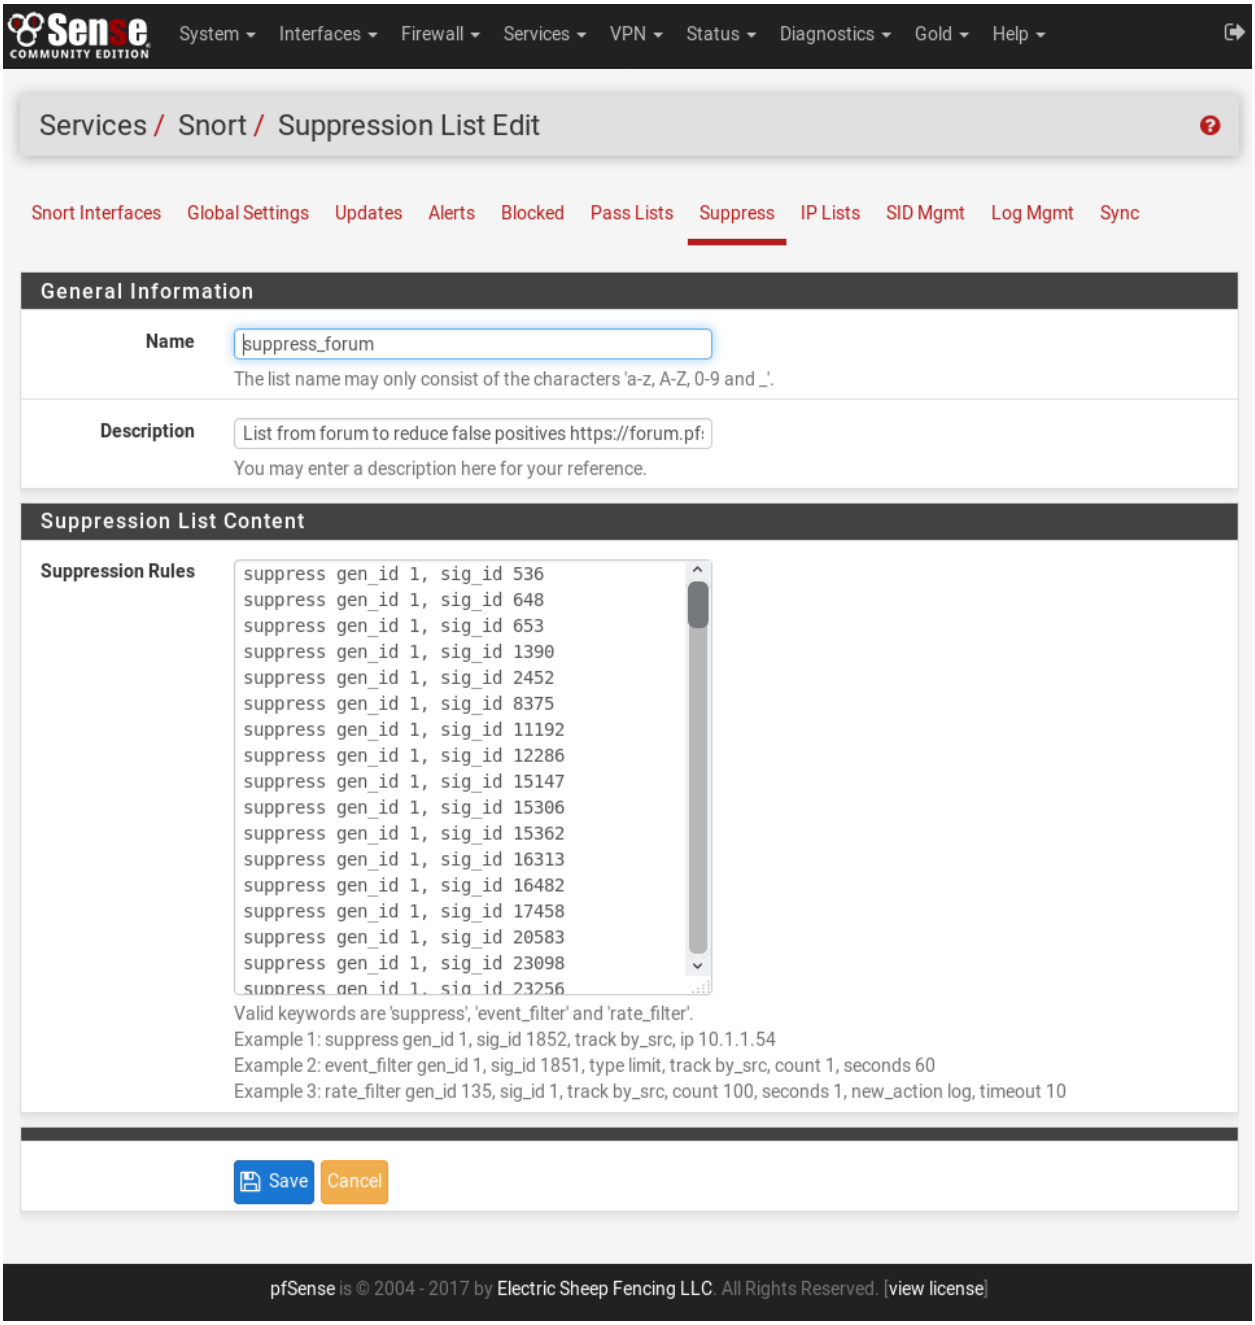
Snort Suppression Lists

Alert Thresholding and Suppression

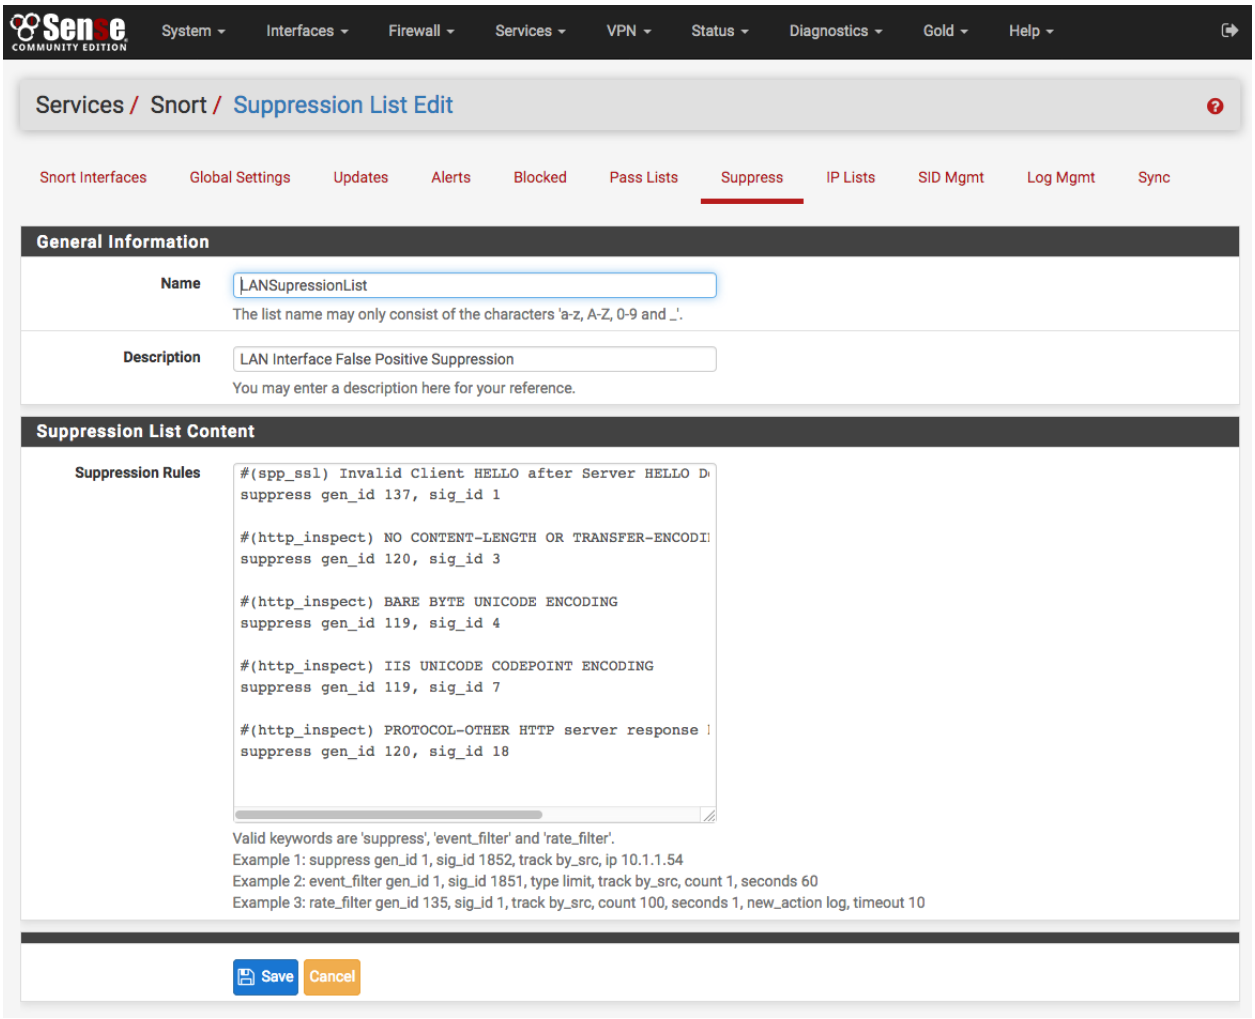
Suppression Lists allow control over the alerts generated by Snort rules. When an alert is suppressed, then Snort no longer logs an alert entry (or blocks the IP address if block offenders is enabled) when a particular rule fires. Snort still inspects all network traffic against the rule, but even when traffic matches the rule signature, no alert will be generated. This is different from disabling a rule. When a rule is disabled, Snort no longer tries to match it to any network traffic. Suppressing a rule might be done in lieu of disabling the rule to stop alerts based on either the source or destination IP. For example, to suppress the alert when traffic from a particular trusted IP address is the source. If any other IP is the source or destination of the traffic, the rule may still be desired. To eliminate all alerts from the rule, then it is more efficient to simply disable the rule rather than to suppress it. Disabling the rule will remove it from the list of match rules in Snort and therefore makes for less work Snort has to do.



On the Suppress List Edit page, suppress lists may be manually added or edited. It is usually easier and faster to add suppress list entries by clicking the + icons shown with the alert entries on the **Alerts** tab. Remember to click the **SAVE** button to save changes when manually editing Suppress List entries.



Lists with comments are easier to manipulate and fine tune. Neither screenshot shows IP address suffix in a suppress entry.



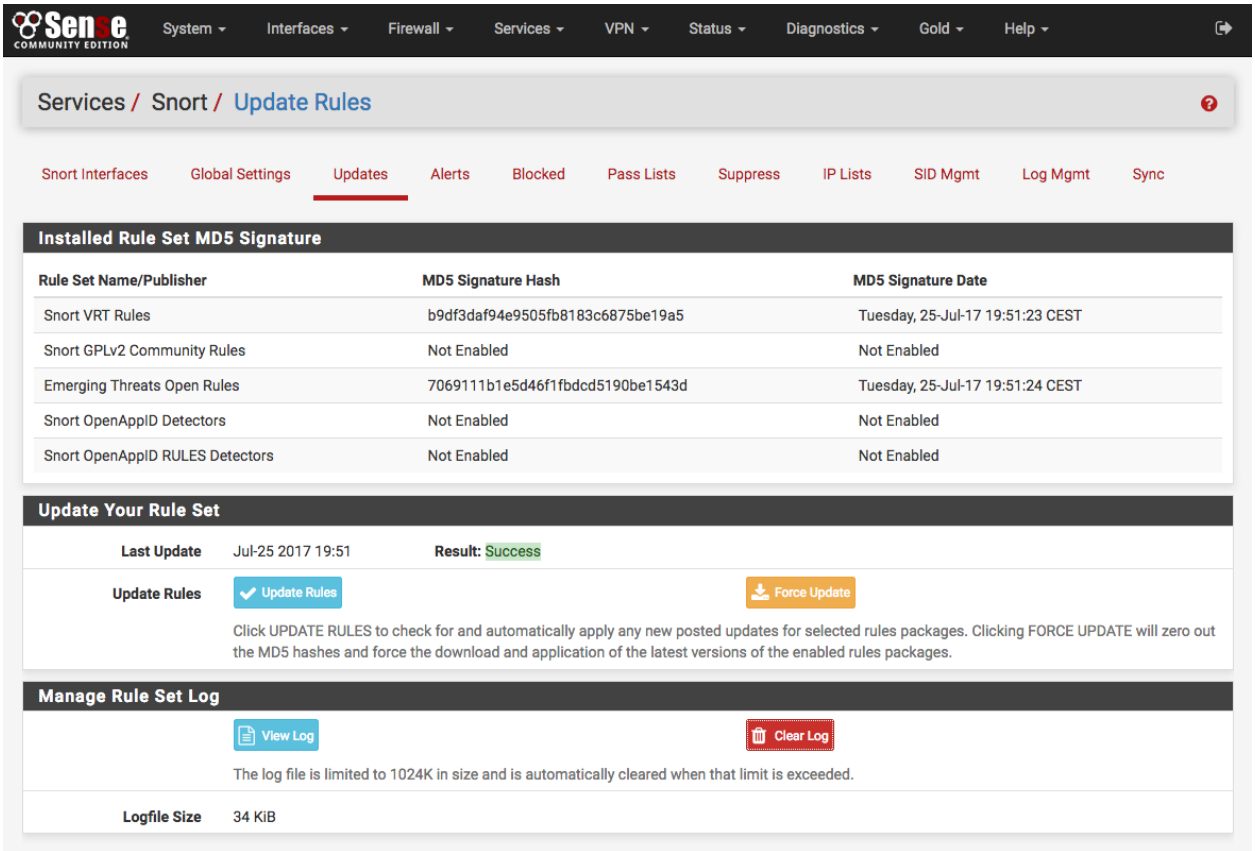
Updating Snort

Update the rules

The **Updates** tab is used to check the status of downloaded rules packages and to download new updates. The table shows the available rule packages and their current status (not enabled, not downloaded, or a valid MD5 checksum and date).

Click on the **Update Rules** button to download the latest rule package updates. If there is a newer set of packaged rules on the vendor website, it will be downloaded and installed. The determination is made by comparing the MD5 of the local file with that of the remote file on the vendor website. If there is a mismatch, a new file is downloaded. The **FORCE** button can be used to force download of the rule packages from the vendor website no matter how the MD5 hash tests out.

In the screenshot below, the Snort VRT and Emerging Threats Open rule packages have been successfully downloaded. The calculated MD5 hash and the file download date and time are shown. Also note the last update time and result are shown in the center of the page.



See also:

Troubleshooting Snort Rule Updates

Tip: For more information or to get help, check out the IDS/IPS category on the Netgate forum.

30.24 Stunnel package

The stunnel program is designed to work as an SSL encryption wrapper between remote client and local (`inetd` managed) or remote servers. It can be used to add SSL functionality to commonly used `inetd` daemons like POP2, POP3, and IMAP servers without any changes in the program's code.

It will negotiate an SSL connection using the **OpenSSL** or **SSLeay** libraries. It calls the underlying crypto libraries, allowing stunnel to use any cryptographic algorithms compiled into the crypto package.

Note: The pfSense® software package implements only a subset of the configuration options available in stunnel. For more advanced configurations, please consider configuring stunnel manually on the firewall, run it in a dedicated jail, or on a different system.

The package has two configuration screens (tabs):

- Tunnel definitions
- Certificates

30.24.1 Tunnels

For each tunnel, the following options are available:

- Listening socket IP address and port.
- Certificate to use for the listening socket.
- Target IP address and port.
- IP address to bind to when connecting to the target.

If no certificate is specified for a tunnel, the default certificate will be used. This is a self-signed certificate which is generated upon package (re)installation, and is not suited for production use.

30.24.2 Certificates

Certificates are managed in the simplest possible way, by requiring the user to provide RSA key and certificates/chains in PEM format. The **Certificates** tab will list the configured certificates along with status information, indicating whether the certificate is valid, will expire soon, or is already expired. A sanity check is also performed to make sure the key and certificate matches.

Note that for private certificates and certain commercial ones (Extended Validation), a complete certificate chain may be required. This is to ensure that the client is able to verify the certificate validity. A chain should be built in the following way:

1. Root certificate of the certificate issuer/CA
2. Any intermediate certificates between the root and the server certificate
3. Server certificate

See also:

Refer to the stunnel documentation for more information on how to format a certificate chain.

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.25 Sudo Package

The sudo package configures basic rules for allowing unprivileged users (i.e. anyone but `root/admin`) to run commands as `root` or another user/group in the shell.

Once the `sudo` package is installed, it is located at **System > sudo** in the GUI.

30.25.1 Sudo Settings

The package allows multiple entries for privileges. To add a new entry, click  **Add**, and fill in the settings:

User/Group
: The user or group name to which this privilege is being granted.

  The list includes users and groups defined in the GUI as well as those from the operating system (e.g. daemon users and groups added by packages).

Run As
The user or group name under which the command will be run.

In most cases this is `root`, so that users may run commands as `root` without knowing the `root/admin` credentials.

No Password
Controls whether the user is not prompted for their own password when executing commands using `sudo`.

This is unset by default, so users are prompted for their password when running `sudo`. `sudo` caches credentials in a login session for at least five minutes after each execution of `sudo` so that the user is not prompted on each attempt. Users can execute `sudo` without reauthenticating inside that time frame, but if they stop for five minutes they will be prompted again on the next run.

When set, the user is not prompted for their password when running `sudo`. This is less secure, but more convenient. If `sudo` is invoked non-interactively, such as from a cron script, this is required as there is no way for a user to enter their credentials.

Command List
A list of commands the **User/Group** can run.

See also:

More information on the full command options may be found in the sudoers manual.

By default, the command is `ALL` meaning the user can run any commands. Leaving the commands field blank assumes `ALL`.

A comma-separated list of one or more commands can be supplied to limit the user to individual binaries. Full paths to commands are required by `sudo` to ensure the user is properly restricted to specific binaries or scripts.

If parameters are specified after a command, they will be required. To disallow running a command with parameters, add `""` after the command.

Custom Configuration
This option controls whether `sudo` will read additional configuration files from `/usr/local/etc/sudoers.d`.

> **Warning:** Including custom configuration files allows options to be set which are not implemented in the GUI, but these files can be a potential security risk, and they are not included in backups.

The setting can be one of:

Do Not Include
`sudo` will **not** include additional configuration files.

Include at Start
`sudo` will include additional configuration files **before** the GUI settings.

Include at End
`sudo` will include additional configuration files **after** the GUI settings.

30.25.2 Sudo Examples

Example 1

Allow `bob` to run `ping` commands only as `root` without a password:

User/Group
User: `bob`

Run As
User: `root`

No Password
checked

Commands
`/sbin/ping`

Example 2

Allow anyone in the admins group to run all commands as any user, but prompted for a password:

User/Group
Group: `admins`

Run As
User: *ALL Users*

No Password
Unchecked

Commands
`ALL`

Command Examples

These examples demonstrate how to specify commands in various ways.

- Run `ping` with any parameters:

```
/sbin/ping
```

- Run `ping` only to `192.168.1.2`:

```
/sbin/ping 192.168.1.2
```

- Run command `blah` without any parameters:

```
/usr/local/bin/blah ""
```

- Run `ping` and `traceroute` and their IPv6 variants with any parameters:

```
/sbin/ping, /sbin/ping6, /usr/sbin/traceroute, /usr/sbin/traceroute6
```

30.25.3 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.26 Status Traffic Totals

This package displays different ways to view the traffic usage generated by the network traffic monitoring tool vnStat.

See also:

Similar tools are available for firewall bandwidth monitoring: *How can I monitor bandwidth usage*

30.26.1 Notes

The package adds every NIC during installation. If a NIC is later added or removed on the firewall, remove the package and re-install it again. If the firewall has data for a NIC vnStat will report the data even if the NIC has been removed.

A reinstallation of the package will not change this as the firewall has data pertaining to the non-existent data and thus other packages such as vnstat2 will report the data it has or has found.

30.26.2 Known Issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

30.26.3 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

30.27 UDP Broadcast Relay

The UDP Broadcast Relay package listens for UDP broadcast packets and retransmits them on additional interfaces. This enables UDP broadcast discovery protocols to cross separate networks.

Important: Firewall rules are not automatically created for configured instances. If hosts on separate networks need to respond to the relayed broadcast packets, firewall rules may be needed to allow traffic back to the original sender.

30.27.1 General Settings

Enable
: Enable the UDP Broadcast Relay service. If checked, enabled instances will be active.

Track CARP Status
: Tracks the CARP status of the selected CARP VIP. The service will only run when the selected VIP is in the MASTER state.

Configured Instances

ID
The instance ID.

Port
UDP port being relayed.

Interfaces
Interfaces to relay packets on.

Description
The user-specified description.

Actions
Administrative actions available for the instance.

Edit
Edit the instance settings.

Add
Configure a new instance.

Delete
Delete the selected instance(s).

30.27.2 Instance Settings

Enable
Enable this instance. If unchecked, the instance will not be active.

Description
A description for administrative reference (not parsed).

Network Interfaces
Interfaces to receive and transmit packets on; must select at least **two**. When a packet is received, it's sent to all specified interfaces except for the one it originated on. Unassigned and special interfaces (including `lo0` and `enc0`) are omitted.

Note: This service is only compatible with interfaces which have the broadcast flag and an IPv4 address. It cannot use Point-to-point interfaces such as those used in WireGuard.

Spoof Source
Spoof the source IP address and/or port when relaying packets.

Keep Original (default)
The source IP address and UDP port of the broadcast packet will remain unchanged.

Use Interface Address and Destination Port
Sets the source IP address to that of the outgoing interface and the source UDP port to the same as the destination port.

Use Interface Address only
Sets the source IP address to that of the outgoing interface and keeps the original source UDP port.

Instance ID
A unique number between instances (1-63). This is used to set the DSCP of outgoing packets to determine if a packet is an echo and should be discarded.

Destination Port
Destination UDP port to listen on (1-65535).

> **Note:** Instances need to bind to the specified port on any/all interfaces. If another service such as UPnP (port `1900`) or Avahi (port `5353`) is running, the instance will fail to start.

Multicast Address
Multicast group to listen for and relay packets on (optional).

30.27.3 Examples

A list of example use cases are given below.

UDP Port	Spoof Source	Multicast Address	Description
5353	Address/Port	224.0.0.251	mDNS: used for e.g. Chromecast and Bonjour
1900	Original	239.255.255.25	SSDP: used for e.g. DLNA, Sonos, and UPnP
1900	Original	239.255.255.25	Windows Network Neighborhood Discovery: uses SSDP,
138	Original	-	NetBIOS-NS, and NetBIOS-SS
137	Original	-	
21027	Original		Syncthing Discovery
19132	Original		Minecraft

CHAPTER

THIRTYONE

VIRTUALIZATION

pfSense® software can be run in a variety of Type-1 (bare metal/native) and Type-2 (hosted) virtualization environments, such as Proxmox VE, VirtualBox, Xen, KVM, Hyper-V, VMware and so on.

Warning: The best practice is to use Type-1 hypervisors for production. Type-2 hypervisors such as VirtualBox or VMware Workstation work fine for testing, but avoid using them for production roles where possible.

Set up and installation is straightforward and similar to physical hardware using the ISO image.

31.1 VirtIO Drivers

The FreeBSD kernel used by pfSense® software includes VirtIO drivers built into the kernel. No special action is necessary to enable the drivers.

31.1.1 Disable Hardware Checksum Offloading

With the current state of VirtIO network drivers in FreeBSD, it is necessary to disable hardware checksum offload to reach systems (at least other VM guests, possibly others) protected by pfSense software **directly from the VM host**. The firewall attempts to do this automatically when it detects `vtnet` interfaces, but the setting may also be changed manually under **System > Advanced** on the **Networking** tab.

Note: After changing the setting manually *reboot the firewall, even though there is no prompt instructing to do so*.

The issue is most likely caused by FreeBSD Bug 165059.

Hardware checksums and other NIC offloading features like TSO may also need to be disabled on the hypervisor system in addition to the pfSense VM.

31.2 Guides

- *Virtualizing with Proxmox® VE*
- *Virtualizing pfSense Software with Hyper-V*
- *Virtualizing pfSense Software with VMware vSphere / ESXi*

CHAPTER

THIRTYTWO

WIRELESS

32.1 Should pfSense software act as an access point?

Historically, the access point functionality in FreeBSD has suffered from serious compatibility or performance problems with some wireless clients. Over time this has improved significantly. pfSense® software access points are used in various locations with no trouble. It is used with various clients such as MacBook Pro, Apple AirTunes, Mac mini, iPod Touch, Android devices, Palm, various Windows laptops, Xbox 360, and FreeBSD clients, and it works very reliably.

There is the possibility of finding incompatible devices with any access point. FreeBSD is no exception, and it can be more common with FreeBSD than other access points. Using pfSense software as an access point can work quite well with the right card and configuration.

In general, the best practice is *Using External Wireless Access Points*, especially if clients require 802.11ac or newer standards. Placing that burden on an external device and allowing pfSense software to focus on the firewall/routing/NAT/etc is simpler in the long run. An access point may still be connected to a dedicated interface or VLAN for isolation purposes.

32.1.1 Incompatible wireless clients

There are no known incompatible devices at this time.

32.2 Recommended Wireless Hardware

A variety of wireless cards can be used in FreeBSD 15.0-CURRENT@bf06074106cf, and pfSense® software includes drivers for every card supported by FreeBSD. Some have better implementations than others. Most development of wireless features on pfSense software uses Atheros hardware, so they are the most likely to work. Users have reported success with other cards as well, with Ralink being another popular choice.

FreeBSD and pfSense software may work with other cards, but those cards may not offer all available features. In particular, some cards manufactured by Intel can be used in infrastructure mode as clients but cannot run in access point mode due to limitations of the hardware itself.

32.2.1 Wireless cards from big name vendors

Linksys, D-Link, Netgear and other major manufacturers commonly change the chipset used in their wireless cards without changing the model number. There is no way to ensure a specific model card from these vendors will be compatible because there is no reliable way of knowing which "minor" card revision and chip a package contains. While one revision of a particular model may be compatible and work well, another card of the same model may be incompatible. For this reason, the best practice is to avoid cards from major manufacturers. If a card is already on hand, it is worth trying to see if it is compatible. Be wary when purchasing because even if the "same" model worked for someone else, a new purchase may result in a completely different piece of hardware that is incompatible.

32.2.2 Status of 802.11n

pfSense software version 2.8.0-RELEASE is based on FreeBSD 15.0-CURRENT@bf06074106cf which can utilize 802.11n on certain hardware such as those based on the Atheros AR9280 and AR9220 chipsets. Netgate has tested cards using those chipsets and they work well. Some other non-Atheros cards are documented by FreeBSD to work on 802.11n, specifically, `mwl(4)` and `iwn(4)`. These may work using the 802.11n standard but experiences with 802.11n speeds may vary.

The FreeBSD Wiki Article for 802.11n Support contains the most up-to-date information about chipsets and drivers that work with 802.11n.

32.2.3 Status of 802.11ac

Currently, 802.11ac in FreeBSD is still under development, so it is not available in pfSense software. Development on FreeBSD can be tracked by checking the FreeBSD Wiki Article for 802.11ac Support.

32.2.4 Radio Frequencies and Dual Band

Some cards can utilize both 2.4GHz and 5GHz bands, such as the Atheros AR9280, but only one band may be used at a time. Currently, there are no cards working in FreeBSD that will operate in both bands concurrently. Using two separate cards in one unit is not desirable as their radios may interfere. In cases which require dual or multiple bands, the best practice is to use an external AP.

32.2.5 Wireless drivers included in pfSense software

This section lists the wireless drivers included in pfSense software and their compatible chipsets. This information was derived from the FreeBSD man pages for the drivers in question. Drivers in FreeBSD are referred to by their driver name, followed by `(4)`, such as `ath(4)`. The `(4)` refers to the kernel interfaces section of the man page collection, in this case specifying a network driver. The drivers are listed in order of frequency of use based on reports from users.

Access Point (hostap) Mode Cards

The cards in this section are capable of acting as an access point to accept connections from other wireless clients. This is referred to as *hostap* mode.

ath(4)

The `ath(4)` driver is for cards based on the Atheros AR5210, AR5211, AR5212, AR5416, and AR92xx APIs which are used by many other Atheros chips of varying model numbers. Most Atheros cards can offer up to four virtual access points (VAPs), stations, or a combination to create a wireless repeater.

Though not explicitly listed in the man page, the FreeBSD Wiki Article for 802.11n Support also states that the driver is compatible with AR9130, AR9160, AR9280, AR9285, AR9287, and potentially other related chipsets.

ral(4) / ural(4) / run(4) / rum(4)

There are several related Ralink Technology IEEE 802.11 wireless network drivers, each for a different set and type of card.

ral(4)
: Cards based on the Ralink Technology RT2500, RT2501 and RT2600, RT2700, RT2800, RT2900, RT3090, and RT3900E chipsets.

ural(4)
: RT2500USB.

run(4)
: RT2700U, RT2800U, RT3000U, RT3900E, and similar.

rum(4)
: RT2501USB and RT2601USB and similar.

Of these, only certain chips compatible with `run(4)` can utilize VAPs.

The RT3090 `ral(4)` chip is the only model listed as capable of 802.11n on FreeBSD. The RT2700 and RT2800 `ral(4)` and the RT3900E `run(4)` hardware are capable of 802.11n, but the drivers on FreeBSD do not currently implement 802.11n features.

mwl(4)

The Marvell IEEE 802.11 wireless network driver, `mwl(4)`, is compatible with cards based on the 88W8363 chipset and fully implements 802.11n. This card offers multiple VAPs and stations, up to eight of each.

Client (station) Mode Only Cards

The cards in this section are not capable of acting as access points, but may be used as clients in station mode, for example as a wireless WAN.

uath(4)

Atheros USB 2.0 wireless devices using AR5005UG and AR5005UX chipsets are compatible with the `uath(4)` driver.

ipw(4) / iwi(4) / iwn(4) / wpi(4)

Intel wireless network drivers cover various models with different drivers.

ipw(4)
Intel PRO/Wireless 2100 MiniPCI adapters.

iwi(4)
Intel PRO/Wireless 2200BG/2915ABG MiniPCI and 2225BG PCI adapters.

iwn(4)
Intel Wireless Wi-Fi Link 4965, 1000, 5000 and 6000 series PCI Express adapters.

wpi(4)
Intel 3945ABG adapters.

Cards compatible with the `iwn(4)` driver are documented by FreeBSD as capable of using 802.11n in client mode.

Several Intel adapters have a license restriction with a warning that appears in the boot log. The `ipw(4)`, `iwi(4)`, and `wpi(4)` drivers have license files that must be read and agreed to. These licenses are located on the firewall in `/usr/share/doc/legal/intel_ipw/LICENSE`, `/usr/share/doc/legal/intel_iwi/LICENSE`, and `usr/share/doc/legal/intel_wpi/LICENSE` respectively. To agree to the license, and *Loader Tunables* indicate the license acknowledgment, such as:

```
legal.intel_ipw.license_ack=1
```

Given the limited use of these adapters as clients only, development of a GUI-based solution to acknowledge these licenses is unlikely.

bwi(4) / bwn(4)

The Broadcom BCM43xx IEEE 802.11b/g wireless driver is split in two depending on the specific models in use.

bwi(4)
BCM4301, BCM4303, BCM4306, BCM4309, BCM4311, BCM4318, BCM4319 using an older v3 version of the Broadcom firmware.

bwn(4)
BCM4309, BCM4311, BCM4312, BCM4318, BCM4319 using a newer v4 version of the Broadcom firmware.

Driver compatibility overlaps for some cards. The `bwn(4)` driver is preferred for compatible cards while the `bwi(4)` driver must be used on the older cards not covered by `bwn(4)`.

malo(4)

Marvell Libertas IEEE 802.11b/g wireless driver, `malo(4)` is compatible with cards using the 88W8335 chipset.

upgt(4)

The Conexant/Intersil PrismGT SoftMAC USB IEEE 802.11b/g wireless driver, `upgt(4)`, compatible with cards using the GW3887 chipset.

urtw(4) / urtwn(4) / rsu(4)

The trio of related Realtek wireless drivers cover several different models:

urtw(4)
: RTL8187B/L USB IEEE 802.11b/g models with a RTL8225 radio

urtwn(4)
: RTL8188CU/RTL8188EU/RTL8192CU 802.11b/g/n

rsu(4)
: RTL8188SU/RTL8192SU 802.11b/g/n

As in other similar cases, though the chips compatible with `urtwn(4)` and `rsu(4)` are capable of 802.11n, FreeBSD does not implement their 802.11n features.

zyd(4)

The ZyDAS ZD1211/ZD1211B USB IEEE 802.11b/g wireless network device driver, `zyd(4)`, compatible with adapters using the ZD1211 and ZD1211B USB chips.

32.3 Working with Virtual Access Point Wireless Interfaces

pfSense® software can utilize virtual wireless interfaces using Multi-BSS. These are known as Virtual Access Point or VAP interfaces, even if they are being used for client mode. VAPs allow multiple access points or clients to be run on the same wireless card, or to use a combination of access point and client mode. The most common use case is for multiple access points with different SSIDs each with unique security requirements. For example, one with no encryption but with captive portal and strict access rules and a separate network with encryption, authentication, and less strict access rules.

Warning: Even if a card is not compatible with multiple VAP instances, the first entry must be created manually before it can be assigned.

VAP compatibility varies by card and driver, see *Recommended Wireless Hardware* for details. Odds are, however, if an Atheros wireless card is in use, it will work. While there is no theoretical limit to the number of VAPs a card may use, driver and hardware implementations vary, so the practical limit is four VAPs on *ath(4)* and eight on *mwl(4)*.

All VAPs on a given card share some common settings, such as the channel, regulatory settings, antenna settings, and wireless standard. Other settings such as the mode, SSID, encryption settings and so on may vary between VAPs.

32.3.1 Creating and Managing Wireless Instances

To create a new wireless instance:

- Navigate to **Interfaces > Assignments** on the **Wireless** tab.
- Click  **Add** to create a new entry
- Select the **Parent Interface**, for example *ath0*
- Pick the **Mode** from one of *Access Point*, *Infrastructure* (BSS, client mode), or *Ad-hoc* (IBSS)
- Enter a **Description**
- Click **Save**

An example is shown in Figure *Adding a Wireless Instance*.

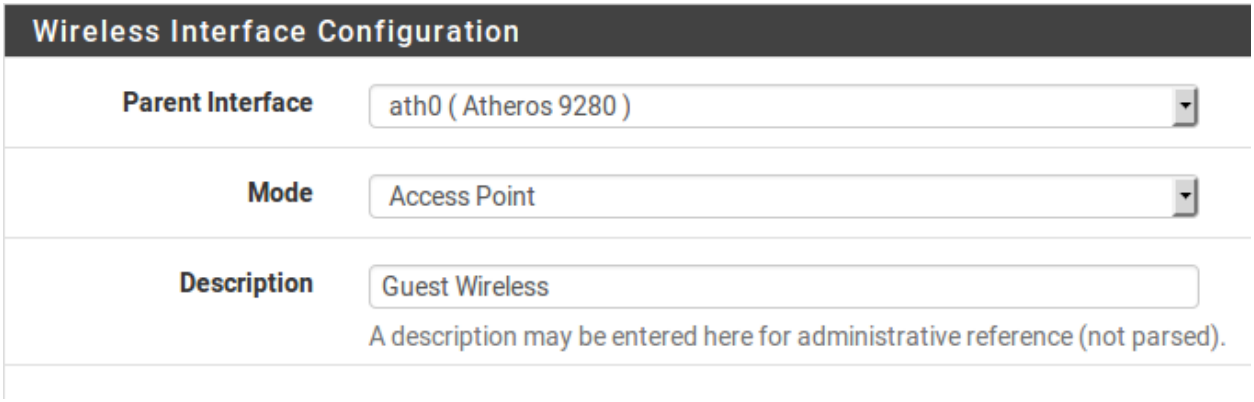


Fig. 1: Adding a Wireless Instance

Once the entry has been saved it is then available for assignment under **Interfaces > Assignments**. From there, assign and then edit the settings like any other wireless interface.

Note: The assigned interface must be configured to use the same mode specified when the VAP was created.

32.4 pfSense Software as a Wireless Access Point

With a wireless card that implements and allows hostap mode (See *Access Point (hostap) Mode Cards*), pfSense® software can be configured as a wireless access point.

32.4.1 Should an external AP or pfSense software be used for an access point?

The access point functionality in FreeBSD, and thus pfSense, has improved dramatically over the years and is considered stable currently for most uses. That said, many use cases behave better with an external access point, especially deployments that have requirements such as 802.11ac, concurrent operation in 2.4GHz and 5GHz, wireless mesh networks, or rare cases with clients that will not associate with an access point run using pfSense software.

Access points on pfSense software have been used with success in small-to-medium deployments, with gear such as a MacBook Pro, Apple AirTunes, iPod Touch, iPad, Android phones and tablets, various Windows laptops, Xbox, and FreeBSD clients, and it works very reliably across all these devices. There is the possibility of finding incompatible devices with any access point, and FreeBSD is no exception.

The main deciding factor these days is whether it is possible to use 802.11n, 802.11ac, or newer wireless standards. Hardware compatible with these standards in pfSense software is somewhat limited, and recent protocols are not implemented. This is a dealbreaker for some, and as such using an external access point would be best for networks requiring 802.11ac and in some cases 802.11n if suitable hardware cannot be obtained.

The next most common factor is location of the antennas or the wireless access point in general. Often, the firewall running pfSense software is located in an area of the building that is not optimal for wireless, such as a server room in a rack. For ideal coverage, the best practice is to locate the AP in an area that is less susceptible to wireless interference and that would have better signal strength to the area where wireless clients reside. If the firewall running pfSense software is located alone on a shelf in a common area or other similar area conducive to good wireless signal, this may not be a concern.

32.4.2 Configuring pfSense software as an access point

The process of configuring pfSense software to act as a wireless access point (AP) is relatively easy. Many of the options will be familiar to anyone who has configured other wireless routers before, and some options may be new unless commercial-grade wireless equipment has been used. There are dozens of ways to configure access points, and they all depend upon the environment in which it will be deployed. In this example pfSense software is configured as a basic AP that uses WPA2 encryption with AES. In this example, ExampleCo needs wireless access for laptops in the conference room.

Preparing the Wireless Interface

Before starting, ensure that the wireless card is installed in the firewall and the pigtails and antennas are firmly attached.

Create the wireless instance as described in *Creating and Managing Wireless Instances* if it does not already exist. When working as an access point, it must use *Access Point* mode. The wireless card must be assigned as an OPT interface and enabled before the remaining configuration can be completed.

Interface Description
: When in use as an access point, naming the interface *WLAN* (Wireless LAN) or *Wireless*, or naming it after the SSID makes it easier to identify. If pfSense software will be driving multiple access points, there should be some way to distinguish them, such as "WLANadmin" and "WLANsales". In this example, it is named `ConfRoom`.

Interface Type
: Since this example will be an AP on a dedicated IP subnet, the IPv4 Configuration Type must be set to *Static IPv4*

IP Address
: An IPv4 Address and subnet mask must be specified. This is a separate subnet from the other interfaces. For this example it can be `192.168.201.0/24`, a subnet that is otherwise unused in the ExampleCo network. Using that subnet, the IPv4 Address for this interface will be `192.168.201.1`.

Common Wireless Settings

These settings are shared for all VAPs on a given physical wireless card. Changing these settings on one interface will change them on all other virtual interfaces using the same physical adapter.

Persist common settings
By checking **Persist common settings**, the configuration values in this section will be preserved even if all the interfaces and VAPs are deleted or reassigned, when they would otherwise be lost.

Wireless Standard
Depending upon hardware, there are several choices available for the wireless **Standard** setting, including *802.11b*, *802.11g*, *802.11g turbo*, *802.11a*, *802.11a turbo*, *802.11ng*, *802.11na*, and possibly others. For this example, choose *802.11ng* for an 802.11n access point operating in the 2.4GHz band.

802.11g OFDM Protection Mode
The **802.11g OFDM Protection Mode** setting is only useful in mixed standard environments where 802.11g and 802.11b have to interact. Its primary use is for avoiding collisions. Given the age of 802.11b and scarcity of working devices that use it, the setting is best left at *Protection mode off*. There is a performance penalty for using it, since it has some overhead on each frame and also requires extra steps when transmitting frames.

Wireless Channel Selection
When selecting a **Channel**, knowledge of nearby radio transmitters in similar frequency bands is required to avoid interference. In addition to wireless access points, there are also cordless phones, Bluetooth, baby monitors, video transmitters, microwaves, and many other devices that utilize the same 2.4 GHz spectrum that can cause interference.

Often any channel will work so long as the AP clients are near the antenna. With 802.11g and before, the safest channels to use were *1*, *6*, and *11* since their frequency bands did not overlap each other. This is no longer true with 802.11n and later or even some 802.11g setups which use wider ranges of frequencies to attain higher speeds. For this network, since there are no others around, channel *1* is a fine choice.

Note: Always pick a specific channel. Do not select *Auto* for the channel of an Access Point. The input validation on current versions of pfSense software prevents this from being selected.

When using other standards, or using wireless in countries other than the US, there may be many more channels available than described here. Cards compatible with 802.11a or 802.11n may also be capable of using channels in the 5 GHz spectrum.

The full list of channels available to the card is shown in the **Channel** drop-down and must agree with the chosen *Standard*. For example, do not choose *802.11ng* for the **Standard** and then pick a **Channel** used only for *802.11na*. The channel list also includes some information about the standard, frequency of the channel, and the maximum transmit power both of the card and in the regulatory domain for that particular channel. Be careful to watch the power when selecting a channel, because some channels, especially in the 5GHz band, vary widely in their allowed power levels.

See also:

Survey tools such as NetSurveyor, InSSIDer, Wi-Spy, and countless other apps for various operating systems, phones, tablets, and so on may help to choose a less busy channel or area of the spectrum. Mileage may vary.

Distance setting
Measured in meters, and only implemented by Atheros cards, The **Distance Setting** field tunes ACK/CTS timers to fit the distance between AP and Client. In most cases it is not necessary to configure this value, but it may help in certain tricky wireless setups such as long-range clients.

Regulatory settings

The **Regulatory settings** section controls how the card is allowed to transmit legally in a specified region. Different countries typically have different regulatory settings, and some countries have none. If unsure, check with the local government to see which laws apply in a given area. The default values are usually OK, as the cards may be set to a specific region already. In some cases **Regulatory settings** must be set manually if the card has a default not understood by the driver. Similar to the previous section, these values are applied to the card itself and cannot vary between VAPs on the card.

While it may be tempting to set the card to *Debug* in order to use settings not otherwise allowed, this action could result in legal trouble should it be noticed. The likelihood of this happening varies greatly by country/area so use that with caution.

Regulatory domain
: The **Regulatory Domain** is the governmental body that controls wireless communications in a region. For example, the US and Canada follow FCC regulations while in the UK it's ETSI. If unsure of the regulatory domain in a region, see the **Country** setting.

Country
: Sometimes specific countries inside a regulatory domain have different restrictions. The **Country** option contains a drop-down list of many countries throughout the world and their associated country codes and regulatory domains.

Location
: Certain restrictions exist for *Indoor* and *Outdoor* transmissions as well. Setting the **Location** of the transmitter will further adjust the allowed transmission power and/or channels.

Network-specific wireless configuration

These settings are unique per interface, even on virtual wireless interfaces. Changing these settings does not affect any other interfaces.

Wireless Mode
: Set the **Mode** field to *Access Point*, and pfSense software will use `hostapd` to act as an AP.

Service Set Identifier (SSID)
: The **SSID** is the "name" of the AP as seen by clients. Set the SSID to something readily identifiable yet unique. Keeping with the example, `ConfRoom` is a good name to use.

Minimum wireless standard
: The **Minimum wireless standard** drop-down controls whether older clients are able to associate with this access point. Allowing older clients may be necessary in some environments if devices are still around that require it. Some devices are only compatible with 802.11g and require a mixed network g/n in order to work. The flip side of this is that slower speeds may be seen as a result of allowing such devices on the network as the access point will be forced to cater to the lowest common denominator when an 802.11g device is transmitting at the same time as an 802.11n device. In our example conference room, users will only be using recently purchased company-owned laptops that are all capable of 802.11n, so *802.11n* is the best choice.

Intra-BSS Communication
: If **Allow intra-BSS communication** is checked, wireless clients will be able to see each other directly. If clients will only need access to the Internet, it is typically safer to uncheck this option. In this scenario, users in the conference room may need to share files back and forth directly between laptops, so this will stay checked.

Enable WME
: Wireless Multimedia Extensions, or **WME**, is a part of the wireless standard that provides some Quality of Service for wireless traffic to ensure proper delivery of multimedia content. It is required

for 802.11n to operate, but is optional for older standards. This feature is not implemented by all cards/drivers.

Hide SSID
Normally the AP will broadcast its SSID so that clients can locate and associate with it easily. This is considered by some to be a security risk, announcing to all who are listening that a wireless network is available, but in most cases the convenience outweighs the (negligible) security risk. The benefits of disabling SSID broadcasting are overblown by some, as it does not actually hide the network from anyone capable of using many freely available wireless security tools that easily find such wireless networks. For the conference room AP, this example leaves the option unchecked to make it easier for meeting attendees to find and use the service.

Wireless Encryption (WPA)

Two types of encryption are available for 802.11 networks: WPA, and WPA2. WPA2 with AES is the most secure. Even when not worrying about encrypting the over-the-air traffic (which should be done), it provides an additional means of access control. All modern wireless cards and drivers implement WPA2.

Warning: Wireless Encryption Weaknesses

WEP has serious known security problems for years, and WEP capability has been removed from pfSense software. It is possible to crack WEP in a matter of minutes at most, and it should never be relied upon for security. If WEP is required, an external AP must be used.

TKIP (Temporal Key Integrity Protocol), part of AES, became a replacement for WEP after it was broken. It uses the same underlying mechanism as WEP, and hence is vulnerable to some similar attacks. These attacks have become more practical and TKIP is no longer considered secure. TKIP should never be used unless devices are present that are incompatible with WPA or WPA2 using AES. WPA and WPA2 in combination with AES are not subject to these flaws in TKIP.

In this example, the ConfRoom wireless must be secured with WPA2.

Enable
This checkbox enables *WPA* or WPA2 encryption, so it should be **checked**

WPA Pre-Shared Key
Enter the desired wireless key, in this example `excoconf213`.

WPA Mode
WPA or WPA2, in this example, *WPA2*

WPA Key Management Mode
Can be *Pre-Shared Key* (PSK) or *Extensible Authentication Protocol* (EAP). In this example, PSK is sufficient.

WPA Pairwise
This should almost always be set to *AES*, due to the weaknesses in *TKIP* mentioned previously.

Group Key Rotation
This option allows setting how often the broadcast/multicast encryption keys (Group Transient Key, GTK) are rotated, in seconds. It can be any value from `1` to `9999`, but it should be shorter than the **Group Master Key Regeneration** value. The default value of `60` seconds (one minute) is adequate. Lower values may be more secure but may bog things down with frequent rekeying.

Group Master Key Regeneration
This parameter controls how often, in seconds, the master key (Group Master Key, GMK) used in-

ternally to generate GTKs is regenerated. It can be any value from `1` to `9999`, but it should be longer than the **Group Key Rotation** value. The default value of `3600` seconds (one hour) is adequate.

Strict Key Regeneration
This option causes the firewall to change the GTK whenever a client leaves the access point, much like changing the passwords when an employee leaves. There may be a slight performance penalty in cases where there is a high turnover of clients. In cases where security is not a primary concern, this can be left disabled.

IEEE 802.1X Authentication (WPA Enterprise)

Another available type of wireless security is known as IEEE 802.1X Authentication, or more commonly referred to as *WPA Enterprise* or *WPA2 Enterprise*. This mode allows using a more traditional username and password entry in order to gain access to the wireless network. The downside is that this authentication must be done via RADIUS servers. If an existing RADIUS server is already present or easily deployed, it may be a viable source of wireless access control. In this example, 802.1X is not used, but the options are explained.

See also:

The FreeRADIUS package (*FreeRADIUS package*) can fulfill this purpose.

Note: Some older operating systems may not properly handle 802.1X or may have long delays after failed authentication attempts, but there are typically workarounds for those issues via OS updates or patches.

Clients must also be configured to properly access the service. Some may pick up the proper settings automatically, others may need to be configured for a specific mode (e.g. *PEAP*) or may need certificates loaded. The specific values depend on the RADIUS server settings.

To get started with 802.1X authentication, first set **WPA Key Management** to *Extensible Authentication Protocol*.

Enable 802.1X Authentication
When checked, 802.1X authentication is enabled and required of clients.

Primary 802.1X Server
The preferred server for 802.1X authentication.

IP Address
The IP address of the preferred RADIUS server to use for 802.1X client authentication.

Port
The port upon which to contact the RADIUS server for authentication requests, typically `1812`.

Shared Secret
The password to use when communicating with the RADIUS server from this firewall. This must match the shared secret defined for this firewall on the RADIUS server.

Secondary 802.1X Server
The same parameters as above, but for a secondary RADIUS server in case the first one is unreachable.

Authentication Roaming Preauth
This option sets up pre-authentication to speed up roaming between access points. This will perform part of the authentication process before the client fully associates to ease the transition.

Finishing AP Settings

The previous settings are enough to get a wireless access point running with 802.11n with WPA2 + AES encryption. When the settings are complete, click **Save**, then **Apply Changes**.

Configuring DHCP

Now that an entirely separate network has been created, DHCP must be enabled to automatically provide associating wireless clients an IP address. Browse to **Services > DHCP Server**, click on the tab for the wireless interface (**ConfRoom** for this example). Check the box to **Enable**, set whatever size range will be needed, and any additional options desired, then click **Save** and **Apply Changes**. For more details on configuring the DHCP service, see *DHCP*.

Adding Firewall Rules

Since this wireless interface is an OPT interface, it will have no default firewall rules. At the very least a rule must be added to allow traffic from this subnet to any destination. Since the conference room users will need internet access and access to other network resources, a rule to allow all traffic by default will be fine in this case. To create the rule:

- Navigate to **Firewall > Rules**
- Click on the tab for the wireless interface (**ConfRoom** for this example).
- Click  **Add** and configure a rule as follows:

  Interface
  ConfRoom

  Protocol
  Any

  Source
  ConfRoom Net

  Destination
  Any

- Click **Save**
- Click **Apply Changes**

See also:

For more information about creating firewall rules, see *Firewall*.

Associating Clients

The newly configured pfSense software AP should appear in the list of available access points from a wireless device, assuming broadcasting of the SSID was not disabled. A client should now be able to associate with it as it would with any other access point. The exact procedure will vary between operating systems, devices, and drivers, but most manufacturers have streamlined the process to make it simple for everyone.

Viewing Wireless Client Status

When a wireless interface is configured for access point mode, the associated clients will be listed on **Status > Wireless**.

Interesting sysctl tunables which cannot be controlled from GUI

dev.ath.0.tpc
`0` = disable 1 = enable

Switch on or off Transmission Power Control. Can be tricky in point to multipoint applications.

dev.ath.0.tpscale
`0,1,2,3,4`

Size of the increment that TPC will use to up/down the power, normally 1 is the best choice. A higher scale value will most likely make the link drop if the signal is close to what it needs to be and the TPC is throttled down.

dev.ath.0.tpack
`0 -> 99`

Controls the ACK power separately. Normally it is the same as `tpcts`

dev.ath.0.tpcts
`0 -> 99`

Controls the CTS power separately. Normally it is the same as `tpack`

Tuning ACK timers manually:

```
Real life values:

range                ack-timeout

5GHz 5GHz-turbo      2.4GHz-G
0km  default default         default
5km  52      30              62
10km 85      48              96
15km 121     67              133
20km 160     89              174
25km 203     111             219
30km 249     137             268
35km 298     168             320
40km 350     190             375
45km 405     -               -
```

32.5 Wireless WAN

A wireless card in a firewall running pfSense® software can be used as the primary WAN interface or an additional WAN in a multi-WAN deployment.

32.5.1 Interface assignment

If the wireless interface has not yet been assigned, there are two possible choices: Add it as an additional OPT interface or reassign it as WAN.

Before starting, create the wireless instance as described in *Creating and Managing Wireless Instances* if it does not already exist. When working as a WAN, it must use *Infrastructure* mode (BSS).

To add the interface as a new OPT interface:

- Browse to **Interfaces > Assignments**
- Select the wireless interface from the **Available network ports** drop-down below the other interfaces
- Click **Add** to add the interface as an OPT interface

To reassign the wireless interface as WAN:

- Browse to **Interfaces > Assignments**
- Select the wireless interface as **WAN**
- Click Save

Figure *Wireless WAN Interface Assignment* shows an Atheros card assigned as WAN.

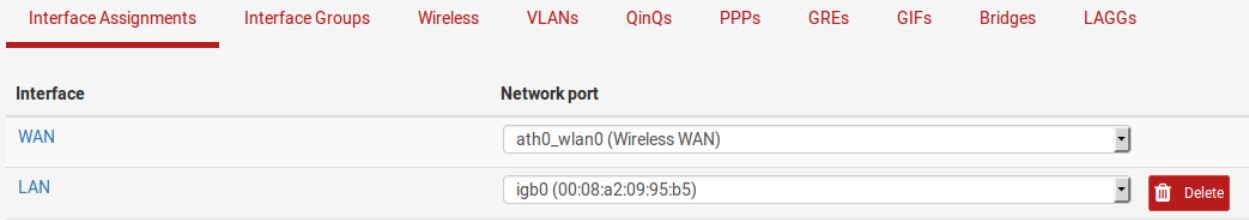


Fig. 2: Wireless WAN Interface Assignment

32.5.2 Configuring the wireless network

Most wireless WANs need only a handful of options set, but specifics vary depending on the Access Point (AP) to which this client interface will connect.

- Browse to the **Interfaces** menu for the wireless WAN interface, for example **Interfaces > WAN**
- Select the type of configuration (*DHCP*, *Static IP*, etc.)
- Scroll down to **Common Wireless Configuration**
- Set the **Standard** to match the AP, for example *802.11g*
- Select the appropriate **Channel** to match the AP
- Scroll down to **Network-specific Wireless Configuration**
- Set the **Mode** to *Infrastructure (BSS)* mode
- Enter the **SSID** for the AP
- Configure encryption such as WPA2 (Wi-Fi Protected Access) if in use by the AP
- Review the remaining settings if necessary and select any other appropriate options to match the AP
- Click **Save**

- Click **Apply Changes**

32.5.3 Checking wireless status

Browse to **Status > Interfaces** to see the status of the wireless interface. If the interface has successfully associated with the AP it will be indicated on the status page. A **status** of **associated** means the interface has connected to the AP successfully, as shown in Figure *Associated Wireless WAN Interface*

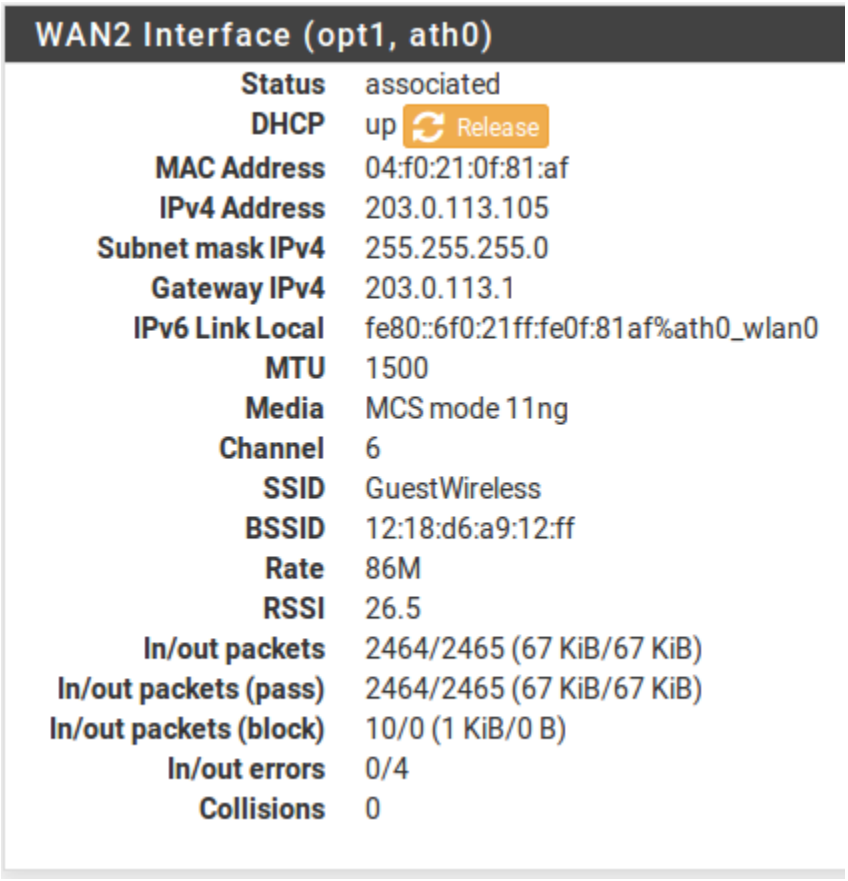


Fig. 3: Associated Wireless WAN Interface

If the interface **status** shows **No carrier**, it was unable to associate. Figure *No carrier on wireless WAN* shows an example of this, where the antenna was disconnected so it could not connect to a wireless network that was some distance away.

32.5.4 Showing available wireless networks and signal strength

The wireless access points visible by the firewall may be viewed by navigating to **Status > Wireless** as shown in Figure *Wireless Status*.

A wireless interface must be configured before this menu item will appear.

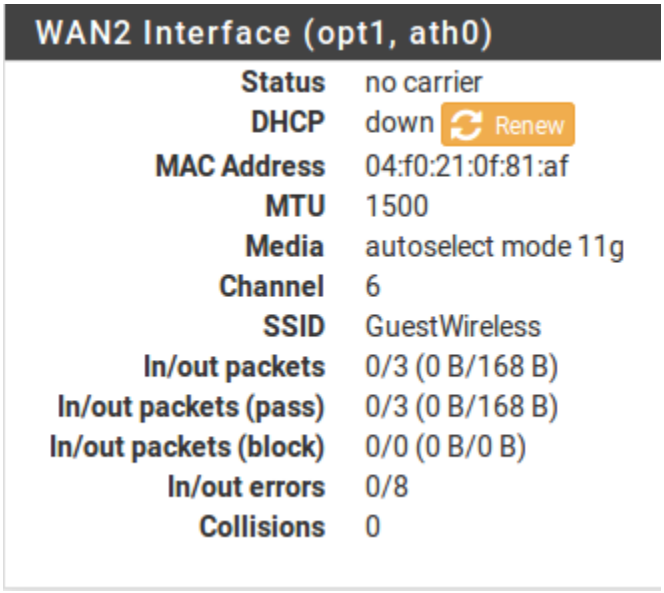


Fig. 4: No carrier on wireless WAN

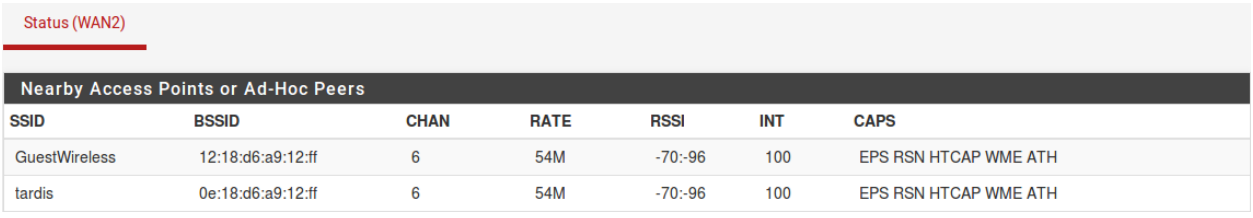


Fig. 5: Wireless Status

32.6 Bridging and wireless

Bridging two interfaces together places them on the same broadcast domain as if they were connected to the same switch. Typically, this is done so that two interfaces will act as though they are on the same flat network using the same IP subnet, in this case a wireless interface and wired interface. When two interfaces are bridged, broadcast and multicast traffic is forwarded to all bridge members.

Certain applications and devices rely on broadcast traffic to function. For example, Apple's AirTunes will not function across two broadcast domains. So if AirTunes is present on the wireless network, and it must be accessed from a system on the wired network, the wired and wireless networks must be bridged. Other examples include media services provided by devices such as Chromecast, TiVo, Xbox, and Playstation. These rely on multicast or broadcast traffic that can only function if the wired and wireless networks are bridged.

32.6.1 Wireless Access Points and Bridging

Only wireless interfaces in access point (hostap) mode will function in a bridged configuration. A wireless interface configured for hostap can be bridged to another interface which combines them on the same broadcast domain. This may be desirable for certain devices or applications that must reside on the same broadcast domain to function properly, as mentioned previously.

32.6.2 BSS and IBSS wireless and Bridging

Due to the way wireless works in BSS mode (Basic Service Set, client mode) and IBSS mode (Independent Basic Service Set, Ad-hoc mode), and the way bridging works, a wireless interface cannot be bridged in BSS or IBSS mode. Every device connected to a wireless card in BSS or IBSS mode must present the same MAC address. With bridging, the MAC address passed is the actual MAC of the connected device. This is normally a desirable facet of how bridging works. With wireless, the only way this can function is if all the devices behind that wireless card present the same MAC address on the wireless network. This is explained in depth by noted wireless expert Jim Thompson in a mailing list post.

As one example, when VMware Player, Workstation, or Server is configured to bridge to a wireless interface, it automatically translates the MAC address to that of the wireless card. Because there is no way to translate a MAC address in FreeBSD, and because of the way bridging in FreeBSD works, it is difficult to provide any workarounds similar to what VMware offers.

32.6.3 Choosing Routing or Bridging

The choice between bridging (using the same IP subnet as the existing LAN) or routing (using a dedicated IP subnet for wireless) for wireless clients will depend on what services wireless clients require. In many home network environments there are applications or devices that require wired and wireless networks to be bridged. In most corporate networks, there are few if any applications that require bridging. Which to choose depends on the requirements of network applications in use, as well as personal preference.

There are some compromises to this, one example being the Avahi package. It can listen on two different broadcast domains and rebroadcast messages from one to the other in order to allow multicast DNS to work (aka Rendezvous or Bonjour) for network discovery and services. If the required services all use protocols that can be handled by Avahi, then using a routed method may be possible.

For services running on the firewall, bridging can also be problematic. Features such as limiters, Captive Portal, and transparent proxies require special configuration and handling to work on bridged networks. Specifically, the bridge itself must be assigned and the only interface on the bridge with an IP address must be the assigned bridge. Also, in order for these functions to work, the IP address on the bridge must be the address used by clients as their gateway.

32.7 Additional protection for a wireless network

In addition to strong encryption from WPA2 with AES, some users like to employ an additional layer of encryption and authentication before allowing access to network resources. The two most commonly deployed solutions are *Captive Portal* and *Virtual Private Networks*. These methods can be employed whether an external access point is used on an OPT interface or an internal wireless card as the access point.

Note: In theory, The PPPoE server could also be used in this role, but implementation would be impossible on some clients and non-trivial on most others, so it is not typically a viable option when combined with wireless.

32.7.1 Additional wireless protection with Captive Portal

By enabling Captive Portal on the interface where the wireless resides, authentication can be required before users may access network resources beyond the firewall. In corporate networks, this is commonly deployed with RADIUS authentication to Microsoft Active Directory, so users can use their Active Directory credentials to authenticate while on the wireless network. Captive Portal configuration is covered in *Captive Portal*.

Note: If the sole requirement is per-user RADIUS authentication, a better solution for RADIUS is 802.1X rather than using Captive Portal, unless there are clients present which are not compatible with 802.1X.

Captive Portal is more likely to be used on open or limited access wireless networks, such as those in a hotel, restaurant, or similar setting where there is either no encryption enabled or a common knowledge/shared key.

32.7.2 Additional protection with VPN

Adding Captive Portal provides another layer of authentication, but does not offer any additional protection from eavesdropping of wireless traffic. Requiring a VPN connection before allowing access to the internal network and Internet adds another layer of authentication as well as an additional layer of encryption for wireless traffic. The configuration for the chosen type of VPN will be no different from a remote access configuration, but the firewall rules must be configured on the pfSense® interface to only allow VPN traffic from wireless clients.

Configuring firewall rules for IPsec

Figure *Rules to allow only IPsec from wireless* shows the minimal rules required to allow only access to IPsec on the WLAN interface IP address. Pings to the WLAN interface IP address are also allowed to assist in troubleshooting.

Configuring firewall rules for OpenVPN

Figure *Rules to allow only OpenVPN from wireless* shows the minimal rules required to allow access only to OpenVPN on the WLAN interface IP address. Pings to the WLAN interface IP address are also allowed to assist in troubleshooting. This assumes the default UDP port 1194 is in use. If another protocol or port was chosen, adjust the rule accordingly.

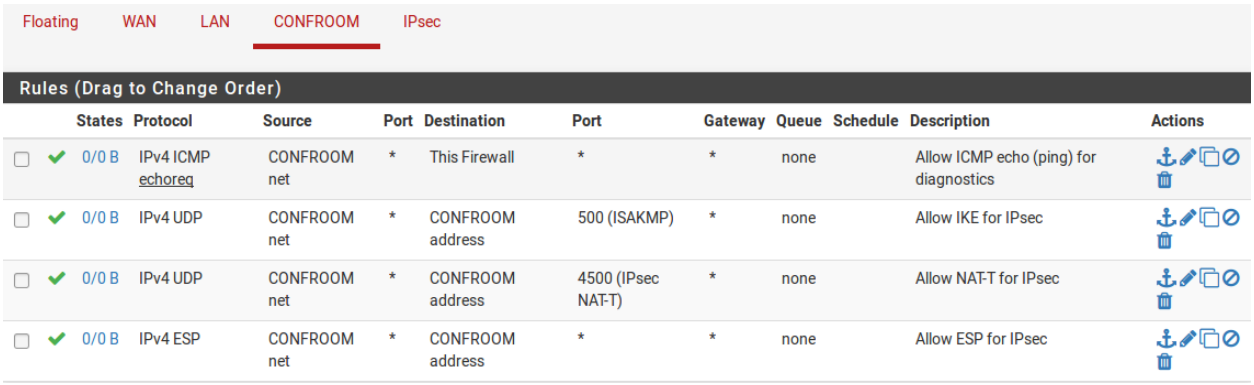


Fig. 6: Rules to allow only IPsec from wireless

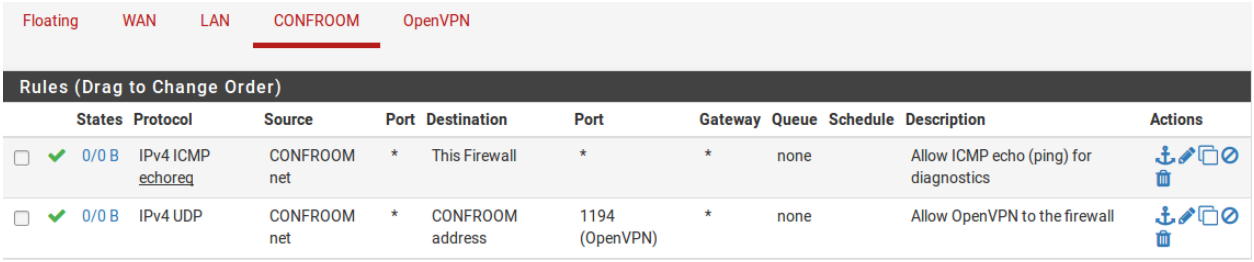


Fig. 7: Rules to allow only OpenVPN from wireless

32.8 Configuring a Secure Wireless Hotspot

A company or organization may wish to provide Internet access for customers or guests using an existing Internet connection. This can be a boon to the customers and business, but can also expose the existing private network to attack if not done properly. This section covers the common means of providing Internet access to guests and customers, while protecting the internal network.

32.8.1 Multiple firewall approach

For the best protection between a private network and a public network, obtain at least two public IP addresses from the ISP and use a second firewall for the public network. To accommodate this, place a switch between the Internet connection and the WAN interface of both firewalls.

This has the added benefit of putting the public network on a different public IP address from the private network, so if a report of abuse is received, it is easy to tell the origin. The firewall protecting the private network will see the public network no differently than any Internet host and vice versa.

32.8.2 Single firewall approach

In environments where the multiple firewall approach is cost prohibitive or otherwise undesirable, the internal network can still be protected by connecting the public network to an OPT interface on a firewall running pfSense® software. Assign a dedicated private IP subnet to this OPT interface, and configure its firewall rules to allow access to the Internet but not the internal network.

In *Rules to allow only Internet access from wireless* the firewall rules allow clients to reach the firewall for DNS requests and ICMP echo request (ping), but prevent all access to other private networks. The RFC1918 alias referenced in the figure includes the RFC1918 private network list, `192.168.0.0/16`, `172.16.0.0/12`, and `10.0.0.0/8`.

Floating | WAN | LAN | EXTERNALAP | IPsec | OpenVPN

Rules (Drag to Change Order)

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
☐ ✔	0/0 B	IPv4 ICMP echoreq	EXTERNALAP net	*	EXTERNALAP address	*	*	none		Allow ping to the firewall	
☐ ✔	0/0 B	IPv4 TCP/UDP	EXTERNALAP net	*	*	53 (DNS)	*	none		Allow DNS -- will be trapped by NAT rule	
☐ ✋	0/0 B	IPv4 *	*	*	RFC1918	*	*	none		Deny access to any other private networks	
☐ ✔	0/0 B	IPv4 *	EXTERNALAP net	*	*	*	*	none		Allow Internet access	

Fig. 8: Rules to allow only Internet access from wireless

32.8.3 Access control and egress filtering considerations

Other than not allowing traffic from the publicly accessible network to the private network, there are additional things to consider in the configuration of a hotspot.

Restrict network access

While many hotspots use open wireless networks with no other authentication, consider additional protections to prevent network abuse. On wireless, consider using WPA or WPA2 and providing the passphrase to guests or customers. Some businesses taking this approach display the passphrase on a placard in the lobby or waiting area, posted in a guest room, or provide it upon request. Also consider implementing Captive Portal on pfSense software (covered in *Captive Portal*). This helps prevent people in other offices and outside the building from using the wireless network even if it is open access.

Disable Intra-BSS communication

If the access point allows, disable intra-BSS communication. This option is also sometimes called "AP Client Isolation". This prevents wireless clients from communicating with other wireless clients directly, which protects users from intentional attacks from other wireless users as well as unintentional ones such as worms.

Intra-BSS communication may be required for certain functions such as wireless printers, Chromecast devices or similar cases when two wireless devices must talk directly to each other, but this is rarely if ever required in the context of a public hotspot.

Egress filtering

Consider what kind of egress policy to configure. The most basic, allowing access to the Internet without allowing access to the private network, is probably the most commonly deployed but consider additional restrictions.

To avoid having the public IP address of the firewall black listed because of infected visiting systems acting as spambots, consider blocking SMTP. Several large ISPs already block SMTP outbound because clients have moved to using authenticated access on the submission port (587) rather than using port 25 directly. An alternative that still lets people use their SMTP e-mail but limits the effect of spambots is to create an allow rule for SMTP and specify **Maximum state entries per host** under **Advanced Options** when editing the firewall rule. Ensure the rule is above any other rules that would match SMTP traffic, and specify a low limit. Because connections may not always be properly closed by the mail client or server, do not set this too low to prevent blocking legitimate users, but a limit of five connections should be reasonable. **Maximum state entries per host** can be set on all firewall rules, but keep in mind that some protocols will require dozens or hundreds of connections to function. HTTP and HTTPS may require numerous connections to load a single web page depending on the content of the page and the behavior of the browser, so don't set the limits too low.

Balance the desires of users against the risks inherent in providing Internet access for uncontrolled systems, and define a policy that fits the environment.

See also:

- *Wireless Logs*
- *Wireless Status*
- *Using External Wireless Access Points*
- *Troubleshooting Wireless Connections*

pfSense® software includes built in wireless capabilities that allow a firewall running pfSense software to be turned into a wireless access point, to use a wireless 802.11 connection as a WAN connection, or both. This chapter covers how to configure pfSense software for these roles as well as suggested means of securely accommodating external wireless access points and how to securely deploy a wireless hotspot. In-depth coverage of 802.11 in general is outside the scope of this documentation. For those seeking such information, see other works such as 802.11 Wireless Networks: The Definitive Guide.

See also:

Hangouts Archive to view the May 2015 Hangout on Wireless Access Points.

CHAPTER

THIRTYTHREE

CELLULAR WIRELESS

pfSense® software can use compatible cellular modems (4G/LTE/5G) as WAN interfaces for connectivity. This can be used as a sole means of connectivity or as a backup.

33.1 Configuring Cellular Modems

To configure a cellular modem in pfSense® software, plug in a *Known Working Modem* and log into the firewall GUI to begin configuration.

See also:

PPP (Cellular Modem)

33.1.1 Example Cellular Configuration

- Navigate to **Interfaces > Assignments**, **PPPs** tab
- Click **+** to create a new entry
- Configure the settings as follows:

  Link Type
  PPP

  Link Interface(s)
  Select the port for the modem from the list. The list contains all available serial ports on the firewall.

  Note: A modem may list several serial ports. Typically, the correct choice is the last entry, but may require trial and error to identify.

  Description
  Text used to reference this PPP configuration in other parts of the GUI. For example, it could contain the service provider and/or modem model.

  Country
  Select the country in which this modem is operating.

  The firewall populates the **Provider** list based on the value of this field.

  Provider
  Select the cellular network provider for the modem.

The firewall populates the **Plan** list based on the value of this field.

Plan
Select the cellular plan used by this modem. Check with the cellular provider to determine which choice is correct.

This populates the remaining fields where possible with values specific to the **Plan**.

Username / Password
Enter login credentials if the provider requires authentication.

The remaining settings can typically be left at their automatic or existing values. For more information on the available settings, see *PPP (Cellular Modem)*.

- Click **Save**.

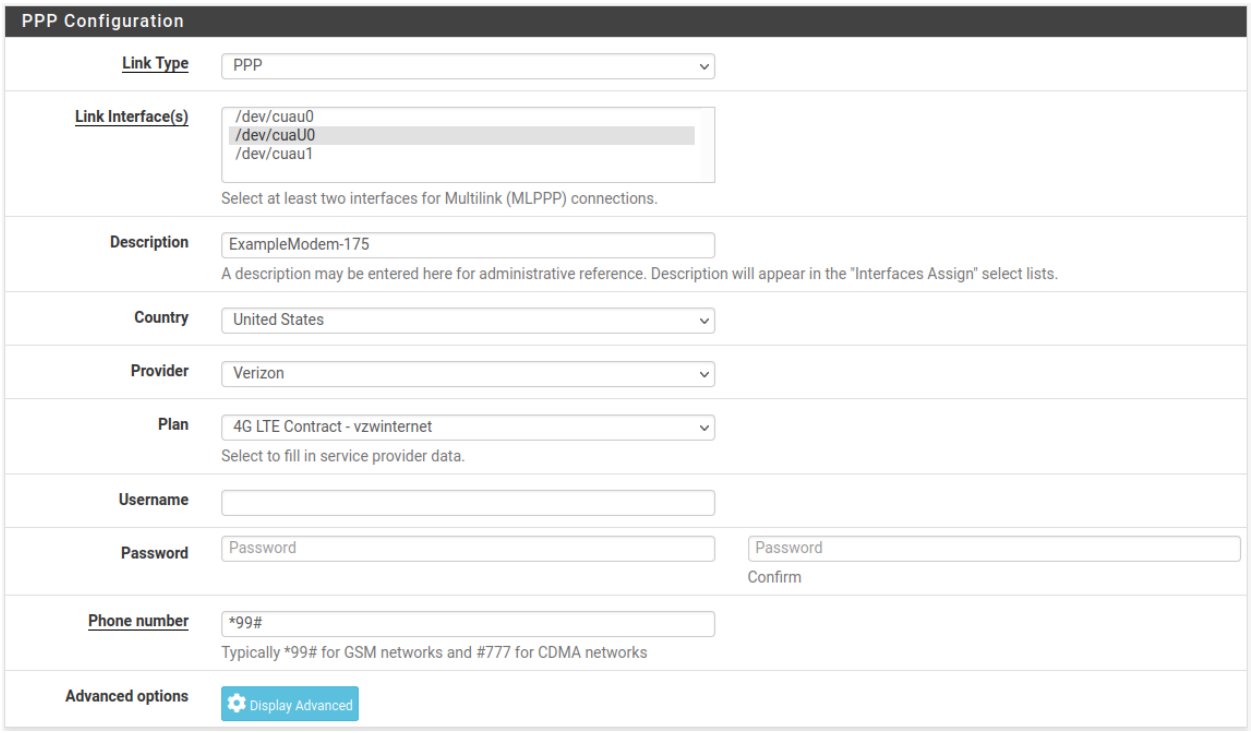


Fig. 1: Example modem configuration

The GUI lists the newly created PPP interface on the **PPPs** tab:



Fig. 2: PPP list entry for a modem

Assigning the PPP Interface

Next, assign the PPP instance to an interface.

- Navigate to **Interfaces > Assignments**
- Set the **Available network ports** field to the newly created PPP interface
- Click  **Add**

The firewall will assign the PPP interface as the next available OPTx interface, for example, OPT1.

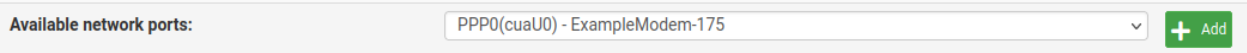


Fig. 3: Assign the PPP interface

Enable the PPP Interface

- Navigate to the menu entry for the new interface name, e.g. **Interfaces > OPT1**
- Check **Enable Interface**
- Enter a new name for the interface in the **Description** field, e.g. `CellModem`
- Click **Save**
- Click **Apply changes**

Check the interface status

Navigate to **Status > Interfaces** and locate the entry for the new PPP interface in the list.

If the interface does not show connected, check the logs under **Status > System logs**, **PPP** tab to see if the firewall has log messages indicating why the connection is failing.

Note: Some connection problems are a lack of signal. If the cellular modem is in a location with poor reception, such as an equipment room, datacenter, rack, etc., then use an antenna and/or longer USB cable to achieve a better signal.

33.2 Known Working 3G-4G Modems

This page lists 3G and 4G modem devices which are known to work with pfSense® software.

- 4G Systems XS Stick P14
- Alcatel Onetouch 4G L850V
- Anydata ADU-635 WA
- Verizon/Pantech UM175
- Verizon/Pantech UML290 and UML295 (see also: *UML290 config info*)
- Verizon USB727
- HP hs2340 HSPA+ Mini Card AMO Ericsson

- Huawei B970/B970B
- Huawei E122
- Huawei E153
- Huawei E156G
- Huawei E160
- Huawei E161
- Huawei E160E
- Huawei E169
- Huawei E172
- Huawei E173
  - Users have reported that some E173 modems are shipping labeled as E173 but have a different chip that is not compatible with existing drivers. See below.
- Huawei E173U-1
- Huawei E176
- Huawei E180
- Huawei E220
- Huawei E226
- Huawei E272
- Huawei E352
- Huawei E353U-2
- Huawei E367
- Huawei E372
- Huawei E392
- Huawei E397u-53
  - Link interface: `/dev/cuaU0.0`
  - Init string: `&F`
- Huawei E398 (E398u-1)
- Huawei E960
- Huawei E1550
- Huawei E1552
- Huawei E1556
- Huawei E1612
- Huawei E1692
- Huawei E1750
- Huawei E1756
- Huawei E1762

- Huawei E1820
- Huawei E3372s LTE USB-stick
  - Link interface: `/dev/cuaU0.1`
  - Init string: `&F&C1&D2E0S0=0`
- Huawei E3372h LTE USB-stick
  - As an Ethernet device: see *Modems reported to work as Ethernet devices*
  - As a modem device: requires manual firmware changes, see this article
- Huawei E3531

  The command to switch to the correct mode depends on the hardware revision.

  Newer versions of the hardware may use this command:

  ```
  usb_modeswitch -v 12d1 -p 1f01 -c /usr/local/share/usb_modeswitch/12d1:1f01
  ```

  Older hardware revisions may use this ID instead:

  ```
  usb_modeswitch -v 12d1 -p 15e7 -c /usr/local/share/usb_modeswitch/12d1:15e7
  ```

  The command has to be executed every time it's detached and reattached, the interface has to be disabled and then enabled again.
  - Link interface: `/dev/cuaU0.0`
  - Init string: `&F0E1Q0 +CMEE=2`
- Huawei K3563
- Huawei E5372
- Huawei E5776
- Huawei VIK K3715 HSU by Vodafone
- Huawei K3765 by Vodafone
- Huawei ME909u-521 4G/LTE Mini-PCIe
  - Link interface: `/dev/cuaU0.0`
  - Init string: `&F0E1Q0 +CMEE=2`
- Huawei ME909s-120 4G/LTE Mini-PCIe
- Nokia Phone E72-1 connected via USB cable
- Sierra Wireless U305
- Sierra Wireless 320u usb LTE
- Sierra Wireless U330
- Sierra Wireless MC7354
- Sierra Wireless MC7355
- Sierra Wireless MC7710
  - May require usb-comp or ID change. See https://forum.netgate.com/post/556751 for more information.
- HP2300 (Sierra Wireless MC8775 3G) Mini-PCIe

- USB Connect Mercury (Sierra Wireless Compass 885 or C885)
- Sierra Wireless Compass 889
- Ovation U727 by Novatel on Sprint CDMA
- Nokia CS-17
- Turkey-TTNET USB Stick 3G Modem. Label says Huawei E173 but its actually Huawei E1800.
- Telstra Maxon bp3-usb (Benchmarked: 2500/350)
- ZTE MF656A
- Vodafone Mobile Connect K3565
- Huawei K4505 (Vodafone Mobile Broadband)
- LTE Yota LiTE LU 156 4G - NOTE: May need nudged in some way out of storage mode. (e.g. boot delay, unplug/re-plug)
- Novatel EU850D (Mini PCIe)
- ZTE MF683 (May need CD-ROM disabled using `AT+ZCDRUN=8` on another system first)
- ZTE MF622
- Ericsson H5321G / Lenovo FRU 60Y3297
- Ericsson F5521GW Gobi3000 / Lenovo
- Ericsson N5321 / Lenovo
  - May need `AT+CFUN=1` in the Init string. Serial port varies from `/dev/cuaU[0-3]`.
- ZTE MF915 LTE modem (T-Mobile)
- ZTE MF190 USB (1&1) using */dev/cuaU0.2*
- ZTE MF669 - May need `camcontrol eject da0` in Shellcmd, uses `/dev/cuaU0.2`
- ZTE MF830 - Can be switched from Ethernet to Modem by accessing the device's web interface, depending on preference.
- ZTE MF861
- ZTE MF985 - Branded as AT&T Velocity 2
- D-Link DWM-157 (3.75HSPA+)
- ONDA MT503HSA Type MF636 (requires eject mode switch, see below)
- Netgear LB1120 (US)
- Netgear LB1121 (US)
- Netgear LB2120 (US)
- Netgear LB1110 (EU)
- Netgear LB1111 (EU)
- Netgear LM1200 (EU)
- And many others

If a modem **DOES WORK** but is not on the list - Please *submit a documentation update*.

If a modem **DOES NOT WORK** - post about it on the Netgate Forum for help, **do not contact Netgate asking for drivers or assistance**.

33.2.1 Modems reported to work as Ethernet devices

- Verizon (Pantech) 295 - Works, but fails if detached and reattached, must reboot.
- ZTE MF60 3g
- ZTE MF90
- ZTE MF823 - Defaults to 192.168.0.1, will need to be sure local system does not have an overlapping network.
- ZTE MF833R - Same
- ZTE MF833V - Same
- ZTE MF915 LTE modem (T-Mobile)
- ZTE MF975S
- ZTE MF79U - Requires Plus 23.01, CE 2.7.0, or later. USB Ethernet device appears after running `camcontrol eject cd0`.
- Huawei E3372-325
- Huawei E3372h - Command to switch to the correct mode:

  ```
  usb_modeswitch -v 12d1 -p 1f01 -c /usr/local/share/usb_modeswitch/12d1:1f01
  ```

  The command has to be executed every time it's detached and reattached, the interface has to be disabled and then enabled again.
- Huawei E5573 - Command to switch to the correct mode:

  ```
  usb_modeswitch -v 12d1 -p 1505 -c /usr/local/share/usb_modeswitch/12d1:1505
  ```

  The command has to be executed every time it's detached and reattached, the interface has to be disabled and then enabled again.
- Huawei E8372h – See *Mode Switching* and #6226
- TP-LINK M7350

33.2.2 Modem variations reported to NOT work

These have the same model numbers as the above, but have different chips and may not be compatible.

- Huawei E173s

```
#Before switching (USB mass storage)
DefaultVendor= 0x12d1
DefaultProduct=0x1c0b
#After switching into modem mode
TargetVendor=  0x12d1
TargetProductList="1c05,1c08"
```

- mPCIe: Sierra Wireless Gobi2000

33.2.3 Mode Switching

Some devices show up as a media device, such as `cd0`, in this case it may be possible to switch modes by executing a command:

```
camcontrol eject cd0
```

If that does switch the modem to the proper mode, it may be added as a *Executing Commands at Boot* using the full path:

```
/sbin/camcontrol eject cd0
```

`usb_modeswitch` is required in order to make certain devices switch to the correct mode.

In some cases, switching to the correct mode is too late, resulting in an error "Network interface mismatch" on the console. In this case, the additional command `&& sleep 10` can be added. For example:

```
/usr/local/sbin/usb_modeswitch -v 12d1 -p 1f01 -c /usr/local/share/usb_modeswitch/
↪12d1:1f01 && sleep 10
```

This package is available in the pfSense software repository, but cannot be installed using the GUI package manager. It can be installed from a shell prompt using `pkg install usb_modeswitch`.

This method should not be used on a production firewall it has not been tested officially.

33.3 Verizon UML290 Cellular Modem

These instructions are specifically for the Verizon 4G modem Pantech UML290.

First, find the phone number associated with the device. One way to find it is to install the VZAccess Manager on a Windows PC and then plug in the device and let it detect the phone number.

Tip: A user also reported that they connected from their Windows PC to ensure the device was activated, but they were unsure if that was necessary.

Follow the instructions here for how to configure the interface in pfSense® software: *Configuring Cellular Modems*.

Some setups may need to modify the instructions because `#777` may not work. Instead, try `*99***3#`. Also for **username** try `devicephone#@vzw4g.com` and for **password** put `vzw`.

33.4 Configuring a Verizon UML295 USB Modem

Connect the UML295 to a computer to do the initial setup and remove the initial SIM PIN.

Test the connection from the computer to ensure that there is 4G connectivity.

If the test succeeded, remove the modem from computer and plug it into the firewall running pfSense® software. There should be no PPP configuration necessary for this device in the firewall as it will register as a USB Ethernet device in the interface list.

Note: Reboot the firewall if the device does not show in the list for assignment.

Assign `ue0` as a DHCP interface, and it should show an IP address in the `192.168.32.0/24` subnet which then should be accessible from the LAN after configuring the necessary rules.

The gateway on the new interface can also be used by the firewall for failover. See *Multiple WAN Connections* for details.

CHAPTER

THIRTYFOUR

TROUBLESHOOTING

34.1 Troubleshooting Asymmetric Routing

Asymmetric routing happens when traffic between two nodes takes a different path in each direction (e.g. A->B->C, C->D->A). This can pose a problem for TCP which has strict state tracking but often does not affect "stateless" protocols such as ICMP or UDP.

34.1.1 Common Scenario

What happens in most cases is this:

- Client sends a TCP SYN packet which arrives at the firewall
- The firewall creates a state table entry
- The firewall sends an ICMP redirect to the client letting it know to reach the target server through an alternate gateway on the same interface
- The target server sends a TCP SYN+ACK packet to the client by a different path that does not pass through the firewall
- The client sends its ACK and later responses to the server using the other gateway, and they are not visible to the firewall
- After 30 seconds the firewall removes its state table entry as the firewall did not see the client and server fully establish a connection
- Some time later the route the client learned via ICMP redirect expires, and the client sends another packet back to the firewall
- Since this packet is not starting a new connection the firewall drops the packet
- The client gets disconnected as it no longer knows how to reach the server

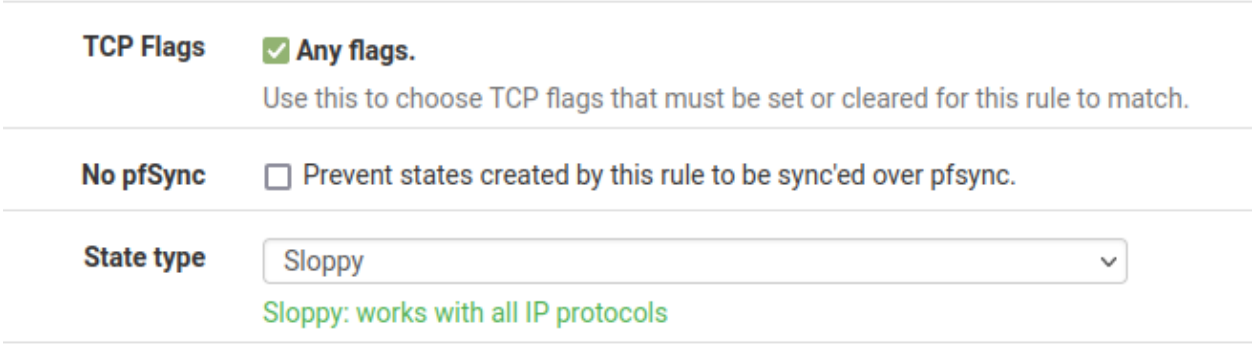
34.1.2 Automatic Fix

The **Bypass firewall rules for traffic on the same interface** option located under **System > Advanced** on the **Firewall & NAT tab** activates rules for traffic to/from the static route networks which are much more permissive when it comes to creating states for TCP traffic and allowing it to pass. The rules allow any TCP packets regardless of their flags to create a state, and also utilize "Sloppy" state tracking which performs a less strict state match.

34.1.3 Manual Fix

The same rules may be created manually by adding one on the affected **interface tab** (e.g. LAN), and a second rule on the **Floating tab** using the same interface (LAN again) to match the traffic in the *out* direction. The rule must be set for a protocol of *TCP*, under **TCP flags** check *Any Flags*, and use a **State Type** of *Sloppy*.

The options for **TCP flags** and **State Type** can be found in the **Advanced Options** when editing a firewall rule.



34.1.4 Alternate Causes

On occasion these issues can be caused by other factors that lead to asymmetric routing, such as issues with `route-to` or `reply-to`, both having to do with gateways on interface settings.

Defining gateways under **System > Routing** does not cause this, but rather these problems can come up when the gateway is improperly configured on the interface pages, **Interfaces > WAN**, **Interfaces > LAN**, and so on.

See also:

WAN vs LAN Interfaces

Gateway set when it should not be set

If a gateway is set on an internal interface, such as LAN, it can cause problematic behavior. Setting a gateway on an internal interface will tag outbound rule on that interface with `route-to`, and inbound rules with `reply-to` which will cause the firewall to forward packets to the defined gateway rather than following their natural path. For WANs this is typically a good thing! For LANs it is not. Among other ill effects, it can lead to a loop of sorts where packets bounce between the firewall and the defined gateway, eventually being blocked or dropped when their TTL expires.

Gateway not set when it should be set

A gateway should usually be set on a WAN or other external-type interface settings (MPLS, IP VPN, etc.) In these cases the `reply-to` and `route-to` behavior is desired and likely required. If it is missing the packets may be blocked or dropped as they attempt to leave the wrong interface. A packet could enter via the alternate WAN, but the reply would leave by the default gateway. Similar to the effect seen when improperly using an *Interface Group* for WAN interfaces.

34.2 Troubleshooting Authentication

Testing authentication servers is possible using the tool located at **Diagnostics > Authentication**.

See also:

- *Authentication Servers*
- *External User Authentication Examples*
- Hangouts Archive to view the August 2015 Hangout on RADIUS and LDAP.

34.2.1 Testing Authentication

Testing user authentication is a simple process:

- Navigate to **Diagnostics > Authentication**
- Select an **Authentication Server**
- Enter a **Username**
- Enter a **Password**
- Click the **Test** button.

Note: This only performs a basic authentication test. Some special use cases, such as EAP, cannot be tested in this manner and may still fail when this test succeeds.

The firewall will attempt to authenticate the user against the chosen server, then it prints the results. The best practice is to try this at least once before attempting to use the server.

If the server returned a set of groups for the user, and the groups exist locally with the same name, the GUI prints the names of the groups in the test results.

If the firewall receives an error when testing authentication, double-check the credentials and the server settings, then make any necessary adjustments and try again.

Debug Option

The **Set debug flag** option enables additional logging for authentication attempts made from this page. Currently, this is only recognized by the LDAP authentication code path.

When debugging is enabled for LDAP authentication, the authentication process will write messages to the main system log with much more detail about the LDAP query and results. This can be a great help for dialing in authentication containers, group membership, and extended query contents. This is especially useful when communicating with an LDAP server over TLS, where a packet capture cannot see those details.

34.2.2 RADIUS Authentication Server Troubleshooting

Missing/Incomplete RADIUS Reply Attributes

There is a limit to the maximum number of attributes the RADIUS client on pfSense® software can receive.

The radius client library in pfSense software does not implement RFC 7499. This restricts the RADIUS request and response payloads to an upper limit of `4096` bytes. Large lists of attributes, such as numerous ACL entries, will be truncated to this limit.

34.2.3 LDAP Authentication Server Troubleshooting

Tip: When troubleshooting LDAP authentication, be sure to check the *Debug Option* for increased LDAP query and response logging.

LDAP DN and Related Settings

For LDAP authentication servers, first ensure the base DN and similar settings match those configured on the LDAP server. Check the LDAP server for more information.

For **Base DN**, it's common to use the root of the LDAP tree but in most cases **Entire Subtree** must also be selected for the **Search Scope**.

Authentication Containers vary by LDAP implementations and setup. On Windows, it is commonly `CN=Users,DC=example,DC=com`, but it may vary. Try using an LDAP browser or similar software to locate the correct container.

LDAP path components are not case-sensitive, so `CN=Administrator` is equivalent to `cn=administrator`.

Bind Credentials

When using authenticated (Not anonymous) binds, the username can be the short name (e.g. `DOMAIN\User` for AD) or a full LDAP specification for a user (e.g. `CN=administrator,CN=Users,DC=example,DC=com`).

Tip: The full DN of a Windows AD bind user can be found by navigating to the user in **ADSI Edit** found under **Administrative Tools** on the Windows server.

For a production setup, an unprivileged user should be used for binding if possible, and not AD Administrator-level account. Such an unprivileged user may need sufficient permissions to attempt binding as other users and access the LDAP directory.

Another common mistake with group membership is not specifying *Entire Subtree* for the **Search Scope Level**.

Active Directory Group Membership

Depending on how the Active Directory groups were made, the way they are specified may be different for things like Authentication Containers and/or Extended Query. For example, a traditional user group in AD is exposed differently to LDAP than a separate Organizational Unit.

ADSI Edit found under **Administrative Tools** on the Windows server can locate the DN for a given group.

Extended Query

The most common mistake with **Extended Query** is that the given directive fails to include **both** the item to search for and how to perform that search, such as:

```
memberOf=CN=VPNUsers,CN=Users,DC=example,DC=com
```

Note that in the above example the DN of the group is given along with the restriction (`memberOf=`).

For users of RFC2307 groups, such as with OpenLDAP, an extended filter might look more like the following:

```
&(objectClass=posixGroup)(cn=VPNUsers)(memberUid=*)
```

Connection-Related Issues (non-SSL/TLS)

Make sure that the LDAP server is listening on the expected port, and that connectivity to the LDAP server network is functional.

Performing a packet capture filtered on the LDAP server IP address and port will help track down the problem. Tools such as **Wireshark** can interpret LDAP packets and help diagnose queries and failures. See *Troubleshooting via Packet Captures*.

Connection-Related Issues (SSL/TLS)

By far the most troublesome connection type is LDAP+SSL/TLS due to its strict security and validation standards.

Restart PHP and the GUI

When making configuration changes to LDAP server entries using SSL/TLS, in some cases PHP and the GUI must be restarted for it to fully utilize the changes. Therefore, it's best to perform this step before others to ensure tests have the best chance to succeed.

To accomplish this, either reboot the device **or** connect to the console or SSH and run menu options `16` then `11`. After restarting PHP and the GUI, run the test again to see if it succeeds.

Hostname Required

When connecting to LDAP with SSL/TLS, the hostname given for the server is also used to verify the server certificate. The server certificate SAN entries and/or CN **must include its hostname**, and that **hostname must resolve to the LDAP server IP address**, e.g. `CN=ldap.example.com`, and `ldap.example.com` is `192.168.1.5`.

If an IP address has been entered for the hostname of the LDAP server, it will not work unless that IP address happens to also be the CN or a SAN of the server certificate.

If this must be worked around, it is possible to create a DNS host override in the DNS forwarder for `<common name of the cert>.<firewall domain name>`. That assumes that the CN is in a format that could actually be a hostname, or that the hostname in question is present in a SAN entry on the server certificate.

Use the Correct Port

When using port `636` for SSL/TLS, the firewall uses an `ldaps://` URL, not STARTTLS. Ensure that the LDAP server is listening on the correct port with the correct mode. For STARTTLS, use port `389`.

Ensure CA Matches

The most important factor in making sure that it is possible to communicate with the LDAP server over SSL/TLS is that the correct CA certificate has been imported into the firewall, and is chosen on the LDAP settings. The key is not required, only the CA certificate.

Nested CAs

If the LDAP server certificate CA is part of a chain, or there is an intermediate CA, every CA certificate in the chain must be imported into the Certificate Manager.

Other Cert/CA Issues

Confirm that the certificates are otherwise valid, for example they are not expired or set to be valid in the future.

Troubleshooting via Server Logs

Authentication failures are typically logged by the target server (FreeRADIUS, Windows Event Viewer, etc.), assuming the request is making it all the way to the authentication host. Check the server logs for a detailed explanation why a request failed. The system log at **Status > System Logs** may also contain information that hints at a resolution.

Troubleshooting via Packet Captures

Packet captures (*Packet Capturing*) can be invaluable for diagnosing errors as well. If an unencrypted method (RADIUS, LDAP without SSL/TLS) is in use, the actual password being used may not be visible, but enough of the protocol exchange can be seen to determine why a request is failing to complete. This is especially true when a capture is loaded in Wireshark, which can interpret the responses, as seen in Figure *Sample LDAP Failure Capture*.

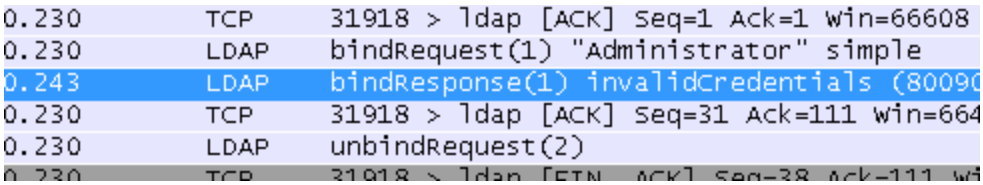


Fig. 1: Sample LDAP Failure Capture

34.3 Troubleshooting Boot Issues

The errors in this article would happen during boot, typically the first boot either from the installation media or immediately after installation. If the system was up and running but then developed a boot issue with no changes in the software/OS, it isn't likely to be related.

34.3.1 Booting with an alternate console

To boot a different console, first get to a loader prompt. Either choose the menu option from the boot menu, or when **Hit [Enter] to boot immediately, or any other key for command prompt.** is seen during the boot process, press space or another key.

Once at the loader prompt, type the following to boot with the serial console active:

```
set console=comconsole
boot -v
```

Similarly, if a serial memstick image is used that prefers the serial console, the video console can be used instead as follows:

```
set console=vidconsole
boot
```

34.3.2 ISA Serial Console not Fully Functional

Older devices with ISA-based serial console ports may not fully detect their console due to changes in how FreeBSD probes serial ports. The console may show boot and kernel messages, but will not show pfSense software boot output or an interactive console.

Devices affected by this are primarily older non-EFI hardware running pfSense CE software. This includes devices such as the APU (1), RCC-DFF 2220, and RCC-VE models such as the 2440, 4860, and 8860.

pfSense plus software attempts to detect known affected models of hardware from Netgate. Other devices may require manual intervention. One workaround is to *add a loader tunable* defining the location of an affected console port:

```
hint.uart.0.at="isa"
```

The value above assumes a standard console on the first serial port. If the device uses the second serial port instead, then change the `0` to `1`.

This value can also be `set` via the loader prompt to test for a single boot. Add the loader tunable after booting successfully, so the setting will be present for future boots.

34.3.3 Booting from USB

If booting fails from a USB 3.0 port and the above does not help, try a USB 2.0 port with the same delay settings. Similarly, if it fails with a USB 3.0 thumb drive, try a USB 2.0 thumb drive. Drive capacity can also make a difference, where some older hardware may not be able to boot large capacity USB drives.

If the boot stops with a `mountroot` error while booting off the installation disc, usually with USB CD/DVD drives, escape to the loader prompt and run the following:

```
set kern.cam.boot_delay="10000"
boot
```

Note: pfSense® software applies this change already when it detects the need, but there may be some edge cases where it should be handled manually.

At which point the boot will continue normally, and a normal installation will be possible.

If running permanently from a medium that requires this delay, add the following as a *Loader Tunable*:

```
kern.cam.boot_delay="10000"
```

34.3.4 Multiple Disk Boot Issues

If a system has multiple disks and pfSense software has been installed on both, it is possible they may conflict in one or more ways.

For details in troubleshooting this type of situation, including identifying which drive the operating system is using, see *Troubleshooting Multiple Disks*.

34.3.5 EFI Boot Issues

Occasionally certain devices or hypervisors may have issues booting pfSense software via EFI. There can be multiple issues here even if they preset using similar symptoms.

A few items to check are:

- Ensure the BIOS is up-to-date from the OEM
- For hypervisors, ensure the hypervisor OS/Software is up-to-date
- Ensure **Secure Boot** is disabled in the BIOS settings.
- For Virtual Machines on Proxmox® VE, check the recommended settings as described in *Booting UEFI*. In particular, adding a **Serial Port** to the VM hardware may improve EFI boot behavior. Shut down/power off the VM and start it back up after adding the port to ensure it's fully added.
- If the BIOS is capable of both Legacy and EFI booting, the boot selection menu may contain two entries for the installation media (Legacy and EFI), try the EFI entry first.
- If the system boots the ISO OK but not the installed OS, dismount the ISO image and/or remove the virtual optical drive from the VM and try again.
- Some EFI-capable BIOS implementations have a method to reset the boot list entries. Resetting or clearing this list may help, though in some rare cases resetting the list may require reinstalling pfSense software afterward.
- Try toggling **Above 4G decoding** in the BIOS settings

- Try the following settings to increase the size of the staging area by adding slop, and to fall back to legacy PCI config access. These must be added as *Loader Tunables*:

```
exec="staging_slop 268435456"
hw.pci.mcfg=0
```

If all else fails, fall back to booting using Legacy mode if possible. The BIOS may need to be adjusted to allow that behavior.

34.3.6 GPT Boot Issues

Some systems may fail to boot a 2.4 memstick because they do not allow booting from GPT, or they are particular about the layout or other attributes.

In these cases, the memstick can be modified using another firewall running pfSense software version 2.3 or later, or with a FreeBSD 10.x or later system.

Insert the memstick into the device with the pfSense or FreeBSD installation and check the system log or `dmesg` output to find the device name, such as `da1`.

First, try to set the partition on the memstick active:

```
gpart recover da1
gpart set -a active da1
```

Remove the memstick and attempt to boot the memstick on the target hardware. If it works, be sure to only use compatible options in the installer. In this case, if the BIOS requires an active GPT partition then either of the following installer choices will work:

- **Auto (UFS)** requires the Partition Scheme **BSD (BSD Labels)**
- **Auto (ZFS)** requires the Partition Scheme **GPT + Active (BIOS)**

Other platforms may need more changes to the memstick, including:

```
gpart bootcode -b /boot/pmbr -p /boot/gptboot -i 2 da1
gpart set -a bootme -i 2 da1
```

In this case the hardware may not be capable of booting ZFS, but use the same option above for UFS with BSD Labels which should work.

In either case, depending on the BIOS and hardware capabilities, similar modifications may be made to the installation target disk to use other partition schemes not listed above, but that has not been tested internally.

34.3.7 BIOS/Disk Errors

- If after installation, a "cannot load kernel" error is observed, or the firewall hangs at the spinner (`/`):
  - Make sure BIOS is up-to-date
  - Reinstall but do not check the packet mode option during boot block installation process
  - Set the HDD mode in BIOS to LBA (don't use "auto", the detected geometry is different if it is set to "auto", which fails)
  - Set the HDD mode in BIOS to CHS if the above fails
  - Set AHCI mode in the BIOS

- Try using multiple partitions, one small one (~4GB) for / and another for /usr using the rest of the disk.
- If the installer fails to start correctly (system reboots), try the following (attempt each sub-step, then rerun the installer each time):
  - Make sure BIOS is up-to-date
  - Change the hard drive ribbon/SATA cable
  - Force a slower speed in the bios for the channel
  - Boot from another disk and run: (note that ad0 is the first ATA hard drive):

```
# dd if=/dev/zero of=/dev/ad0 bs=8k count=16
# fdisk -I ad0
```

- If all those fail…
  - Partition from a FreeBSD Live CD or Linux
  - Different hard drive
  - Different hardware entirely that isn't prone to legacy booting issues

34.3.8 DVD Errors

- If the installer disc starts to boot, but hundreds of errors are printed, try:
  - Make sure BIOS is up-to-date
  - Install from a USB memstick instead
  - Burn the disc at a slower speed
  - Change the DVD drive cable
  - Different DVD drive

34.3.9 Boot Blocks/Loader Issues

- If a read error occurs during boot, please see this Boot Error.
- If FreeBSD will boot but not pfSense, try booting from a FreeBSD Live CD and running the following (More Info):

```
# fdisk -B -b /boot/boot0 /dev/ada0
# bsdlabel -B /dev/ad0as1
```

Note: ada0 is the first hard drive in that example

34.3.10 Vendor-Specific Issues

- Certain Dell Blade servers may hang at boot if the system's virtual USB media is enabled. Disable the virtual media in the BIOS and then it should boot normally.
- Certain systems running Hyper-V on AMD processors may need to do the following:

  Escape to the loader prompt during boot and run:

  ```
  set hw.clflush_disable=1
  boot
  ```

  At that point, boot the rest of the way and install pfSense software. After installation, add the following line as a *Loader Tunable*:

  ```
  hw.clflush_disable=1
  ```

If all else fails, Netgate offers Netgate TAC and hardware through the Netgate Store that is preloaded with pfSense Plus software.

34.3.11 "Fake" RAID cards with a GRAID error

Certain "fake" RAID cards, driver/software-based RAID adapters that are not true hardware RAID, may fail to mount properly with the following error:

```
Root mount waiting for: GRAID
mountroot>
```

Another symptom can be that "Intel RAID" messages are shown during the boot sequence, and typing `?` at the `mountroot>` prompt it only shows the drive itself and no partitions:

```
Mounting from ufs:/dev/ada0s1a failed with error 19
mountroot> ?
[...]
ada0
```

- Escape to a loader prompt during boot and run:

  ```
  set kern.geom.raid.enable="0"
  boot
  ```

- After a successful install/boot, add that settings permanently as a *Loader Tunable*:

  ```
  kern.geom.raid.enable="0"
  ```

34.3.12 Conflicting Hardware

If the system stops with an error such as:

```
run_interrupt_driven_hooks: still waiting after 60 seconds for xpt_action
```

Disable any Firewire/1394 controllers and built-in USB Card Readers in the BIOS.

34.4 Troubleshooting Multiple Disks

If a system has multiple disks and pfSense software has been installed on both, it is possible they may conflict in one or more ways. For example, this can happen if an older disk was left in place after adding a new disk of a different type and reinstalling to the new disk.

In these situations best practice is to remove the unused disk, but that is not always possible. For example, if the original installation was using an embedded disk such as eMMC. If the disk cannot be removed, then the next best solution is to clear the metadata from the unused disk.

A common way multiple disks conflict is if they both use the same ZFS label. In that case it is unpredictable which ZFS pool will be used by the OS, and it may change depending on the boot order. Another way is if the OS boots the kernel from one disk but mounts the other disk in the operating system, leading to a situation where the installed OS appears up-to-date but is booting with an outdated kernel.

34.4.1 Identify the Disk

To clear the metadata safely, first identify the unused disk. This may take some investigation, but typically disks are listed in the full boot log output in `/var/log/dmesg.boot`, the output of `sysctl kern.disks`, along with other OS commands such as `geom list disk`, `gpart list`, and `geom -t`. In some cases it's clear which is which, such as when using an add-on SSD instead of eMMC, where the eMMC disk is named `mmcsdX` and the SSD is `ndaX` or `adaX`.

It's also possible that the drive that loaded the kernel at boot time is different from the drive mounted as the root of the filesystem (`/`). Once booted, it's not possible to determine which drive loaded the kernel, but it is possible to determine which drive holds the root filesystem.

Note: If the unused disk cannot be definitively identified, take a backup, clear the data from *all* disks, and then reinstall.

When a system has multiple disks, odds are high that the disk holding the live root filesystem is the intended disk and whichever disk is *not* used for the root filesystem is the one that should be wiped to avoid conflicts.

UFS

On UFS systems, look at the output of `df /` in the first column (`Filesystem`):

- If it's a disk device (e.g. `/dev/ada0s2a`), then note it and move on. The filesystem device name will include a slice or partition identifier at the end (e.g. `s2a` in the previous example) but it should be possible to match the disk name against the list in `sysctl kern.disks`.

  In this case, it is located on `ada0s2a` which is the first filesystem in the second slice of the disk `ada0`.

- If it contains a label (e.g. `/dev/diskid/`, `/dev/gpt/`, or `/dev/ufs/`), look at the output of `glabel status` and match the start of the filesystem label with the `Name` column and find the disk device in the corresponding `Components` column.

```
$ df /
Filesystem                   1K-blocks    Used   Avail Capacity  Mounted on
/dev/diskid/DISK-9D1CEC59s2a   7353532 1664148 5101104    25%    /
$ glabel status
                  Name  Status  Components
diskid/DISK-9D1CEC59     N/A  ada0
```

In this case, the root filesystem is located on `/dev/diskid/DISK-9D1CEC59s2a` which is the first filesystem in the second slice of disk ID `DISK-9D1CEC59`, which corresponds to the disk `ada0`.

Other label types may not always include a slice or partition identifier.

ZFS

For systems using ZFS, check output of `zpool status` and look at the disk names in the output:

```
$ zpool status
  pool: pfSense
 state: ONLINE
  scan: scrub repaired 0B in 00:00:17 with 0 errors on Wed Feb 22 11:03:52 2023
config:

     NAME        STATE     READ WRITE CKSUM
     pfSense     ONLINE       0     0     0
       nda0p4    ONLINE       0     0     0

errors: No known data errors
```

In this output, the ZFS pool is located on `nda0p4` which is the fourth partition on the disk `nda0`.

Using the Geom Tree

If there is any doubt about the devices in question based on the filesystem device, run `geom -t`. That command outputs a tree style view of all disks and their components, such as slices/partitions. This can make it relatively simple to narrow down the disk which contains a given ID, partition, or slice with minimal searching through command output:

```
$ geom -t
Geom                                Class      Provider
nda0                                DISK       nda0
  nda0                              PART       nda0p1
    nda0p1                          LABEL      gpt/efiboot0
      msdosfs.gpt/efiboot0          VFS
      gpt/efiboot0                  DEV
    nda0p1                          DEV
  nda0                              PART       nda0p2
    nda0p2                          LABEL      gpt/gptboot0
      gpt/gptboot0                  DEV
    nda0p2                          DEV
  nda0                              PART       nda0p3
    swap                            SWAP
    nda0p3                          DEV
  nda0                              PART       nda0p4
```

(continues on next page)

(continued from previous page)

```
    nda0p4                          DEV
    zfs::vdev                       ZFS::VDEV
  nda0                              DEV
mmcsd0                              DISK       mmcsd0
  mmcsd0                            DEV
  mmcsd0                            LABEL      diskid/DISK-9D1CEC59
    diskid/DISK-9D1CEC59            DEV
    diskid/DISK-9D1CEC59            PART       diskid/DISK-9D1CEC59s1
      diskid/DISK-9D1CEC59s1        DEV
      msdosfs.diskid/DISK-9D1CEC59s1 VFS
    diskid/DISK-9D1CEC59            PART       diskid/DISK-9D1CEC59s2
      diskid/DISK-9D1CEC59s2        DEV
      diskid/DISK-9D1CEC59s2        PART       diskid/DISK-9D1CEC59s2a
        diskid/DISK-9D1CEC59s2a     DEV
        ffs.diskid/DISK-9D1CEC59s2a VFS
mmcsd0boot0                         DISK       mmcsd0boot0
  mmcsd0boot0                       DEV
mmcsd0boot1                         DISK       mmcsd0boot1
  mmcsd0boot1                       DEV
```

34.4.2 Clear the Disk

In these examples the unused disk is `mmcsd0`.

The commands in these examples must be run from a console or SSH shell prompt. Do not attempt to execute these commands from the GUI. The best practice is to run them from the console and to have installation media on hand in case a reinstallation is necessary.

Tip: If any of the commands generate an error, boot the *Netgate Installer* and perform the commands from a shell launched through the installer menu (AMD64 and AARCH64). When booted from install media, the disks in the device will not be mounted and can be safely cleared. For ARMv7 devices, boot the recovery installer and use `Ctrl-Z` to suspend the recovery process and reach a shell prompt to run the commands.

Wipe Metadata

The quickest and easiest way to wipe a disk is to clear its metadata.

The following commands clear the disk partition metadata, ZFS metadata, and also wipe the start of the disk to clear the partition table and other data at the beginning of the disk. Depending on the situation it may only be necessary to clear the ZFS metadata, but it's safer to clear it all.

```
### Stop a legacy style GEOM mirror and clear its metadata from all disks
### Mirror name may vary, check "gmirror status" output.
# gmirror destroy -f pfSenseMirror

### Clear the ZFS label (exact partition may vary)
# zpool labelclear -f /dev/mmcsd0p4

### Clear the partition metadata
# gpart destroy -F mmcsd0
```

(continues on next page)

(continued from previous page)

```
### Wipe the first 1MB of the disk
# dd if=/dev/zero of=/dev/mmcsd0 bs=1M count=1 status=progress
```

Note: Alternately, skip the first two commands and omit the `count=1` on dd to wipe the entire target disk from start to end.

Wipe Start and End of Disk

Another tactic is to wipe only the start and end of the disk. However, this approach is a much more complicated process as it involves calculations based on the sector size and number of sectors on the disk:

```
### Wipe the first 1MB of the disk
# dd if=/dev/zero of=/dev/mmcsd0 bs=1M count=1 status=progress

### Wipe the last 1MB of the disk
# dd bs=`diskinfo mmcsd0 | awk '{print $2}'` \
     if=/dev/zero \
     of=/dev/mmcsd0 \
     count=`diskinfo mmcsd0 | awk '{print ((1024 * 1024) / $2)}'` \
     seek=`diskinfo mmcsd0 | awk '{print $4 - ((1024 * 1024) / $2)}'` \
     status=progress
```

Note: Be sure to replace **every** instance of the target disk in each command, as the disk is referenced numerous times to obtain the necessary calculation numbers.

34.5 Troubleshooting "No buffer space available" Errors

On occasion traffic on a NIC may have trouble getting out with an error similar to:

```
ping: sendto: No buffer space available
```

Or:

```
dpinger WANGW x.x.x.x: sendto error: 55
```

The most common causes of this are:

- No route to the target network (or no default route)
- Missing link route for a local target
- Stale state in PF sending the connection out an invalid path (reset states)
- Network memory buffer exhaustion

  See also:

  See *Hardware Tuning and Troubleshooting*

- Faulty NIC and/or driver issue

  Sometimes resetting the NIC can bring it back again:

```
# ifconfig ix3 down; ifconfig ix3 up
```

- Faulty cable
- Traffic shaping (ALTQ or Limiters) dropping the packet
- Virtual NIC being throttled by the hypervisor or host, such as an AWS instance using more throughput than an instance size can handle

  In this case, change the throttling in the host (not guest) or upgrade to a larger instance/higher tier on a hosted platform such as AWS.
- Virtual NIC being disconnected/disabled in certain hypervisors
- An otherwise overloaded NIC exhausting its send/receive buffers
- Other various switch/buffer/connectivity issues

Trying to bounce the NIC with `ifconfig` is the easiest thing to try first. After that, save/apply the interface settings on each interface (or at least WANs and the LAN in question). Check/(re)set the default route if it has been lost. *Reset States*. Replacing the cable may also help. Removing traffic shaping if it is enabled is also a good test.

Otherwise, investigate the traffic on the NIC and look for other buffer-related causes. Contact Netgate TAC for assistance in diagnosing the issue, or post on the forum/mailing list.

34.6 Troubleshooting Captive Portal

This section contains troubleshooting tips for the most common problem with captive portal.

34.6.1 Authentication failures

Authentication failures are normally the result of users entering an incorrect username or password. In the case of RADIUS authentication, these can occur because of connectivity problems to the configured RADIUS server(s), or problems on the RADIUS server itself. Check the RADIUS server logs for indications of why access was denied, and ensure the firewall can communicate with the RADIUS server.

For local users, if the option to require the Captive Portal login privilege is enabled, ensure the users have the privilege directly or are members of a group with the privilege.

34.6.2 Captive Portal Does not Redirect

If clients are not being redirected to the portal page when attempting to browse on an interface with captive portal enabled, it's typically one of the following causes:

DNS resolution not functioning
: Clients on the captive portal interface must either be using the DNS resolver or forwarder on pfSense® software, on the IP address of the interface where the client resides (which is the default configuration), or if using another IP address for DNS, it must be in an allowed IP address entry. If DNS fails, the browser never issues the HTTP request, hence it cannot be intercepted and redirected.

Firewall rules on the captive portal interface do not allow the initial HTTP request
: If the user is trying to browse to google.com, but HTTP connections are not allowed to google.com, the HTTP request will be blocked and hence cannot be redirected. Under **Firewall > Rules**, on the

interface where captive portal is enabled, the traffic to be redirected must be allowed to pass. This is most commonly HTTP to any destination.

The client has an HTTPS home page
The request must be to an HTTP site in order for the portal to redirect the client. If HTTPS is enabled for the portal, this may still work, but it depends upon the client browser or operating system automatic portal detection to work.

Many modern browsers and operating systems detect portals now which makes this less of an issue.

Client is using IPv6
Captive Portal is not currently compatible with IPv6, so IPv6 traffic is not able to traverse the portal interface. If the interface has IPv6 configured and the client attempts to use it, it may encounter network timeouts.

34.6.3 Apple devices are unable to load the portal page or login

Certain versions of Safari on iOS do not properly handle the login form for the Captive Portal page. The most common resolution is to disable autofill for forms in Safari on iOS.

In some cases, Apple devices will not automatically prompt for a Captive Portal login or test for its presence if the wireless network uses encryption. In these cases, manually open a browser and navigate to an HTTP site to get the login redirect.

There have also been reports that on older version of macOS, a Mac would refuse to load any HTTPS sites, including an HTTPS portal, until it could load a CRL and OSCP URL for the certificate. This has been fixed in current versions of OS X.

Some users have had to add `www.apple.com` to their allowed hostnames so that Apple's call to their test page succeeds.

34.6.4 Port Forwards Behind Portal Only Work When Target Logs In

This is a side effect of how the portal operates. No traffic is allowed to reach a host behind the portal unless it has been authenticated or passed through the portal. If a port forward must always work to a device behind the portal, then it must be setup to bypass the portal with either a Pass-through MAC entry (*MAC Address Control*) or an Allowed IP Address entry (*Allowed IP Address*) to allow traffic *To* the target.

34.6.5 Captive Portal Rules

On pfSense Plus software version 22.05 or CE version 2.7.0 or later, Captive Portal uses PF features for L2 `ether` processing under the hood. When having issues with the captive portal, it is possible to inspect the rules for debugging purposes.

To see rules for Captive Portal look in `/tmp/rules.debug` at the multiple sections starting with a comment `# Captive Portal`.

Anchors

The entries for various client and other entries for Captive Portal are kept under special anchors in the ruleset.

The easiest way to see all the anchors and rules is to use the script `pfanchordrill` playback script.

```
# pfSsh.php playback pfanchordrill
```

The relevant anchors are named: `cpzoneid_<id number>_<purpose>` and entries are nested underneath there in subordinate anchors.

The purpose names include:

cpzoneid_<id number>_auth
: Entries for authenticated clients in a zone.

cpzoneid_<id number>_allowedhosts
: Entries for permanently allowed hosts in a zone.

cpzoneid_<id number>_passthrumac
: Entries from the pass-through MAC list.

For example the output of the `pfanchordrill` script includes the following entry for an authenticated captive portal client on `10.7.0.10`:

```
cpzoneid_2_auth/10.7.0.10_32 rules/nat contents:
ether pass in quick proto 0x0800 from 5e:13:3b:ef:4d:24 l3 from 10.7.0.10 to any tag␣
↪cpzoneid_2_auth dnpipe 2000
ether pass out quick proto 0x0800 to 5e:13:3b:ef:4d:24 l3 from any to 10.7.0.10 tag␣
↪cpzoneid_2_auth dnpipe 2001
```

See `/etc/inc/captiveportal.inc` for information on other anchors and rules.

See also:

For assistance in solving problems, post on the Captive Portal category of the Netgate Forum.

34.7 Troubleshooting Cisco VPN Pass Through

If trouble is encountered when attempting a connection from an internal Cisco VPN client to an external host, (e.g. a workstation with the Cisco client is trying to get out through pfSense® software to connect to a remote site), then try the following.

34.7.1 Workaround

- In the Cisco VPN client software, Modify the connection and turn off transparent tunneling completely in the **Transport** tab.
- In the pfSense software GUI, under **Firewall > NAT** on the **Outbound** tab:
  - Enable **Manual Outbound NAT**.
  - Remove any NAT rules that perform static port NAT on UDP/500.

34.8 Troubleshooting Network Connectivity

The following list covers most causes of outbound connectivity failure in common usage scenarios.

Each test assumes the items above it have been checked. This document assumes a single WAN, but most of the advice is relevant to multiple WANs.

34.8.1 WAN Interface

- Check the WAN IP address (**Interfaces > WAN**)
  - This is only relevant to static WANs, dynamic WANs handle addresses automatically
  - Using the wrong address could prevent the ISP from delivering traffic to/from the firewall, among other issues
- Check that the WAN IP address has the correct subnet mask (**Interfaces > WAN**)
  - This is only relevant to static WANs, dynamic WANs handle subnet masks automatically
  - An improper subnet mask such as `/1` could cause connectivity issues to large portions of the Internet, using `/32` for a mask could prevent the firewall from contacting its gateway
- Check that WAN has a gateway and that the gateway IP address is correct (**Interfaces > WAN**)
  - This is only relevant to static WANs, dynamic WANs handle gateways automatically
  - This interferes with automatic outbound NAT and `route-to`/`reply-to`
- Check the default gateway configuration (**System > Routing**)
  - Without a default gateway traffic has no exit path
  - If it is set to *Automatic*, the automatic selection process may have chosen a non-viable gateway
- Check that the default gateway shows **Online** (**Status > Gateways**)
  - If it is not, verify the WAN settings and gateway settings, or use an alternate monitor IP address
- Check the default gateway in the routing table (**Diagnostics > Routes**)
  - Another source such as a VPN may have changed the default gateway

34.8.2 LAN Interface

- Check the LAN IP address (**Interfaces > LAN**)
  - Using an invalid IP address (e.g. `.0` or `.255` in a `/24`) will cause problems reaching addresses locally.
- Check the LAN subnet mask (**Interfaces > LAN**)
  - Using an incorrect subnet mask, such as `/32`, will prevent other hosts in the LAN subnet from finding the firewall LAN address to use as a gateway and vice versa
- Check that LAN does NOT have a gateway set (**Interfaces > LAN**)
  - This will interfere with automatic outbound NAT
- Check that LAN does NOT have **Block Private Networks** set (**Interfaces > LAN**)
  - If the LAN subnet is using a private network, this will block local traffic.
- Check that LAN does NOT have **Block Bogon Networks** set (**Interfaces > LAN**)

  - If the LAN subnet is using a private network, this will block local traffic.

34.8.3 Firewall/Rules

- Check the firewall log for blocked connections from hosts on LAN (**Status > System Logs**, **Firewall** tab)
  - If the log contains entries showing blocked connections, check the rule that triggered the block and adjust rules accordingly (**Firewall > Rules**, **LAN** tab)
- Check that the LAN rule allows all protocols, or at least TCP and UDP ports for reaching DNS and HTTP/HTTPS, and allows ICMP for testing. (**Firewall > Rules**, **LAN** tab)
  - Not allowing UDP would make DNS fail, among other things.
  - Similarly, on a DNS rule, using UDP only and not TCP/UDP will cause larger queries to fail.
  - Not allowing ICMP would cause ping to fail, but other protocols may work
  - Not allowing TCP would cause HTTP, HTTPS, and other protocols to fail.
- Check that the LAN rules allow to a destination of *any* (**Firewall > Rules**, **LAN** tab)
  - Using the wrong destination would not allow traffic to reach the Internet. For example, *WAN net* is only the subnet of the WAN interface, **NOT** the Internet, so typically the correct setting is *any*.
- Check that the LAN rule does not have an improper gateway set (**Firewall > Rules**, **LAN** tab)
  - If it is set to leave by another (possibly broken) non-WAN gateway it would cause the connections to fail

34.8.4 Outbound NAT

- Check **Outbound NAT**, ensure it is set for *Automatic* or *Hybrid* outbound NAT (**Firewall > NAT**, **Outbound** tab)
  - If the firewall requires manual outbound NAT, skip to the next test
  - Incorrect NAT settings will prevent traffic from reaching WAN
- Check manual outbound NAT rules, if in use, to ensure that they match local traffic sources
  - Incorrect NAT settings will prevent traffic from reaching WAN

34.8.5 Diagnostic Tests

- Check connectivity from the firewall itself: Try to ping `8.8.8.8` (**Diagnostics > Ping**)
  - If this does not work, ensure proper WAN settings, gateway, etc.
- Check DNS: Try to look up `pfsense.org` (**Diagnostics > DNS Lookup**)
  - If this does not work, fix/change the DNS configuration (*Troubleshooting DNS Resolution Issues*)
- Test NAT: Try to ping `8.8.8.8` using *LAN* as the **Source Address** (**Diagnostics > Ping**)
  - If this fails but the other tests work, then the problem is likely outbound NAT (See the WAN/LAN gateway checks above)

34.8.6 Client Tests

- Test if the client can ping the LAN IP address of the firewall
  - If this fails, check the LAN rules, client IP address/subnet mask, LAN IP address/subnet mask, etc.
- Test if the client can ping the WAN IP address of the firewall
  - If this fails, check the client subnet mask and gateway
- Test if the client can ping the WAN Gateway IP address of the firewall
  - If this fails, check the client subnet mask and gateway, and double check outbound NAT on the firewall
- Test if the client can ping an Internet host by IP address (e.g. `8.8.8.8`)
  - If this fails, check the client subnet mask and gateway, and triple check outbound NAT on the firewall
- Test if the client can ping an Internet host by Host name (e.g. `www.google.com`)
  - If this fails, check the client DNS settings, and/or the DNS Resolver or Forwarder on the firewall (**Services > DNS Resolver**, **Services > DNS Forwarder**, **Diagnostics > DNS Lookup**)

34.8.7 Miscellaneous Additional Areas

- If Captive Portal is enabled, disable it temporarily (**Services > Captive Portal**).
  - See *Captive Portal Troubleshooting*.
- Check for packages such as pfBlockerNG, Snort, or Suricata, that might interfere with connectivity and disable them if necessary
  - Improperly configured packages could allow certain traffic such as ICMP ping to work but might prevent access to HTTP and/or HTTPS sites.

34.9 Troubleshooting GUI Connectivity

If the GUI is not accessible from the LAN, the first thing to check is cabling. If the cable is a hand-made cable or *shorter* than 3 feet/1 meter, try a different cable. If the client PC is directly connected to a network interface on the firewall, a crossover cable may be needed on older hardware that does not have network cards with Auto-MDIX. This is not a concern on gigabit or faster hardware.

Once there is a link light on both the client network card and the firewall LAN interface, check the TCP/IP configuration on the client PC. If the DHCP server is enabled on the firewall, as it is by default, ensure that the client is also set for DHCP. If DHCP is disabled on the firewall, hard code an IP address on the client residing in the same subnet as the firewall LAN IP address, with the same subnet mask, and use the firewall LAN IP address as the gateway and DNS server.

If the cabling and network settings are correct, the client will be able to ping the LAN IP address of the firewall. If the client PC can ping, but not access the GUI, there are still a few more things to try. First, if the error on the client PC is a connection reset or failure, then either the server daemon that runs the GUI is not running or the client is attempting to use the wrong port. If the error is instead a connection timeout, that points more toward a firewall rule.

If the client receives a connection timeout, refer to *Troubleshooting Access when Locked Out of the Firewall*. With a properly configured network connection, this shouldn't happen, and that section offers ways to work around firewall rule issues.

Double check that WAN and LAN are not on the same subnet. If WAN is set for DHCP and is plugged in behind another NAT router, it may also be using `192.168.1.1`. If the same subnet is present on WAN and LAN, unpredictable results

may happen, including not being able to route traffic or access the GUI. When in doubt, unplug the WAN cable, reboot the firewall, and try again.

If the client receives a connection reset, first try to restart the GUI server process from the system console, typically option 11, followed by option 16 to restart PHP-FPM. If that does not help, start a shell from the console (option 8), and type:

```
# sockstat | grep nginx
```

The firewall will return a list of all running `nginx` processes, and the port upon which they are listening, like this:

```
root     nginx      41948 5  tcp4   *:443                 *:*
root     nginx      41948 6  tcp6   *:443                 *:*
root     nginx      41948 7  tcp4   *:80                  *:*
root     nginx      41948 8  tcp6   *:80                  *:*
```

In that output, it shows that the process is listening on port 443 of each interface on both IPv4 and IPv6, as well as port 80 for the redirect, but that may vary based on the firewall configuration.

Try connecting to the firewall LAN IP address by using that port directly, and with both HTTP and HTTPS. For example, if the LAN IP address was `192.168.1.1`, and it was listening on port `82`, try `http://192.168.1.1:82` and `https://192.168.1.1:82`.

34.10 Troubleshooting OS Issues with a Debug Kernel

Starting with pfSense® Plus software version 23.01 there is a new package containing an operating system kernel and module set built with additional debugging enabled. This is known as a "Debug Kernel".

Note: In the past, there was a similar package which installed debug symbols, but it did not include a full debug kernel or modules.

34.10.1 Installing the Debug Kernel

The debug kernel can be installed from the command line:

```
# pkg update
# pkg install -y pfSense-kernel-debug-pfSense
```

> **Warning:** The debug kernel consumes a significant amount of disk space, ensure there is at least 1GB free.

Once installed, the debug kernel and associated module files are located in `/boot/kernel.debug` and debug symbols are under `/usr/lib/debug/boot/kernel/`.

34.10.2 Booting the Debug Kernel

After installing the debug kernel it can be activated in one of several ways:

Loader Menu

From the loader menu:

- Select menu option `6` until it says `kernel.debug`
- Select menu option `1` or press the `Enter` key to boot

This will boot the debug kernel one time and the next boot will revert to the default kernel.

Loader Prompt

To reach a loader prompt, select menu option `3` from the loader menu or press the space bar during the kernel spinner on systems without a loader menu.

At the loader prompt, enter the following at the `OK` prompt:

```
boot kernel.debug
```

This will boot a debugging kernel one time and the next boot will revert to the default kernel.

Loader Configuration

To persistently boot from the debug kernel, add the following line as a *Loader Tunable*:

```
kernel="kernel.debug"
```

Save and exit, then reboot the firewall as usual from the GUI or console menu.

The next boot and all subsequent boots will use the debug kernel.

> **Warning:** Do not leave the debug kernel persistently active when making a firmware upgrade. The upgrade process relies on booting from the default non-debug kernel. The package may remain installed, but it **must not be active** as a loader tunable during the upgrade process.

To temporarily switch back to the default non-debug kernel, use the loader menu or loader prompt as described in the previous sections, but boot `kernel` instead of `kernel.debug`.

To permanently switch back to the non-debug kernel, remove the `kernel` line from the *Loader Tunables*.

34.10.3 Removing the Debug Kernel

To remove the debug kernel, first boot from the non-debug kernel. Then, check for and remove any reference to the debug kernel in *Loader Tunables*.

Finally, remove the debug kernel package from the command line:

```
pkg delete pfSense-kernel-debug-pfSense
```

34.11 Troubleshooting Offline DHCP Leases

The **Status > DHCP Leases** page only reports clients as "online" if the MAC address for a given system appears in the ARP table on the firewall. This can be verified by checking **Diagnostics > ARP Table**.

Systems which have not communicated with or via the firewall in the past few minutes will appear as offline.

To check a system, try to ping it from **Diagnostics > Ping**. Even if the system does not respond to ping, that action will cause the system to appear in the ARP table if it is on the network, and would thus show online in the DHCP Leases list.

The *Arping Package* may also be of interest. It is available under **System > Packages**. The *Nmap package* can also be used to perform an ARP-based subnet scan to locate online hosts even if they don't respond to ICMP.

34.12 Troubleshooting DHCPv6 Client XID Mismatches

An IPv6 WAN configured to obtain is address via DHCPv6 can suddenly find itself without an IPv6 address if the transaction ID for the IPv6 DHCP client does not match.

When this happens, a log message similar to the following may appear in the DHCP and/or System logs:

```
dhcp6c[xxxxx]: client6_recvadvert: XID mismatch
```

When this happens, the most common cause is having multiple running copies of the DHCPv6 client (`dhcp6c`) on the same interface. Both clients send out a request with different transaction IDs and then get confused by the responses.

When this happens, the quickest way to ensure the clients are reset is to kill the duplicates and start the client again.

From the shell or from **Diagnostics > Command Prompt**, first check for duplicate clients:

```
# ps uxawww | grep dhcp6c
root xxxxx 0.0 0.0 5780 1488 ?? INs Sat09PM 0:00.90 /usr/local/sbin/dhcp6c -d -c /var/
↪etc/dhcp6c_wan.conf -p /var/run/dhcp6c_re1.pid re1
root xxxxy 0.0 0.0 5780 1524 ?? Is Tue07AM 0:00.30 /usr/local/sbin/dhcp6c -d -c /var/etc/
↪dhcp6c_wan.conf -p /var/run/dhcp6c_re1.pid re1
```

Once it has been confirmed that a duplicate client is running, `kill` them:

```
# killall -9 dhcp6c
```

Then navigate to **Interfaces > WAN**, click **Save**, then click **Apply Changes**.

The WAN should now have its IPv6 address once again.

34.13 Troubleshooting DMA and LBA Errors

34.13.1 Non-Fatal Errors

SETFEATURES Error

pfSense® software attempts to configure APM for an ATA hard drive at boot. Sometimes the detection yields a non-fatal error when trying to detect if the feature is available, or when the drive claims it's available but can't set the value.

That error is:

```
ad0: FAILURE - SETFEATURES 0x85 status=41<READY,ERROR> error=4<ABORTED>
```

The error should only appear once at boot on such systems, and can safely be ignored.

34.13.2 Fatal Errors

The following errors would indicate more serious problems such as a faulty HDD/SSD, faulty cable/controller, a faulty SATA/SD/IDE converter, a device out of space, or other similar condition.

```
ad0: FAILURE - READ_DMA status=51<READY,DSC,ERROR> error=40<UNCORRECTABLE> LBA=3919
```

```
ad0: WARNING - READ_DMA UDMA ICRC error (retrying request) LBA=207
ad0: WARNING - READ_DMA UDMA ICRC error (retrying request) LBA=207
ad0: FAILURE - READ_DMA status=51<READY,DSC,ERROR> error=84<ICRC,ABORTED> LBA=207
g_vfs_done():ad0s1a[READ(offset=65536, length=8192)]error = 5
```

```
ad0: FAILURE - WRITE_DMA status=51<READY,DSC,ERROR> error=4<ABORTED> dma=0x06 LBA=1129359
```

34.13.3 Other Errors

The following errors have been observed on certain hardware platforms running pfSense software. While they do not appear to be fatal, the cause appears to be a disk driver issue in FreeBSD (9.2 and later) and it may degrade performance. If these errors are encountered, attempt to reproduce the problem with a stock FreeBSD 10.1 or later installation and report the problem directly to FreeBSD.

```
(ada0:ata0:0:1:0): READ_DMA. ACB: c8 00 3f 95 07 40 00 00 00 00 08 00
(ada0:ata0:0:1:0): CAM status: ATA Status Error
(ada0:ata0:0:1:0): ATA status: ff (BSY DRDY DF SERV DRQ CORR IDX ERR), error: 00 ()
(ada0:ata0:0:1:0): RES: ff 00 46 95 07 00 00 00 00 00 00
(ada0:ata0:0:1:0): Retrying command
```

```
(ada0:ata0:0:1:0): WRITE_DMA. ACB: ca 00 ef e8 1a 40 00 00 00 00 08 00
(ada0:ata0:0:1:0): CAM status: ATA Status Error
(ada0:ata0:0:1:0): ATA status: ff (BSY DRDY DF SERV DRQ CORR IDX ERR), error: 00 ()
(ada0:ata0:0:1:0): RES: ff 00 ef e8 1a 00 00 00 00 08 00
(ada0:ata0:0:1:0): Retrying command
```

This appears to further be limited to onboard CF-to-SATA sockets. Using another disk type (mSATA, SATA, etc.) may also be a viable workaround.

34.14 Troubleshooting DNS Resolution Issues

Working DNS resolution is critical for functional access to the Internet.

34.14.1 Test connectivity

Before diagnosing DNS issues with pfSense® software specifically, start with *Troubleshooting Network Connectivity* to ensure the firewall has a proper networking configuration and working connectivity. Specifically, ensure the firewall can reach hosts on the Internet by IP address and that clients can reach the both the firewall and hosts on the Internet by IP address.

34.14.2 Check DNS service

First check which DNS service is enabled on the firewall and how it is configured.

The default configuration uses the *DNS Resolver* in resolver mode (*DNS Resolver Mode*). This mode does not require specific *DNS Servers*, it queries the root DNS servers and other authoritative servers directly (*DNS Resolution Process*).

Installations upgraded from versions before the DNS Resolver became the default may be using the *DNS Forwarder*, which requires *DNS Servers* to be entered under **System > General Setup** or to be acquired from a dynamic WAN such as DHCP or PPPoE. The DNS Resolver can also operate in this manner if set to forwarding mode.

Whichever service is active, check if it is running under **Status > Services**.

If the DNS Resolver is active, but the firewall is unable to resolve hostnames, the problem is usually a lack of working WAN connectivity. Aside from that, one possibility is that the WAN or upstream network gear does not properly pass DNS traffic in a way that is compatible with DNSSEC. Disable DNSSEC in the *DNS Resolver Configuration* to see resolution functions without DNSSEC. It is also possible that the ISP filters or rate limits DNS requests and/or requires the use of specific DNS servers. In that case, configure *DNS Servers* and then activate forwarding mode in the *DNS Resolver Configuration*.

34.14.3 Check DNS Servers

If the DNS Resolver is in forwarding mode, or the DNS Forwarder is active, then check if the firewall has DNS servers defined and ensure it can reach its DNS servers.

The firewall *DNS Server Settings* are under **System > General Setup**, and DNS servers obtained from dynamic WANs are also visible at **Status > Interfaces**. The best practice is to define at least two DNS servers. If there are multiple WANs, there should be at least one DNS server per WAN with an appropriate gateway set (*Interface and DNS Configuration*).

Perform a *Ping* test to check if the firewall can reach the DNS servers.

Note: Not all DNS servers respond to ICMP ping requests, so a failure here does not necessarily indicate a problem. Proceed to the next test to check if they respond to a DNS query.

34.14.4 Check Firewall DNS

Perform a *DNS Lookup* test to check if the firewall can resolve a hostname. The page will report the results of the query, which servers responded, and how fast they responded.

If using the DNS Resolver in resolver mode without DNS servers configured, then only `127.0.0.1` may be listed. So long as the query received the expected response, that is normal. If no response was received, ensure the DNS Resolver service is running. If it is running, disable DNSSEC and try again, or try forwarding mode. If either of those work, then the ISP may be restricting or redirecting DNS queries.

If DNS Servers are configured on the firewall or obtained from dynamic WANs, the DNS lookup page lists them and whether they responded.

If using the DNS Resolver in forwarding mode or the DNS Forwarder, the individual DNS server responses are important. If any of the servers did not respond, investigate them and potentially replace them with working servers.

If none of the servers respond, check the WAN connectivity (*Troubleshooting Network Connectivity*) and double-check the DNS server IP addresses. If the firewall can reach the gateway address at the ISP, but not the DNS servers, double-check the server IP addresses. If the DNS servers are obtained via DHCP or PPPoE and the firewall cannot reach them, contact the ISP. If all else fails, consider using a public DNS service such as Google public DNS, Quad9, or CloudFlare on the firewall instead of the DNS servers provided by the ISP. If those are already in use, The ISP may be restricting DNS queries so the only choice may be to use the ISP DNS servers.

34.14.5 Check Client DNS

If DNS works from the firewall but not from a client PC, it could be the DNS Resolver or Forwarder configuration on the firewall, the client configuration, or firewall rules.

Out of the box, the DNS Resolver handles DNS queries for clients behind the firewall. Older upgraded configurations may have the DNS Forwarder active in the same capacity.

If the client PCs are configured with DHCP, they will receive the IP address of the firewall interface to which they are connected as a DNS server, unless that is manually changed. For example, if a PC is on the LAN interface, and the firewall LAN IP address is `192.168.1.1`, then the client DNS server should also be `192.168.1.1`.

If the DNS Resolver and DNS Forwarder are disabled, adjust the DNS servers which get assigned to DHCP clients under **Services > DHCP Server**. Normally when the DNS Resolver and DNS Forwarder are disabled, the system DNS servers are assigned directly to the clients, but if that is not the case in practice for this setup, define them in the DHCP settings. If the client PC is not configured for DHCP, be sure it has the proper DNS servers set in its local configuration. This could be the LAN IP address of the firewall or an alternate set of working internal or external DNS servers.

Another possibility for DNS working from the firewall but not a local client is an overly strict firewall rule on the LAN. Check **Status > System Logs**, on the **Firewall** tab. If blocked connections appear in the log from the local client trying to reach a DNS server, then add a firewall rule at the top of the LAN rules for that interface which will allow connections to the DNS servers on **TCP and UDP** port `53`. If the client uses DNS over TLS, allow port `853` as well.

34.15 Troubleshooting the DNS Cache

34.15.1 DNS Resolver

To fully clear the DNS Resolver cache, restart the unbound daemon:

- Navigate to **Status > Services**
- Find unbound in the list

- Click  (restart) or click  (stop) then  (start)

Restarting the daemon will clear the internal cache, but client PCs may still have cached responses.

Alternately, issue a `reload` command using the CLI which will flush the cache without stopping the daemon:

```
unbound-control -c /var/unbound/unbound.conf reload
```

To selectively clear the DNS Resolver cache at the command line, run:

```
unbound-control -c /var/unbound/unbound.conf flush <name>
```

Where <name> is a domain name or hostname to be cleared.

Inspect the contents of the cache from the command line using:

```
unbound-control -c /var/unbound/unbound.conf dump_cache
```

34.15.2 DNS Forwarder

To clear the DNS Forwarder cache, restart the `dnsmasq` daemon:

- Navigate to **Status > Services**
- Find `dnsmasq` in the list
- Click  (restart) or click  (stop) then  (start)

Restarting the daemon will clear the internal cache, but client PCs may still have cached responses.

34.15.3 Client DNS Cache

The DNS cache on a Windows PC may be cleaned from a command prompt or **Start > Run**:

```
ipconfig /flushdns
```

This may need to be executed from an Administrator command prompt on Windows Vista and later.

Other operating systems will surely have other means to clear the DNS resolver cache. For example, Ubuntu-based distributions also use *dnsmasq*, and it may be restarted using:

```
sudo service network-manager restart
```

Browsers also have their own internal DNS caches separate from the OS. Close and re-open the browser if none of the above help.

As a last resort, rebooting/restarting a client will surely clear any locally cached data.

34.16 Troubleshooting DNS Queries

An administrator may need to troubleshoot issues with certain queries to the DNS Resolver (Unbound) or DNS Forwarder (dnsmasq). In such cases it can be helpful to view the queries received by the firewall and to see the responses generated.

For the DNS Resolver this can be accomplished by adding the following keyword to the **Custom Options** box on a new line:

```
server:
log-queries: yes
```

For the DNS Forwarder, add this line to the **Advanced Options** box:

```
log-queries
```

When saved, the DNS Resolver or Forwarder will begin logging the received queries and their replies, along with information about the result. The messages vary depending on the daemon. The DNS Forwarder logs whether an answer was pulled from the cache, but the DNS Resolver does not log extra data for queries answered from the cache.

Here are some examples of exchanges that might find in the query log:

A query using the DNS Resolver in forwarding mode to a system DNS server using DNS over TLS (not answered from the cache):

```
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: 192.168.1.100 daisy.ubuntu.com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] debug: validator[module 0] operate:␣
↪extstate:module_state_initial event:module_event_new
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: validator operate: query daisy.
↪ubuntu.com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] debug: iterator[module 1] operate:␣
↪extstate:module_state_initial event:module_event_pass
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: resolving daisy.ubuntu.com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: processQueryTargets: daisy.ubuntu.
↪com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: sending query: daisy.ubuntu.com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] debug: sending to target: <.> 9.9.9.9#853
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] debug: cache memory msg=16528 rrset=16528␣
↪infra=3485 val=16644
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] debug: iterator[module 1] operate:␣
↪extstate:module_wait_reply event:module_event_reply
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: iterator operate: query daisy.ubuntu.
↪com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: response for daisy.ubuntu.com. A IN
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: reply from <.> 9.9.9.9#853
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: query response was ANSWER
Oct  5 15:16:46 fw1 unbound[96103]: [96103:0] info: finishing processing for daisy.
↪ubuntu.com. A IN
```

A query to the DNS Forwarder where the response was given from the DNS cache:

```
Dec  3 08:56:46 dnsmasq[1068]: query[A] dnl-14.geo.kaspersky.com from 10.0.10.128
Dec  3 08:56:46 dnsmasq[1068]: cached dnl-14.geo.kaspersky.com is 4.28.136.39
```

A cached negative response from the DNS Forwarder:

```
Dec  3 08:56:49 dnsmasq[1068]: query[A] wpad.example.com from 192.0.2.5
Dec  3 08:56:49 dnsmasq[1068]: cached wpad.example.com is NXDOMAIN-IPv4
```

A query to the DNS Forwarder where the reply cannot be sent because of an improper client IP address (subnet ID, invalid IP address):

```
Dec  3 08:49:21 dnsmasq[1068]: query[A] teredo.ipv6.microsoft.com from 192.0.2.0
Dec  3 08:49:21 dnsmasq[1068]: forwarded teredo.ipv6.microsoft.com to 8.8.8.8
Dec  3 08:49:21 dnsmasq[1068]: forwarded teredo.ipv6.microsoft.com to 8.8.4.4
Dec  3 08:49:21 dnsmasq[1068]: reply teredo.ipv6.microsoft.com.nsatc.net is 157.56.144.
↪215
Dec  3 08:49:21 dnsmasq[1068]: failed to send packet: Permission denied
```

34.17 Troubleshooting Disk and File System Issues

pfSense® software will run a filesystem check (`fsck`) at boot when it detects an unclean UFS filesystem, typically from after a power outage or other sudden unclean reboot or shutdown. In rare cases, that isn't always enough, as a filesystem can become corrupted in other ways that may not always leave the drive marked unclean.

Note: This is not necessary for ZFS. There is no `fsck` equivalent for ZFS, and it is not prone to the issues for which UFS and other filesystem types require checks and repairs.

The command `zpool scrub <pool name>` can validate the contents of a pool, identify potential issues, and attempt to repair damage where possible. The scrub operation is not the same as `fsck`; it is not necessary in cases where `fsck` is typically needed, and it does not require a read-only filesystem so it can be run at any time.

In these cases, perform one of the following repair methods.

34.17.1 Automatic File System Check

These methods force a file system check during the boot sequence even if the drive is considered clean.

Note: This option is not present on all firewalls as it is not compatible with certain hardware. To run a manual check instead, see *Manual File System Check*.

GUI

- Navigate to **Diagnostics > Reboot**
- Set **Reboot Method** to *Reboot with File System Check*
- Click  Submit

The firewall will reboot and run the check. Monitor the console output for errors.

Console

- Connect to the console
- Choose the menu option to reboot from the console menu (`5`)
- Enter `F` (uppercase "f")

The firewall will reboot and run the check. Monitor the console output for errors.

34.17.2 Manual File System Check

If an automatic file system check is not possible, run a manual check instead:

- Reboot the firewall into single user mode as described in *Entering Single User Mode*
- Press `Enter` when prompted for a shell
- Enter `fsck -fy /`
- Repeat the command at least **five** times, or until no errors are found nor fixed, even if the file system is reported clean.
- Exit single user mode as described in *Exiting Single User Mode*.

Example:

```
/boot/kernel/kernel data=0x19e4818+0x777e8 syms=[0x4+0x9a3b0+0x4+0xdc388]
|
Hit [Enter] to boot immediately, or any other key for command prompt.

Type '?' for a list of commands, 'help' for more detailed help.
loader> boot -s
[lots of boot output]
Enter full pathname of shell or RETURN for /bin/sh:
# fsck -fy /
[...]
# fsck -fy /
[...]
# fsck -fy /
[...]
# fsck -fy /
[...]
# fsck -fy /
[...]
# reboot
```

See also:

The Netgate Resource Library contains a video which walks through the process of running a filesystem check.

34.18 Troubleshooting Full File System or Inode Errors

Error messages may appear on the console or main system log indicating that the hard drive in the firewall is full, such as:

```
/: write failed, filesystem is full
/: create/symlink failed, no inodes free
Warning: session_start(): open(/tmp/sess_XXXXXX, O_RDWR) failed: No space left on device
```

The typical cause of such errors is rarely that the drive is full, but that the operating system is unable to contact the drive. In short, the disk (HDD, SSD, CF, etc.) is dead or dying.

In cases when the drive is dying, the OS tries to write to the drive, and receives back an error code that there aren't any inodes left. Typically, other messages about the underlying problem are logged to the console but not passed back to PHP so they can be seen by the GUI. Usually those refer to things like `g_vfs_done` or *Troubleshooting DMA and LBA Errors*.

In order to eliminate the possibility that the drive is actually full, such as from a package that went crazy eating up space, try to get to a shell from the console or ssh, and run the following command:

```
: df -hi
```

The output shows disk space usage for both capacity and inodes, using human-readable numbers. The **System Information** widget on the Dashboard also shows the usage for all mounted partitions. The following output is from a system running with UFS:

```
Filesystem                     Size    Used   Avail Capacity iused ifree %iused  Mounted␣
↪on
/dev/ufsid/5ec430415d66e7cc     29G    1.9G     25G     7%     25k  4.0M    1%   /
devfs                          1.0K    1.0K      0B   100%       0     0  100%   /dev
/dev/md0                       3.4M    120K    3.0M     4%      43   979    4%   /var/run
devfs                          1.0K    1.0K      0B   100%       0     0  100%   /var/
↪dhcpd/dev
```

Note: The `devfs` lines do NOT indicate an actual problem; The `devfs` filesystem is virtual and used for housing device nodes not for files.

Of special concern is the space on `/var` and `/tmp` on systems configured to use RAM disks for those directories. Large files, such as an abnormally large DHCP leases file, can in fact fill up the `/var` memory disk and that is one way to encounter the problem.

If the root (`/`) slice has space and inodes remaining, and so do `/var` and `/tmp` and so on, then the problem is most likely a failing disk.

If disk space has been exhausted, find a way to free it up. This usually involves uninstalling or removing packages or changing the settings so they use less space.

If the disk is failing, swap it out as soon as possible.

Often when the drive is failing or full the system will continue routing packets indefinitely until it would need to access the hard drive, at which point it would quit. Users have reported firewalls running for months with a dead drive unnoticed before, though of course that is not advisable.

Another possible cause is mild filesystem corruption, which could be helped by *Troubleshooting Disk and File System Issues*.

34.19 Troubleshooting File System Capacity Shrinking

If the capacity of a ZFS dataset appears to shrink or reduce in size, the most likely reason is because the "missing" space is consumed by *ZFS Boot Environments* or ZFS snapshots created by the user manually.

ZFS Boot Environments/Snapshots consume space based on how different the current disk contents are compared to when the snapshot was taken. Rather than appear as filesystem usage like a traditional file, the apparent capacity of the filesystem is reduced instead. This allows traditional filesystem utilities to accurately account for space used by actual files vs filesystem data which is fundamentally different.

Tip: On systems with ZFS Boot Environments, use the ZFS Dashboard Widget to see the ZFS pool size as a total, rather than the size and consumption of individual filesystems.

Cleaning up older ZFS Boot Environments/snapshots will return the space they consumed and the capacity of the filesystem will increase accordingly.

See also:

Boot Environment Disk Space Usage

The ZFS Boot Environment list in the GUI displays a total amount of usage, but it is possible to see a more detailed breakdown of usage with the `zfs list -t snapshot` command in the CLI:

```
: zfs list -t snapshot
NAME                                                    USED  AVAIL  REFER  MOUNTPOINT
pfSense/ROOT/default@2023-03-27-13:48:48-0               991M      -  1.21G  -
pfSense/ROOT/default@2023-05-15-10:36:04-0               724M      -  1.29G  -
pfSense/ROOT/default@2023-05-16-09:38:51-0               715M      -  1.29G  -
pfSense/ROOT/default@2023-05-19-08:44:32-0               715M      -  1.29G  -
pfSense/ROOT/default@2023-05-22-07:59:47-0               720M      -  1.30G  -
pfSense/ROOT/default/cf@2023-02-20-12:39:12-0            2.82M     -  2.88M  -
pfSense/ROOT/default/cf@2023-03-27-13:48:48-0            2.59M     -  2.65M  -
pfSense/ROOT/default/cf@2023-05-15-10:36:04-0            1.16M     -  2.90M  -
pfSense/ROOT/default/cf@2023-05-16-09:38:51-0             216K     -  2.90M  -
pfSense/ROOT/default/cf@2023-05-19-08:44:32-0             216K     -  2.89M  -
pfSense/ROOT/default/cf@2023-05-22-07:59:47-0            1.25M     -  2.89M  -
pfSense/ROOT/default/var_cache_pkg@2023-02-20-12:39:12-0   216M     -   216M  -
pfSense/ROOT/default/var_cache_pkg@2023-03-27-13:48:48-0   174M     -   174M  -
pfSense/ROOT/default/var_cache_pkg@2023-05-15-10:36:04-0   184M     -   193M  -
pfSense/ROOT/default/var_cache_pkg@2023-05-16-09:38:51-0   184M     -   198M  -
pfSense/ROOT/default/var_cache_pkg@2023-05-19-08:44:32-0   184M     -   199M  -
pfSense/ROOT/default/var_cache_pkg@2023-05-22-07:59:47-0   185M     -   200M  -
pfSense/ROOT/default/var_db_pkg@2023-02-20-12:39:12-0     5.61M     -  6.81M  -
pfSense/ROOT/default/var_db_pkg@2023-03-27-13:48:48-0     3.25M     -  6.95M  -
pfSense/ROOT/default/var_db_pkg@2023-05-15-10:36:04-0     2.39M     -  7.04M  -
pfSense/ROOT/default/var_db_pkg@2023-05-16-09:38:51-0     2.38M     -  7.04M  -
pfSense/ROOT/default/var_db_pkg@2023-05-19-08:44:32-0     2.52M     -  7.04M  -
pfSense/ROOT/default/var_db_pkg@2023-05-22-07:59:47-0     2.89M     -  7.06M  -
```

34.20 Troubleshooting Disk Lifetime

An important part of keeping a firewall running reliably is to ensure its storage is in good condition.

If the disk in a firewall fails, it may continue to run in a reduced capacity until the system restarts. Exactly which parts may fail depends on the services and packages in use and what roles the firewall is performing. Packet filtering may continue to function indefinitely, but it may not be able to update rules, for example. Certain types of disks, such as SSD and eMMC disks, may fail into a read only state where disk writes fail or are discarded, but data can still be read.

34.20.1 Checking Disk Health & Lifetime

Contrary to popular rumors, the vast majority of modern flash storage found in SSD and eMMC drives is very resilient. Early SSDs were more prone to failure or only capable of a comparatively small number of writes, but the technology has improved vastly over the years.

Traditional spinning platter hard drives have their own problems, such as failure of mechanical moving parts, which are harder to predict.

Over the lifetime of any disk, they typically will encounter failing spots and remap them to spares. This happens on traditional HDDs as well as SSDs. It's normal to see a few these over time, but if the pool of spares is almost consumed, the disk should be replaced.

Depending on the hardware it may be possible to query the disk for information about its health. Not all platforms implement each method, and disk OEMs track data in different ways. When in doubt, contact the manufacturer of the disk for details.

SSD

Many SSDs can report their health through S.M.A.R.T. data as described in *S.M.A.R.T. Hard Disk Status* as well as perform disk tests. Disks connected using SATA, mSATA, M.2, and other similar methods can typically use the same methods as traditional HDDs, but may have SSD-specific health information in their S.M.A.R.T. data. For example, some SSDs will include a lifetime estimate or media wear indication level.

NVMe disks require special handling, but can also be queried through S.M.A.R.T., though the data they report may be in a different format from typical SSDs.

eMMC

eMMC disks are unique in that they do not implement S.M.A.R.T., but hardware which implements the correct revisions of the eMMC specification are capable of reporting health in their own way.

Install MMC Utilities

The first step is to install the `mmc-utils` package from an SSH or console shell prompt:

```
# pkg install -y mmc-utils; rehash
```

Note: This package is currently only available on pfSense® Plus software and does not have a GUI component. It must be run from an SSH or console shell prompt.

Check MMC Health Status

To check the health of the first MMC disk, run the following command:

```
# mmc extcsd read /dev/mmcsd0rpmb
```

If the disk offers health reporting, it should return output. For additional MMC disks, increase the device number in the command.

34.20.2 Interpreting MMC Health Data

The primary fields to look at in MMC health are the lifetime estimations and Pre-EOL estimation.

Note: Not all disks implement all of these fields.

```
: mmc extcsd read /dev/mmcsd0rpmb | egrep 'LIFE|EOL'
eMMC Life Time Estimation A [EXT_CSD_DEVICE_LIFE_TIME_EST_TYP_A]: 0x01
eMMC Life Time Estimation B [EXT_CSD_DEVICE_LIFE_TIME_EST_TYP_B]: 0x02
eMMC Pre EOL information [EXT_CSD_PRE_EOL_INFO]: 0x01
```

Type A
An estimate for lifetime of SLC (and pseudo-SLC) erase blocks in steps of 10%.

Type B
An estimate for lifetime of MLC erase blocks in steps of 10%.

Type A and B Values
The values of the A and B lifetime estimations are in 10% increments based on the hexadecimal value returned by the disk. This is only an estimate and the value can exceed 100%.

Possible values include:

Value	Meaning
0x00	Not defined
0x01	The disk has used 0%-10% of its estimated life time
0x02	The disk has used 10%-20% of its estimated life time
0x03	The disk has used 20%-30% of its estimated life time
0x04	The disk has used 30%-40% of its estimated life time
0x05	The disk has used 40%-50% of its estimated life time
0x06	The disk has used 50%-60% of its estimated life time
0x07	The disk has used 60%-70% of its estimated life time
0x08	The disk has used 70%-80% of its estimated life time
0x09	The disk has used 80%-90% of its estimated life time
0x0a	The disk has used 90%-100% of its estimated life time
0x0b	The disk has used 100%-110% of its estimated life time

Warning: This is only an estimation. Though useful as a general guideline, it does not necessarily indicate that a disk will fail at a given time.

Pre-EOL
Pre EOL information is an overall status for reserved blocks on the disks.

Possible values are:

Value	Severity	Meaning
0x00		Not defined.
0x01	Normal	The disk has consumed less than 80% of its reserved blocks
0x02	Warning	The disk has consumed more than 80% of its reserved blocks
0x03	Urgent	The disk has consumed more than 90% of its reserved blocks

HDD

Most HDDs can report their health through S.M.A.R.T. data as described in *S.M.A.R.T. Hard Disk Status* as well as perform disk tests. S.M.A.R.T. may need to be enabled both in the system BIOS and on the disk, though in modern systems these tend to both be enabled by default.

USB Disks

Disks connected through USB controllers are unlikely to be capable of responding to health queries, no matter what type of disk they are.

34.20.3 Taking Action

If a disk is showing signs that it might be failing, the safest action is to replace the disk. If a device has a built-in disk like an eMMC disk that cannot be replaced, it may be capable of taking an additional drive using another means such as NVMe, M.2, mSATA, or SATA. When in doubt, contact the device OEM for guidance.

If the disk is showing some wear but still has a lot of life left, consider making changes to *reduce disk writes* to potentially extend its remaining lifetime.

34.21 Troubleshooting Disk Writes

Certain tasks can make the firewall write lots of data to the disk, which could impact the health of the hardware over time. This is not as large a concern on modern disks, even SSDs, but can still be a factor over long time spans.

There are ways to reduce the amount of writes which happen on the disk, depending on the needs of the firewall and its environment.

34.21.1 RAM Disks

The `/var` and `/tmp` directories on the firewall contain most of the files which are highly volatile. The firewall has an option to keep these volatile areas in RAM disks under **System > Advanced** on the **Misc** tab.

See also:

RAM Disk Settings

Enabling RAM disks for `/var` and `/tmp` does have some caveats, which are noted in *RAM Disk Settings*. For instance, it requires sufficient RAM to hold them comfortably without filling up, and it can potentially lead to loss of logging and monitoring data if the firewall suffers a power loss.

Overall, however, if there is enough RAM to spare, using RAM disks will drastically reduce disk writes over time.

34.21.2 Disable Write-Heavy Features

One method to limit writes is by disabling features which cause lots of disk writes.

Logging
: It is possible to disable local logging, and optionally use only remote logging (*Log Settings*). This eliminates all writing of logs to the local disk. Logging is one of the primary sources of disk writes on an ongoing basis.

  This can make troubleshooting on the firewall more difficult, so it's not a best practice.

RRD Graphs
: It is possible to disable the system monitoring RRD graphs which are frequently updated with new monitoring data (*Graph Settings*).

  Instead of monitoring this data locally, most of this data can be monitored remotely by an NMS using SNMP.

DHCP Server
: On a busy network the DHCP lease database can be large and is rewritten frequently. Disabling the *DHCP server* on all interfaces and moving DHCP service to another device will result in decreased load on the firewall disk.

  This tends to be prohibitively inconvenient in most deployments, so in practice this is rare.

Note that all of these features write data in `/var` so if `/var` is in a RAM disk, they can safely remain enabled.

34.21.3 Avoid Write-Heavy Packages

Another way to reduce disk writes is to minimize use of packages that can cause heavy disk writes.

pfBlockerNG, Snort, Suricata, HAProxy
: These can write a lot due to logging and rule updates.

nmap, darkstat, other monitoring
: These use lots of disk writes to maintain databases and reports.

See also:

The package list at *Package List* also notes when specific packages require or work better with an SSD or HDD.

34.21.4 Reinstall

If it has been a while since the firewall OS was installed, reinstalling and restoring from backup can help as well.

File system properties are sometimes optimized in a new release, such as ZFS dataset layouts and attributes. Installing again will ensure the firewall is using the most optimal disk layout.

34.22 Troubleshooting Thread Errors with Hostnames in Aliases

If numerous hostnames are present in aliases, the following error may appear in the system log:

```
filterdns: Unable to create monitoring thread for host myhost.example.com! It
will not be monitored
```

This error indicated that the `filterdns` daemon hit the limit of the number of concurrent threads it could launch in order to resolve and monitor the hostnames properly.

The cure for this is to raise the number of allowed threads per process. This can be done by navigating from the pfSense® software GUI to **System > Advanced** on the **System Tunables** tab. There, create an entry for `kern.threads.max_threads_per_proc` and set it to `4096`.

After raising that value, `filterdns` must be restarted. This can be done by any method that triggers a filter reload, such as navigating to **Status > Filter Reload** and clicking **Reload Filter**.

34.23 Troubleshooting Firewall Rules

This section provides guidance for troubleshooting issues with firewall rules.

34.23.1 Check The Firewall Logs

The first step when troubleshooting suspected blocked traffic is to check the firewall logs (**Status > System Logs**, on the **Firewall** tab).

By default, pfSense® software logs all dropped traffic and will not log any passed traffic. Unless block or reject rules exist in the ruleset which do not use logging, all blocked traffic will be logged. If there are no log entries with a red  in the firewall logs which match the traffic in question, pfSense software is not likely to be dropping the traffic.

34.23.2 Check the State Table

Attempt a connection and immediately check the state table at **Diagnostics > States** and filter on the source or destination to see if a state exists. If a state table entry is present, the firewall has passed the traffic.

If the rule in question is a pass rule, the state table entry means that the firewall passed the traffic through and the problem may be elsewhere and not on the firewall.

If the rule is a block rule and there is a state table entry, the open connection will **not** be cut off. To see an immediate effect from a new block rule, the states must be reset. See *Firewall States* for more information.

34.23.3 Review Rule Parameters

Edit the rule in question and review the parameters for each field. For TCP and UDP traffic, remember the source port is almost never the same as the destination port, and should usually be set to *any*.

If the default deny rule is to blame, craft a new pass rule that will match the traffic to be allowed. If the traffic is still blocked, there may be some other special aspect of the packets which require additional handling in the rule configuration. For example, certain multicast traffic may need to have **Allow IP Options** enabled, or the log entries may be due to asymmetric routing, or the packets may have an invalid combination of parameters such as a fragmented packet with "Don't Fragment" set inside.

See also:

- *Bypass Firewall Rules for Traffic on Same Interface*
- *Static Route Filtering*
- *Troubleshooting Asymmetric Routing*
- *Troubleshooting Blocked Log Entries for Legitimate Connection Packets*

In such advanced cases, running a packet capture for the traffic in question can help diagnose the problem. Refer to *Packet Capturing* for more information on how to capture and analyze packets.

Protocol

The protocol to which the rule will apply must be specified. Most often, this is TCP, UDP, or ICMP, but other protocols such as ESP, AH, and GRE are regularly encountered when dealing with VPNs.

Confusion arises when a firewall administrator is unsure of what protocol to use. A rule set with TCP may not work because the application being filtered may actually use UDP instead. When in doubt, try using TCP/UDP.

34.23.4 NAT Confusion

When crafting rules for firewalls involving inbound NAT connections, remember to use the **private IP address** as the Destination. This applies for port forwards as well as 1:1 NAT

34.23.5 Port Forward *pass* action

When creating a port forward, the *pass* action will bypass firewall rules and pass the traffic directly through without filtering. Change the setting to create an associated rule and then arrange the block rule above the resulting pass rule.

34.23.6 Source and Destination Ports

When crafting rules, bear in mind that typically only a source *or* a destination port needs to be specified, and rarely both. In the majority of cases, the source port does not matter at all. For example, to allow ssh access to the firewall, only specify a **destination** port of `22`. The source port of the client will be random.

34.23.7 Review Rule Ordering

Firewall rules are generally processed as follows:

- Floating Rules
- Interface Group rules
- Interface tab rules

See also:

See *Ordering of NAT and Firewall Processing* for more details.

If a floating rule with **quick** checked passed the traffic, then a block rule on an interface would have no chance to match the traffic.

34.23.8 Rules and Interfaces

Ensure rules are on the correct interface to function as intended. Traffic is filtered only by the ruleset configured on the interface where the traffic is *initiated*. Traffic coming from a system on the LAN destined for a system on any other interface is filtered by **only** the LAN rules. The same is true for all other interfaces.

34.23.9 Enable Rule Logging

Determine which rule is matching the traffic in question. The hit counters in the rule list can help with this to some degree. By enabling logging on pass rules, the firewall logs will show an individual entry specifically to determine which rule passed the connection.

34.23.10 Troubleshooting with packet captures

Packet captures can be invaluable for troubleshooting and debugging traffic issues. With a packet capture, it is easy to tell if the traffic is reaching the outside interface or leaving an inside interface, among many other uses. See *Packet Capturing* for more details on troubleshooting with packet captures.

34.23.11 New Rules Are Not Applied

If a new rule does not appear to apply, there are a couple possible explanations.

First, If the rule is a block rule and there is a state table entry, the open connection will **not** be cut off. See *Check the State Table*.

Second, the ruleset may not be reloading properly. Check **Status > Filter Reload** to see if an error is displayed. Click the  **Reload Filter** button on that page to force a new filter reload. If an error is displayed, resolve the problem as needed.

If none of the above causes are to blame, it's possible that the rule is not matching at all, so review the traffic and the rule again.

34.23.12 Unfilterable Traffic

Certain traffic cannot be filtered. Not because the pfSense software isn't capable, but because they actually do not touch the firewall at all. A prime example of this is trying to keep one device on the LAN from accessing another device on the same LAN. This is not possible if both clients are on the same subnet and switch; In that case, the routing of packets is handled at the switch level (layer 2), and the firewall has no knowledge of the traffic. If there is a need to control access in this way, the devices in question must be on separate firewall interfaces. When on different "legs" of the network, their traffic will route through the firewall, the firewall will have full control of the flow.

34.23.13 UPnP IGD & PCP passed traffic

If *UPnP IGD & PCP* is enabled and a LAN device opens a port to the world, the traffic may still get in even if it appears it should otherwise be blocked.

34.23.14 Asymmetric Routing

If reply traffic such as `TCP:A`, `TCP:SA`, or `TCP:RA` is shown as blocked in the logs, the problem could be asymmetric routing. See *Troubleshooting Asymmetric Routing* for more info.

34.23.15 Ruleset Failing to Load

It is also possible that the rules are not being loaded properly. Typically, this would result in a notification in the GUI, however manual tests can be performed to check.

From the GUI, visit **Status > Filter Reload**. Click  **Reload Filter** wait for the process to stop, then scroll to the bottom of the page to see if the last line says `Done.` or if it stops. If it stops, for example in a particular package, then there may be a problem with that package.

The ruleset can also be verified from the console or **Diagnostics > Command** in the **Shell Execute** box by running:

```
pfctl -f /tmp/rules.debug
```

If an error is displayed, it may have an obvious fix, or search for that error to find possible resolutions.

34.23.16 Other Causes

There are other pitfalls in firewall rules, NAT, routing, and network design that can interfere with connectivity. See *Troubleshooting Network Connectivity* for more suggestions.

See also:

Hangouts Archive to view the June 2016 hangout on Connectivity Troubleshooting which contains much more detailed troubleshooting procedures.

34.24 Troubleshooting Bogon Network List Updates

Make sure the firewall can resolve DNS host names and can reach the bogons host, otherwise the update will fail.

To ensure the firewall can resolve the bogon update host via DNS, perform a *DNS Lookup*:

- Navigate to **Diagnostics > DNS Lookup**
- Enter `files.pfsense.org` in the **Hostname** field
- Click **Lookup**

If that fails, *troubleshoot DNS resolution* for the firewall itself.

If that works, then perform a *port test* as demonstrated in Figure *Testing Connectivity for Bogon Updates*:

- Navigate to **Diagnostics > Test Port**
- Enter `files.pfsense.org` in the **Hostname** field
- Enter `80` in the **Port** field
- Click **Test**

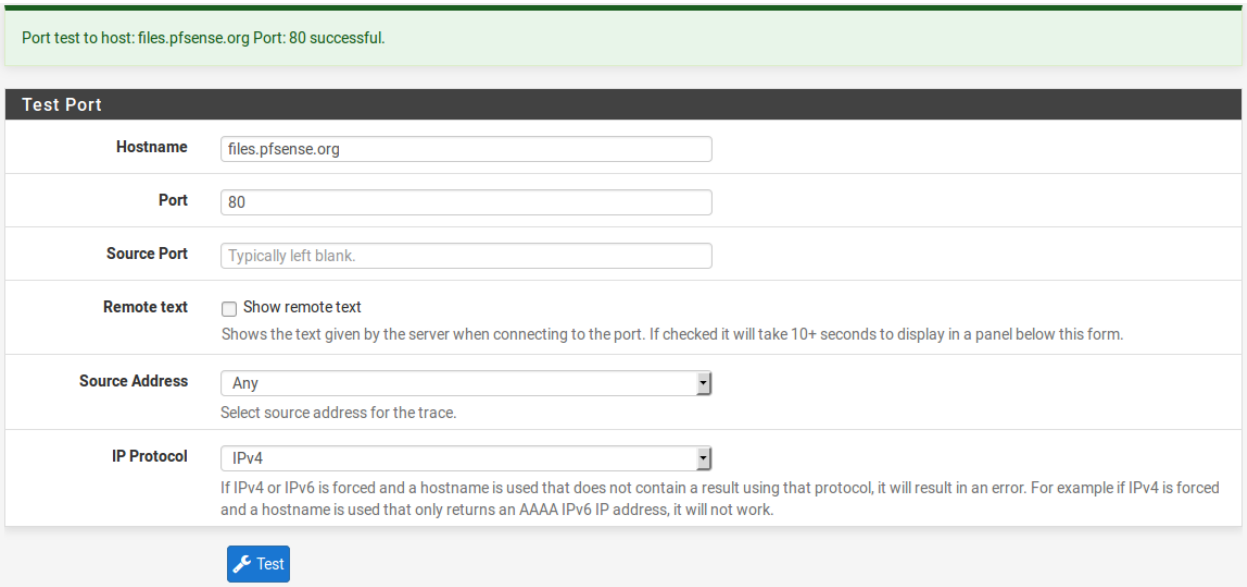


Fig. 2: Testing Connectivity for Bogon Updates

If that fails, *troubleshoot connectivity from the firewall*.

34.24.1 Forcing a Bogon Network List Update

With the relatively infrequent changes to the bogons list, and advance notice of new public IP assignments, a monthly bogons update is adequate. However, there may be scenarios where a manual bogon update can help, such as if the bogon updates have been failing because of an incorrect DNS configuration. Execute an update via the GUI:

- Navigate to **Diagnostics > Tables**
- Select *bogons* or *bogonsv6* from the **Table** list
- Click  **Update**

34.25 Troubleshooting FTP Connections

pfSense® software does not have a built-in FTP proxy, but one is available as an add-on package.

See also:

- *Using NAT and FTP without a Proxy*

34.25.1 Rules to allow FTP

If problems are encountered with FTP, check the rules to/from FTP devices, ensure that both the control port and PASV range are allowed.

34.25.2 Troubleshooting/Alternatives

- Disable the FTP Proxy and attempt the connection again
- Use SCP/SFTP which only needs one port to traverse the firewall since it is wrapped in SSH (a safe AND simple way of traversing a firewall!)
- Don't use FTP (highly recommended option)

34.26 Troubleshooting Gateway Monitoring

In some cases, the dpinger gateway monitoring daemon will output numeric error codes in the Gateways log indicating a problem reaching the monitored target IP address. The errors on this page are the most common.

34.26.1 sendto error: 55

```
55 ENOBUFS
No buffer space available.
An operation on a socket or pipe was not performed because the system lacked sufficient␣
↪buffer space or because a queue was full.
```

Several possible conditions can cause this. For a list of possible causes and solutions, see *Troubleshooting "No buffer space available" Errors*.

34.26.2 sendto error: 64

```
64 EHOSTDOWN
Host is down.
A socket operation failed because the destination host was down.
```

In this case, the firewall is unable to reach a target host directly connected at layer 2 (No ARP response), or it received a similar error response from an upstream source. Generally this only happens due to remote problems, indicating that the target is actually down or the L1/L2 link to the target is down.

34.26.3 sendto error: 65

```
65 EHOSTUNREACH
No route to host.
A socket operation was attempted to an unreachable host.
```

Either there is no possible route to the target locally, or status information was received from an upstream router that indicated the same condition elsewhere along the path to the target.

This can happen due to a lack of default route, missing interface link route, or similar conditions.

34.27 Troubleshooting High Availability DHCP Failover

There are several potential scenarios that can cause problems with DHCP service for a high availability cluster. This document contains items to check as well as potentially problematic scenarios and workarounds.

The issues and limitations encountered here vary depending on which DHCP daemon backend is active.

34.27.1 Common Issues

These issues may affect high availability DHCP failover no matter which backend is active.

Time Not Synchronized

The system time on both cluster nodes must be within 90 seconds of each other. Otherwise, the time difference is too large and the DHCP daemon processes will not communicate.

34.27.2 Kea DHCP Daemon

These issues may affect high availability DHCP failover when the Kea DHCP backend is active.

Kea High Availability Node Status

Check the DHCP Leases or DHCPv6 Leases to see the current status of failover for each daemon. The High Availability status for Kea DHCP operates identically for both DHCP and DHCPv6. For details on how the DHCP HA status operates, see *High Availability Status – Kea DHCP Only*.

If any of the peers are in a state other than `hot-standby` it may indicate a problem. The solution is likely in one of the other answers in this document.

Incorrect Failover Peer Address

The **Settings** tab for DHCP and DHCPv6 service defines the failover peers for HA in Kea.

The address in this field **must** be an actual **interface IP address** on the peer – **not** a CARP VIP, other shared address. Most often this is the IP address of the Sync interface on each node.

Note: The address family does not need to match the family of the DHCP daemon, for example, DHCPv6 HA can be performed using IPv4 peer addresses.

When XMLRPC synchronization is enabled the primary node will adjust this automatically when copying settings to the secondary node, swapping the local and remote values appropriately.

Failover Peer Unreachable

If the peers are both active but cannot communicate with each other, they may both consider the peer to be in an unreachable state. If this happens, both nodes may hand out addresses at the same time, causing a conflict. Double-check the IP addresses and ports used for failover as well as the firewall rules on the interfaces involved.

Firewall Rules

For two DHCP peers to exchange failover data they must be able to reach each other on the configured ports through the interfaces containing the configured local and remote addresses. These addresses are typically on the Sync interface.

It may be necessary to add firewall rules to pass the sync traffic, for example when using strict firewall rules on the Sync interface.

The primary node must be allowed to reach the secondary node on the defined port (default `8765` for DHCPv4 and `8766` for IPv6) using the **Remote Address** defined in the DHCP settings, and vice versa. Check the firewall rules on the Sync interface and add rules as needed to pass this traffic.

Secondary Does Not Enter `partner-down` State

In certain cases with customized **Advanced** settings, a secondary node may not enter the `partner-down` state even when the primary is unreachable. Usually this happens because of the **Max Unacked Clients** setting, which instructs the daemon to wait for that many clients to send unanswered DHCP requests before it will take over.

While this setting can help prevent a secondary node from taking over too quickly in certain cases, a high value for **Max Unacked Clients** will prevent that number (minus 1) clients from getting addresses until another client also requests an address and goes unanswered. On a large and busy network this may not take long, but on a small or quiet network this could take a significant amount of time and leave some clients stranded. If too many clients are going unanswered for too long, lower the value until clients are served in a timely manner during a failure of the primary node.

The default value is `0` which causes the secondary to take over immediately after the primary has been unresponsive (default time: 60 seconds).

Mismatched TLS Settings

The TLS settings for Kea DHCP and DHCPv6 HA do not synchronize via XMLRPC as they may be different on each node. If the TLS settings are not configured in appropriate alignment for both nodes, then the two nodes cannot communicate, and DHCP HA will fail.

When using TLS transport, ensure each node has the **Server certificate** option set to use a **Server** type certificate signed by a CA trusted by both peers (e.g. both using the same CA).

When using mutual TLS, ensure each node is using an appropriate **User** certificate signed by the same CA in the **Client certificate** option.

34.27.3 ISC DHCP Daemon

These issues may affect high availability DHCP failover when the ISC DHCP backend is active.

Interface Order Mismatch

The interfaces must be assigned identically on both nodes, for example: wan=WAN, lan=LAN, opt1=Sync, opt2=DMZ. Check the `config.xml` contents directly to ensure a match. If the interfaces are not assigned in the same order, the automatically generated failover pool names will not match, which prevents DHCP failover from working.

Pool Status

Look at the pool status section at **Status > DHCP leases**. All defined pools (often 1 per interface) are listed at the top of that page.

If any of the pools are in a state other than "normal", then debug the problem further. The solution is most likely found in one of the other items in this document.

If the pool status on both nodes is "normal" for a given interface, then issues with clients obtaining leases on that interface is not likely a problem in the DHCP configuration or failover, but elsewhere. For example, an L2 or switch problem, client problem, etc.

Incorrect Failover Peer Address

Each interface tab in the DHCP server options has a separate field for the Failover Peer IP address. This field must be filled in for each interface participating in DHCP with HA. The address in this field **must** be the actual **interface IP address** on the peer corresponding to the chosen tab – **not** a CARP VIP, other shared address, or an address from an unrelated interface (e.g. The SYNC interface).

When XMLRPC synchronization is enabled the primary node will adjust this automatically when copying settings to the secondary node, filling in its own IP address for the interface.

Failover Peer Unreachable

If one failover peer cannot contact the other peer when it starts up, it will stop itself from handing out leases intentionally. It does this as a fail-safe to prevent itself from handing out conflicting lease data.

This can happen if, for example, both nodes suffer a power loss and only one recovers. Another common scenario is if one node suffers a hardware failure and the working node must be rebooted before the failed node can be repaired.

Correcting this can be tricky. The simplest way to correct it is to bring the other peer online if possible. If that is not possible, then the only way may be to remove the failover peer IP addresses from each DHCP interface configuration so the node no longer believes it should be part of a failover pool. When the other node recovers, the configuration can be put back in place.

Firewall Rules

For two DHCP peers to exchange failover data, they must be able to reach each other on an interface. There are typically automatic firewall rules which handle this, but there have been issues in the past where these automatic rules did not cover every possible scenario.

If the firewall log shows this traffic being blocked, then it may be necessary to add manual rules to pass the traffic. Ensure the two nodes are allowed to communicate on every relevant interface. The primary node must be allowed to receive traffic on TCP port 519 from the secondary node and the secondary node must be allowed to receive traffic on TCP port 520 from the primary node.

Restart DHCP Daemons

Stop and restart the DHCP daemon from **Status > Services** on both nodes and check the status after a few moments. This may correct the issue or at least provide better detail in the logs during the startup procedure for the daemons.

Check CARP VIP Configuration

Check the CARP VIP configuration for VIPs on interfaces used for DHCP failover. The primary node must have an **Advertising Frequency Skew** value *below* `20`, the secondary node must have an **Advertising Frequency Skew** value *above* `20`.

Mismatched Versions

Both nodes must be running the same version of pfSense® software. Update both nodes to the newest available stable release if they do not match. Older versions may have problems with various aspects of DHCP failover that have already been corrected.

Reset Lease Database

If the two nodes cannot agree on the pool status, it may be due to the contents of the lease databases. This can sometimes happen when first setting up failover or after reinstalling an HA node without backing up and restoring its DHCP lease database.

If all else fails, perform the following:

- Stop the DHCP daemon on both nodes
- Remove the DHCP lease database files from `/var/dhcpd/var/db/dhcpd.leases*` on both nodes
- Start the DHCP daemon on both nodes

It's important to perform each step on both nodes individually and not do the whole procedure at once on each node; the important part is that they both start at the same time with a new empty lease database.

Inconsistent Client Hostnames

The DHCP servers on each node in a failover configuration work in coordination with one another. Each server will handle a portion of the total pool and relay lease information to the failover peer.

The lease information the nodes exchange, however, does not include client hostnames. This means that features which rely on DHCP hostnames, such as DNS resolution of DHCP client hostnames, will not work consistently when using DHCP failover.

This is a limitation of the ISC DHCP daemon and not something that can be changed or corrected in pfSense software.

The most viable workaround is to define static DHCP mappings for each host that must be resolved via DNS.

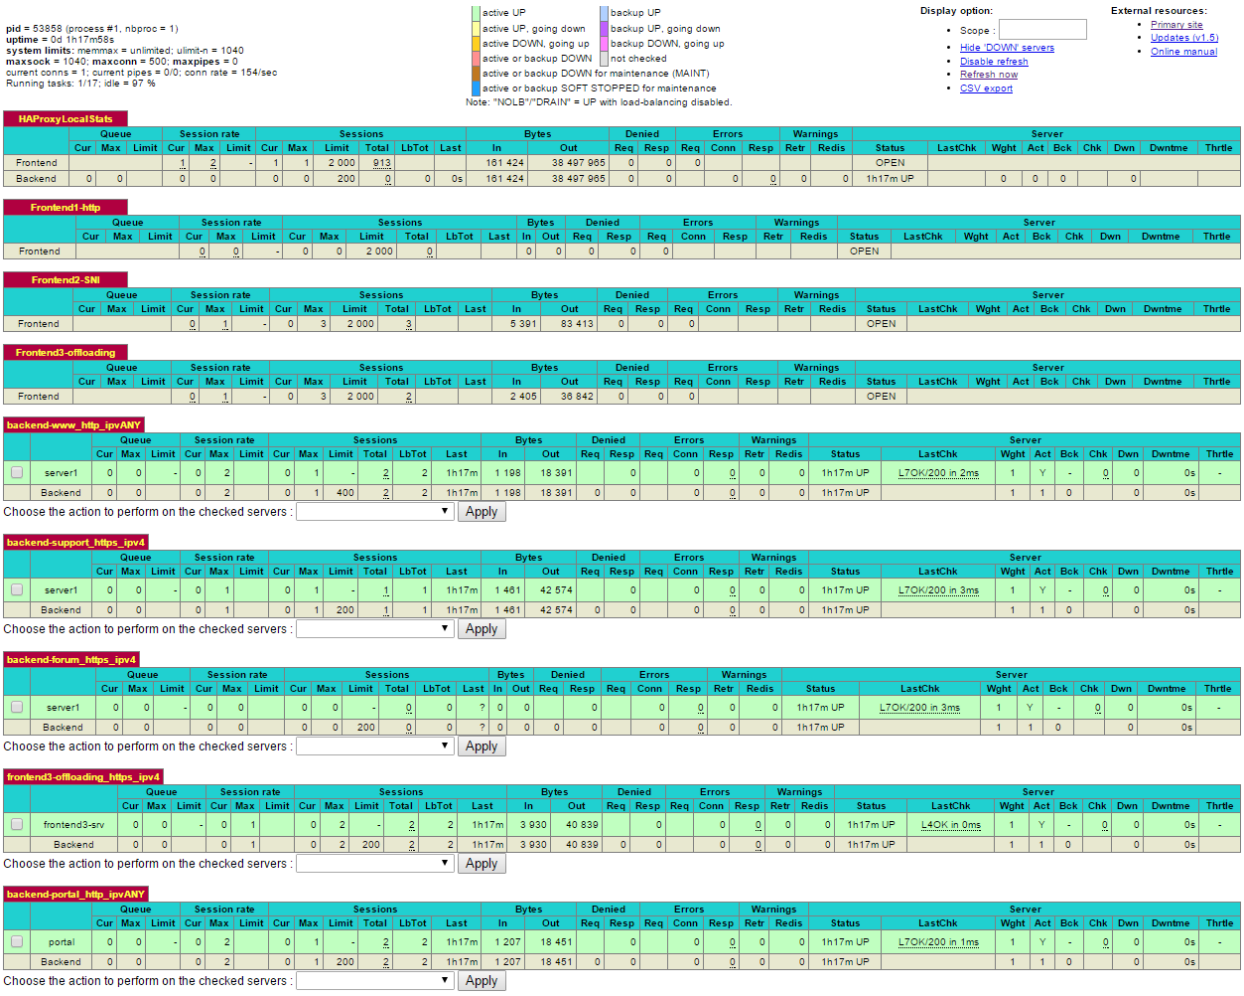
34.28 Troubleshooting the HAProxy Package

Troubleshooting steps for *HAProxy package*.

For troubleshooting there are 2 parts are helpful, depending on the issue:

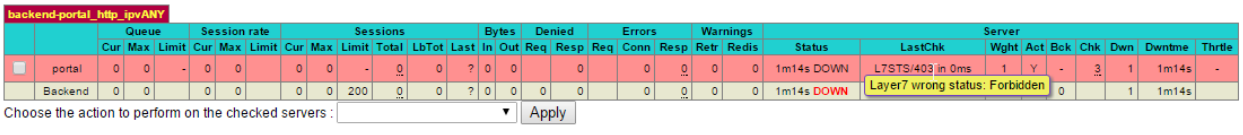
- Stats page
- Syslog logging

34.28.1 Stats

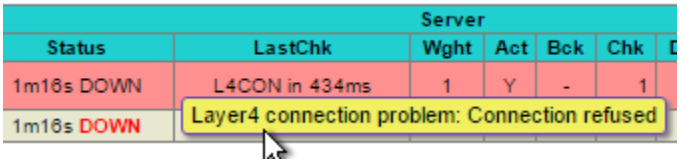


If health checks have been configured on the servers, the backend will show what servers are up or down. Layer 7 checks provide the most information about this, but a layer 6 or 4 check can also be useful.

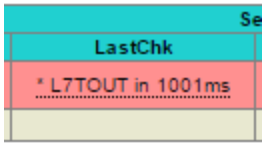
If a server is shown in red like here, hover over the check result for a second. It will tell what the short error code means in a more readable description:



There are different error codes that ask for different resolution.



or



A layer 4 issue might indicate that a wrong server IP address or port was filled in, or that the server is not running / accepting connections. A firewall on the server itself, or a missing route could all cause these kinds of issues. A layer 6 issue indicates a problem with the SSL certificates. A layer 7 issue would generally be due to an unexpected or no status returned by the web server, the web server might take to long to present the checked URL. Or the web server does not allow the configured check method/options.

For configuring the health check, the following options should be accepted by most web server, but are more difficult to filter out from normal web server logs:

```
option httpchk GET / HTTP/1.1\r\nHost:\ www.yourdomain.com\r\nAccept:\ */*
```

If the backend requires authentication, another option could be to change the URL to a different page that does not need authentication, perhaps specifically added to the web server for this purpose:

```
option httpchk GET /healthcheck.php
```

Or accept that authentication error as the 'valid' result:

```
http-check expect status 401
```

34.28.2 Syslog

If all backend servers are 'up' in the stats, but 'sometimes' users are reporting problems, then logging is important to configure and collect.

HAProxy allows for configuring syslog server destination on the settings tab. The actual logging to files must be configured on that destination syslog server. The default log format is rather detailed if configured for the appropriate format. As such the 'Detailed logging' option in the frontend edit page should be checked. For mode HTTP servers the following will be configured in the config file:

```
option httplog
```

For mode TCP servers the following will be configured in the config file:

```
option tcplog
```

If needed, it's also possible to capture additional traffic headers which will be added into the syslog messages.

More information about these options can be found in the official documentation: Official HAProxy manual

34.29 Troubleshooting VPN Connectivity to a High Availability Secondary Node

If there is a VPN connection to a High Availability cluster (site-to-site or remote access/mobile), often remote devices can communicate with the active node, but the backup node is unreachable.

The reason for this is that the VPN is configured and active on both firewalls. The packet from the client goes to the primary over the VPN tunnel that is up and connected, then goes from the primary to the secondary, and the secondary attempts to send it back over its own copy of the VPN which is down because it is the backup. The response never makes it back to the original client.

To address this situation, configure Hybrid or Manual outbound NAT and add rules such that the firewall performs NAT on the traffic from the VPN subnet going to IP addresses on the secondary node, and vice versa. That way these connections appear to originate from the opposing firewall and not the VPN, so the traffic returns as expected. Hybrid or Manual Outbound NAT is likely already enabled set since it is typically a requirement for HA using CARP VIPs.

For example, set the mode to **Hybrid Outbound NAT** and add an outbound NAT rule on the LAN interface. Configure the rule with the source being the VPN subnet, destination being an alias containing both the primary and secondary node LAN IP addresses. Translation would be **Interface Address** (NOT the CARP VIP!).

With the NAT rule present, when attempting to access the opposing node over the VPN the traffic will appear to originate from the node to which the VPN is currently connected, and the return traffic will go back as expected.

34.30 Troubleshooting High Availability

High availability configurations can be complex, and with numerous different ways to configure a failover cluster, it can be tricky to get things working properly. This section discusses common problems and solutions for the majority of cases. If issues are still present after consulting this section, there is a dedicated HA/CARP/VIPs board on the Netgate Forum.

Before proceeding, take the time to check all members of the HA cluster to ensure that they have consistent configurations. Often, it helps to walk through the example setup, double-checking all the proper settings. Repeat the process on the secondary node, and watch for any places where the configuration must be different on the secondary. Be sure to check the CARP status (*Check CARP status*) and ensure CARP is enabled on all cluster members.

Errors relating to HA will be logged in **Status > System Logs**, on the **System** tab. Check those logs on each system involved to see if there are any messages relating to XMLRPC sync, CARP state transitions, or other related errors.

See also:

The issues on this page are for HA in general. For issues specific to using HA in virtual environments, see *Troubleshooting High Availability Clusters in Virtual Environments*

34.30.1 Common Misconfigurations

There are several common misconfigurations that happen which prevent HA from working properly.

Incorrect Interface Order

As mentioned on *pfSense Software XMLRPC Config Sync Overview*, the interface assignment order and internal identifiers must match identically on both nodes.

If the interface order does not match, the configuration synchronization process will copy rules and other settings such as DHCP failover to different/unexpected interfaces on the secondary node.

Use a different VHID on each CARP VIP

Every CARP VIP on a given interface or broadcast domain must use a different VHID. The VHID determines the virtual MAC address used by a CARP IP address, thus different clusters attempting to use the same VHID on the same L2 segment cause a MAC address conflict.

Within a single HA pair, input validation prevents configuring duplicate VHIDs. Unfortunately, it isn't always that simple. CARP is a multicast technology, and as such anything using CARP on the same network segment must use a unique VHID, even if it is a different subnet. Other protocols such as VRRP and HSRP also use protocols similar enough to CARP that they can conflict, so ensure there are no other devices using similar VHIDs with other protocols, such as if the ISP or another router on the local network is using VRRP.

The best way around this is to use a unique set of VHIDs. If a known-safe private network is in use, start numbering at 1. On a network where VRRP or CARP are conflicting, consult with the administrator of that network to find a free block of VHIDs. Performing a packet capture on the segment and analyzing it with tools such as Wireshark may also reveal other hosts sending similar multicast messages with IDs to avoid.

Incorrect CARP VIP Settings

Inspect the settings for CARP VIPs (**Firewall > Virtual IPs**) to ensure they are correct and consistent on both nodes.

The **Advertising Frequency** values must be appropriate for each VIP and node:

Base
: Values should be the same on both nodes. In some situations where the secondary node is on a slow or non-local link, users have increased this value on only the secondary, but that can lead to problems with each node assuming their expected roles at the proper times.

Skew
: Values must be different on the primary and secondary nodes. The primary is typically `1` or `0`, and the secondary is typically `100`.

Incorrect Times

Check that all nodes involved are properly synchronizing their clocks using NTP and have valid time zones, especially if running in a Virtual Machine. If the clocks are too far apart, some synchronization tasks like DHCP failover will not work properly.

Incorrect Subnet Mask

The real subnet mask **must** be used for a CARP VIP, **not** `/32`. This **must** match the subnet mask for the IP address on the interface to which the CARP VIP is assigned.

Both Nodes in Maintenance Mode

If both nodes have activated **Persistent CARP Maintenance Mode** at **Status > CARP (failover)**, they each will advertise a skew of `254` and the actual status will be unpredictable. Ensure only one node is in maintenance mode at a time.

34.30.2 Incorrect Hash Error

There are a few reasons why this error turns up in the system logs, some more worrisome than others.

If CARP is not working properly when this error is present, it could be due to a configuration mismatch. Ensure that for a given VIP, that the VHID, password, and IP address/subnet mask all match.

If the settings appear to be proper and CARP still does not work while generating this error message, then there may be multiple CARP instances on the same broadcast domain. Disable CARP and monitor the network via packet capturing to check for other CARP or CARP-like traffic (*Packet Capturing*), then adjust VHIDs appropriately.

If CARP is working properly and this message is in the logs when the node boots up, it may be disregarded. It is normal for this message to be seen during boot, as long as CARP continues to function properly (primary shows **MASTER** status, secondary shows **BACKUP** status).

34.30.3 Both Nodes Appear as MASTER

This will happen if the secondary node cannot see the CARP heartbeat advertisements from the primary. Check the firewall rules, connectivity between the nodes on that segment, and switch configurations. Also check the system logs for any relevant errors that may lead to a solution. If this is encountered in a Virtual Machine (VM) hypervisor environment such as VMware ESXi, see *Troubleshooting High Availability Clusters in Virtual Environments*.

34.30.4 Primary Node is Stuck as BACKUP

In some cases this may happen normally for a short period after a node comes back online. However, certain hardware failures or other error conditions can cause a server to silently take on a high `advskew` of `240` in order to signal that it still has a problem and should not become master. This can check be checked from the GUI, or via the shell or **Diagnostics > Command**.

In the GUI, this condition is printed in an error message on **Status > CARP**.

From the shell or **Diagnostics > Command**, run the following command to check for a demotion:

```
# sysctl net.inet.carp.demotion
net.inet.carp.demotion: 240
```

If the value is greater than `0`, the node has demoted itself.

In that case, isolate the firewall, check its network connections, and perform further hardware testing.

If the demotion value is `0` and the primary node still appears to be demoting itself to `BACKUP` status, or it is flapping, check the network to ensure there are no layer 2 loops. If the firewall receives its own heartbeats back from the switch, it can also trigger a change to `BACKUP` status.

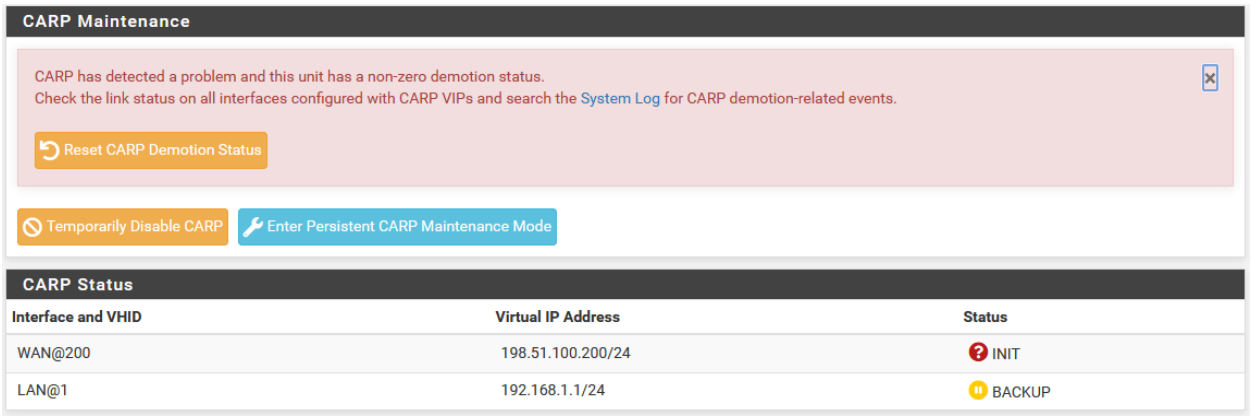


Fig. 3: CARP Status when Primary is demoted

34.30.5 Other Switch and Layer 2 Issues

- If the nodes are plugged into separate switches, ensure that the switches are properly trunking and passing broadcast/multicast traffic.
- Some switches have broadcast/multicast filtering, limiting, or "storm control" features that can break CARP.
- Some switches have broken firmware that can cause features like IGMP Snooping to interfere with CARP.
- If a switch on the back of a modem/CPE is use, try a dedicated switch instead. Switches built into routers and other similar CPE devices often do not properly handle CARP traffic. In these cases, plugging the firewalls into a proper switch and then uplinking to the CPE will eliminate problems.

34.30.6 Configuration Synchronization Problems

Double-check the following items when problems with configuration synchronization are encountered:

- The XMLRPC synchronization user must be configured properly in the user manager. The account must have the "System - HA node sync" privilege. If users are synchronized, the account must be added on both nodes initially, once the first synchronization happens, the primary will copy its entry the secondary.
- The password in the configuration synchronization settings on the primary node must match the synchronization user password on the secondary node.
- The GUI must be on the same port on all nodes.
- The GUI must be using the same protocol (HTTPS or HTTP) on all nodes.
- Traffic must be permitted to the GUI port on the interface which handles XMLRPC synchronization traffic.
- Verify that only the primary node has configuration synchronization options enabled.
- Ensure no IP address is specified in the **Synchronize Config to IP** on the secondary node.
- Ensure the clocks on both nodes are current and are reasonably accurate.

34.30.7 State Synchronization Problems (pfsync)

If the State Creator Host IDs do not line up under **Status > CARP** in the **State Synchronization Status** section, that can indicate that the states have not been synchronized. The status should include the **Filter Host ID** of both nodes if states are synchronizing correctly. If the filter host ID has been changed recently, additional values may be in the list until the older states expire.

- Ensure that **Synchronize States** is enabled on **both** nodes.
- Ensure both nodes have the correct **Synchronize interface** selected.
- Ensure the two nodes can communicate directly on the chosen synchronize interface (e.g. Verify with `ping` that they can both reach each other.)
- Check the firewall logs for blocked traffic using the `pfsync` protocol. If the traffic is blocked, make sure it is present on the correct interface. If the interface is correct, then adjust the firewall rules to allow the traffic to pass.
- Ensure the interface assignment order matches.
- If state synchronization does not work with **Synchronize Peer IP** left empty, fill in the Sync interface IP address of each peer on both nodes.

34.30.8 HA and Multi-WAN Troubleshooting

If clients have trouble reaching CARP VIPs when using with Multi-WAN, double check that a rule is present like the one mentioned in *Firewall Configuration*.

34.31 Troubleshooting High Availability Clusters in Virtual Environments

34.31.1 Hypervisor users (Especially VMware ESX/ESXi)

The below settings are specifically for VMware ESX/ESXi, but similar settings may be present on Hyper-V, VirtualBox, and other similar hypervisors.

Note: These notes all apply to CARP VIPs in multicast mode. Unicast mode CARP on pfSense Plus software may not require these settings, but experiences may vary by hypervisor and environment.

- Enable promiscuous mode on the vSwitch
- Enable **MAC Address changes**
- Enable **Forged transmits**
- If multiple physical ports exist on the same vSwitch, the `Net.ReversePathFwdCheckPromisc` option must be enabled to work around a vSwitch bug where multicast traffic will loop back to the host, causing CARP to not function with "link states coalesced" messages. (See below)

ESX VDS Promisc Workaround

If a Virtual Distributed Switch is in use, a port group can be made for the firewall interfaces with promiscuous mode enabled, and a separate non-promiscuous port group may be used for other hosts. This has been reported to work by users on the forum as a way to strike a balance between the requirements for letting CARP function and for securing client ports.

ESX VDS Upgrade Issue

If a VDS (Virtual Distributed Switches) is used in ESX 4.0 or 4.1 and an **upgrade** from 4.0 to 4.1 or 5.0 is performed, the VDS will not properly pass CARP traffic. If a new VDS is **created** on 4.1 or 5.0, it will work, but the upgraded VDS will not.

It is reported that disabling promiscuous mode on the VDS and then re-enabling it will resolve the issue.

ESX VDS Port Mirroring Issue

If port mirroring is enabled on a VDS, it will break promiscuous mode. To fix it, disable promiscuous mode, then re-enable promiscuous mode.

Client Port Issues

If a bare metal HA cluster is connected to a switch with an ESX host using multiple ports on the ESX host (lagg group or similar), and only certain devices or IP addresses are reachable by the target VM, then the port group settings in ESX may need adjusted to set the load balancing for the group to hash based on IP address, not the originating interface.

Side effects of having that set incorrectly include:

- Traffic only reaching the target VM in promiscuous mode on its NIC
- Inability to reach the CARP VIP from the target VM when the "real" IP address of the primary firewall is reachable
- Port forwards or other inbound connections to the target VM work from some IP addresses and not others.

Changing Net.ReversePathFwdCheckPromisc

Login to the VMware vSphere Client

For each VMware host

- Click on host to configure and select the **Configuration** Tab
- Click **Software Advanced Settings** in left pane
- Click on **Net** and scroll down to **Net.ReversePathFwdCheckPromisc** and set to 1
- Click **OK**

Promiscuous Mode interfaces need to be set now or toggled off and then back on. This is done per host by clicking **Networking** in the **Hardware** section

- For each vSwitch and/or Virtual Machine Port Group:

> **Note:** If Promiscuous is already enabled it must be disabled, saved and then re-enabled and saved again.

  - Click on **Properties** of the vSwtich

    By Default Promiscuous Mode is **Reject**.
  - Click the **Edit > Security** Tab
  - Select **Accept** from the drop-down
  - Click **OK**
- However, this setting is usually applied per Virtual Machine Port Group (More Secure) where the vSwitch is left at default to Reject.
  - Navigate to **Edit > Security > Policy Exceptions**
  - Uncheck Promiscuous Mode
  - Click **OK**
  - Navigate to **Edit > Security > Policy Exceptions**
  - Check Promiscuous Mode and select **Accept**.

More information available from VMware

ESX Physical NIC Failure Fails to Trigger Failover

Self-demotion of a CARP VIP relies on the loss of link on a switch port. As such, if a primary and secondary node instance are on separate ESX host and the primary ESX host loses a switch port link and does not expose that to the VM, CARP will stay MASTER on all of its VIPs and the secondary will also believe it should be MASTER. One way around this is to script an event in ESX that will take down the switch port on the VM if the physical port loses link. There may be other ways around this in ESX as well.

VMware Workstation

If using VMware workstation on Linux for testing/modeling and CARP failover does not function, it is likely because VMware workstation is running non-root and cannot set the `vmnet` adapter in Promiscuous mode.

The permissions on `/dev/vmnet*` should be changed such that the user running VMware workstation is allowed to modify the `/dev/vmnet*` devices. See the VMware KB for details.

To make the change permanent, edit `/etc/init.d/vmware`, and in function `vmwareStartVmnet()`, add commands to `chgrp` and `chown` the `vmnet` devices to a group which contains user running VMware Workstation.

34.31.2 Proxmox VE, KVM, and QEMU Issues

Use VirtIO (`vtnet(4)`) or e1000 NICs (`em(4)`), not the `ed(4)` NICs or CARP VIPs will never leave the INIT state.

34.31.3 VirtualBox Issues

From this thread:

Setting **Promiscuous mode: Allow All** on the relevant interfaces of the VM allows CARP to function on any interface type (Bridged, Host-Only, Internal)

34.32 Troubleshooting High CPU Load

First, open a shell from SSH or the serial/VGA console (option 8).

Typically one of these commands will include some obvious consumer of large amounts of system resources. For example, if the system CPU usage is high, it may be the packet filter. If a VPN process is using lots of CPU, the hardware may not be able to process more VPN traffic. If it is a NIC, see *Hardware Tuning and Troubleshooting*, etc.

If the solution is not obvious based on the output, post the collected information on the forum or contact Netgate TAC for further assistance.

34.32.1 View CPU Processes

To view the top processes, including interrupt processing CPU usage and system CPU:

```
top -aSH
```

34.32.2 View Interrupt Counters

To view the interrupt counters and other system usage:

```
systat -vmstat 1
```

34.32.3 View mbuf Usage

To view the mbuf usage:

```
netstat -m
```

Note: Alternately, check the dashboard mbuf counter, and the graph under **Status > Monitoring** on the **System** tab.

34.32.4 View I/O Operations

To view I/O operations:

```
systat -iostat 1
```

Or:

```
top -aSH
```

Then press `m` to switch to **I/O mode** to view disk activity.

34.33 Troubleshooting Installation Issues

The vast majority of the time, installations of pfSense® software finish with no problems. The following sections describe the most common problems and the steps to resolve them.

See also:

Troubleshooting Boot Issues

34.33.1 Installer Network Connectivity Problems

As the installer requires network connectivity getting the WAN settings correct is critical to its success.

If the installer is unable to contact Netgate servers it will display an error saying "Cannot verify the eligibility of this system, please try again." This could be due to a network configuration or connectivity issue, for example. Double-check the WAN settings before attempting the installation again.

If the installer is still unable to achieve outbound connectivity, it may need to be relocated behind a different connection or on a different network through which it can directly reach the Internet.

34.33.2 Errors During Installation

Errors may occur during the installation, for example if the network connection is interrupted or if the installer encounters a problem with the hardware.

The installer saves a log containing all the installation output to a file named `/tmp/install-log.txt`.

After the installer encounters an error, it displays a notice stating the installation failed and then exits to a shell prompt.

From that shell prompt, it's possible to copy that log file off either over the network with SCP or by copying it to a USB drive, for example.

34.33.3 Boot from Install Media Fails

Due to the wide array of hardware combinations in use, it is not uncommon for a memstick or CD to boot incorrectly (or not at all). Given the unpredictable nature of commodity hardware, using hardware from the Netgate Store is the only guaranteed path to success.

That said, the most common problems and solutions are:

USB Flash Drive Booting
: Some BIOS implementations can be picky about booting from USB media. If booting from one device fails, try a different one.

USB 3 Ports
: Certain combinations of USB drives and ports, especially USB 3.0 ports, may not work correctly. Try a USB 2.0 memstick and/or a USB 2.0 port.

BIOS Issues
: Update to the most recent BIOS, and disable any unneeded peripherals such as Firewire, Floppy Drives, and Audio.

Dirty Optical Drive
: Clean the drive with a cleaning disc or a can of compressed air, try another drive, or use USB media instead.

Bad Optical Media
Burn another disc and/or burn the disc at a lower speed. Perhaps try another brand of optical media or use USB media instead.

SATA/IDE Cable Issues
Try a different cable between the DVD drive and the controller or motherboard or use USB media instead.

34.33.4 Boot from hard drive after installation fails

After the installation completes and the firewall restarts, there are conditions which may prevent the operating system from fully booting. The most common reasons are typically BIOS-related. For example, a BIOS implementation may not boot from a disk using GPT or ZFS, or may require UEFI.

Some of these may be worked around by choosing different options for the partition layout during the installation process. Upgrading the BIOS to the latest version available may also help.

Altering the SATA options in the BIOS has improved booting in some situations as well. If a SATA hard drive is being used, experiment with changing the SATA options in the BIOS for settings such as AHCI, Legacy, or IDE. AHCI is the best mode to use with current versions of pfSense software.

See also:

For additional troubleshooting techniques, see *Troubleshooting Boot Issues*.

34.33.5 Interface link up not detected

If the firewall complains that it did not detect an interface link up event during automatic assignment, first make sure that the cable is unplugged and that the interface does not have a link light prior to choosing the link detection option. After selecting the option, plug the cable back into the interface and ensure it has a link light prior to pressing `Enter`. Test or replace the cable in question if it does not show a link light on the switch and/or NIC port once it is connected.

If a network cable is connected directly between two computers and not to a switch, and one of those pieces of hardware is older (e.g. 10/100 NIC) ensure that a crossover cable is being used. Gigabit adapters all implement Auto-MDIX and will handle this internally, but many older 10/100 adapters do not. Similarly, if connecting a firewall running pfSense software to a switch without Auto-MDIX, use a straight-through patch cable.

If the interface is properly connected, but the firewall still does not detect the link event, the network interface may not properly detect or report link status to the operating system or driver. In this case, manually assigning the interfaces is necessary.

34.33.6 Hardware Troubleshooting

The following suggestions will help resolve general hardware issues.

Booting from USB

If the boot stops with a `mountroot>` prompt while booting off the installer image, usually with USB CD/DVD drives, escape to the loader prompt from the boot menu and run the following:

```
set kern.cam.boot_delay=10000
boot
```

At which point the boot will continue normally.

If the firewall is running permanently from a medium that requires this delay, add the following line as a *Loader Tunable*:

```
kern.cam.boot_delay=10000
```

Remove unnecessary hardware

If the firewall contains hardware that will not be used, remove or disable it. This normally isn't an issue, but can cause problems and has the potential to reduce performance. If an unused piece of hardware is removable, take it out of the firewall or disable it in the BIOS.

Upgrade the BIOS

The second most common fix for hardware problems is upgrading the BIOS to the latest revision. People seem to have a hard time believing this one, but trust us, do it. BIOS updates commonly fix bugs in hardware. It isn't uncommon to hit problems induced by hardware bugs on systems that have been stable running Windows for long periods of time. Either Windows doesn't trigger the bug, or has a workaround, as Netgate TAC has encountered this on multiple occasions. Things that BIOS updates can fix include: Failing to boot, time keeping problems, general instability, and other issues like hardware compatibility.

Reset BIOS settings to factory defaults

Recycled systems may have an atypical BIOS configuration. Most contain an option allowing factory default options to be loaded. Use this option to get a fresh start on the BIOS settings.

Other BIOS settings

If the BIOS allows power management configuration, try toggling that option. Look for anything else that seems relevant to whichever aspect of the installation is failing. If it gets to this point, the target hardware is probably a lost cause and alternate hardware may be necessary. Also check to see if the BIOS has an event log that may list hardware errors such as memory test failures.

If the hardware uses a new or recent chipset, a development version of pfSense software may work. Depending on the development cycle of pfSense software, snapshots may be available via upgrade or potentially for installation from within the installer.

Other Hardware Issues

The target hardware may be faulty, which testing with diagnostic software may reveal. Test the hard drive with diagnostic software from the OEM, and test the memory with a program such as memtest86+. These and more tools are available on the "Ultimate Boot CD", which is preloaded with numerous free and/or open source hardware diagnostic utilities.

Also ensure that all the fans are spinning at speed, and that no components are overheating. If this is older reused hardware, compressed/canned air cleaning of the fans and heat sinks can work wonders.

34.34 Troubleshooting IPsec VPNs

Due to the finicky nature of IPsec it is not unusual for trouble to arise with tunnels when creating them initially or over time.

Follow the troubleshooting advice in this section to diagnose and solve most common problems with IPsec tunnels on pfSense® software.

34.34.1 Troubleshooting IPsec Connections

IPsec connection names

IPsec tunnels follow a consistent naming pattern when forming connection names used in the strongSwan configuration. These names are printed in the IPsec status and can also be found in the IPsec configuration file (`/var/etc/ipsec/swanctl.conf`), the IPsec log, and the output of various `swanctl` commands.

Non-mobile tunnels all use an IKE connection named `conX` where `X` is the phase 1 IKE ID.

Phase 2 child definitions use slightly different names based on the tunnel settings:

For normal IKEv2 tunnels without **Split Connections** enabled all phase 2 entries are combined into a single child definition. In this case the connections are named `conX` where `X` is the phase 1 IKE ID and this is identical to the name of the IKE portion of the connection.

For IKEv1 tunnels and for IKEv2 tunnels with **Split Connections** enabled each phase 2 entry is defined as a separate child. In this case the child definitions are named `conX_Y` where `X` is the phase 1 IKE ID and `Y` is the phase 2 `reqid`.

Note: The phase 1 IKE ID and phase 2 `reqid` are printed in the IPsec tunnel list and on the page when editing those entries.

To see a list of current connections, run the following command from the shell:

```
# swanctl --list-conns
```

The output of that command lists the IKE connection name first (e.g. `con1`) with no indentation. Child definitions are listed at the end of a tunnel entry and are indented.

Manually connect IPsec from the shell

Connections can be manually initiated and terminated from the shell using the `swanctl` command.

Tip: When initiating a tunnel in this way, `swanctl` will output only the relevant logs to the terminal. This is much easier than attempting to follow the log file contents in other ways.

The connection name for a tunnel must be used in this case, such as `con1` or `con2_1`.

Note: To locate the correct `con` identifier, see *IPsec connection names*.

The following command will attempt to initiate the IKE portion of a tunnel (phase 1):

```
# swanctl --initiate --ike conX
```

The following command will attempt to initiate the child SA portion of a tunnel (phase 2) as well as IKE if it is not already connected:

```
# swanctl --initiate --child conX
```

Terminating a tunnel uses similar syntax.

Terminate IKE connection (also terminates all child connections):

```
# swanctl --terminate --ike conX
```

Terminate a child connection:

```
# swanctl --terminate --child conX
```

Tunnel does not establish

First check the service status at **Status > Services**. If the IPsec service is stopped, check if there is at least one configured and enabled IPsec tunnel (*IPsec Tunnels Tab*).

If the service is running, check the firewall logs at **Status > System Logs**, **Firewall** tab. Look for entries that indicate that the connection is being blocked. If the tunnel is not establishing, check for UDP entries for ports `500` and `4500`. Rules are normally added automatically for IPsec (*IPsec and firewall rules*), but that feature can be disabled or there may be edge cases where the firewall cannot identify the remote IPsec gateway. Add rules to pass traffic if needed.

The single most common cause of failed IPsec tunnel connections is a configuration mismatch. Often it is something small, such as a DH group set differently, or perhaps a subnet mask of /24 on one side and /32 on the other in the phase 2 networks. Some routers (Linksys, for one) also like to hide certain options behind "Advanced" buttons or make assumptions. A lot of trial and error may be involved, and a lot of log reading, but ensuring that both sides match precisely will help the most.

Depending on the Internet connections on either end of the tunnel, it is also possible that a router involved on one side or the other does not properly handle IPsec traffic. This is a larger concern with mobile clients and networks where NAT is involved outside the actual IPsec endpoints. The problems are generally with the ESP protocol and problems with it being blocked or mishandled along the way. NAT Traversal (NAT-T) encapsulates ESP in UDP port `4500` traffic to work around these issues. Typically, this situation is detected automatically but in some edge cases it can help to force NAT traversal for IKEv1 tunnels.

"Random" tunnel disconnects/DPD failures on low-end routers

If IPsec tunnels are dropped on low-end hardware that is pushing the limits of its CPU, DPD on the tunnel may need disabled. Such failures tend to correlate with times of high bandwidth usage. This happens when the CPU on a low-power system is tied up with sending IPsec traffic or is otherwise occupied. Due to the CPU overload it may not take the time to respond to DPD requests or see a response to a request of its own. As a consequence, the tunnel will fail a DPD check and be disconnected. This is a clear sign that the hardware is being driven beyond its capacity. If this happens, consider replacing the firewall with a more powerful model.

Tunnels establish and work but fail to renegotiate

In some cases a tunnel will function properly but once the phase 1 or phase 2 lifetime expires the tunnel will fail to renegotiate properly. This can manifest itself in a few different ways, each with a different resolution.

DPD is not available and one side drops while the other remains

Consider this scenario, which DPD is designed to prevent, but can happen in places where DPD is not enabled or implemented:

- A tunnel is established from Site A to Site B, from traffic initiated at Site A.
- Site B expires the phase 1 or phase 2 before Site A.
- Site A will believe the tunnel is up and continue to send traffic as though the tunnel is working properly.
- Only when the Site A phase 1 or phase 2 lifetime expires will it renegotiate as expected.

In this scenario, the likely things resolutions are:

- Check to make sure all the settings match on both sides, especially the phase 1 **DH Group** and phase 2 **PFS** values.
- Enable DPD, or Site B must send traffic to Site A which will cause the entire tunnel to renegotiate. The easiest way to make this happen is to enable a keep alive mechanism on both sides of the tunnel.
- Enable the periodic check keep alive method on one end (*Configuring IPsec Keep Alive*).

Tunnel establishes when initiating but not when responding

If a tunnel will sometimes establish, but not always, generally there is a settings mismatch. The tunnel may still establish because if the settings presented by one side are more secure the other may accept them, but not the other way around.

Lifetime mismatches do not cause a failure in phase 1 or phase 2.

To track down these failures, configure the logs as shown in *Troubleshooting IPsec Logs* and attempt to initiate the tunnel from each side, then check the logs.

Tunnel establishes at start but not when disconnected

An IPsec tunnel can be disconnected for a variety of reasons. For example, connectivity being interrupted to the far side, the remote being down or offline for an extended time, or even a manual or policy action on the far side.

Note: This is not the same scenario as a rekey or reauthentication event, which will rebuild the appropriate parts of the tunnel and remain active.

A tunnel mode IPsec instance will connect at start and when it disconnects, will connect again on demand. This happens due to trap policies which trigger initiation when traffic attempts to use the tunnel. A tunnel mode IPsec connection can be reconnected without manual intervention by the automatic ping keep alive function on a phase 2 entry.

VTI mode IPsec is not compatible with trap policies, so it is not capable of using this tactic. As such, a VTI tunnel may need help to stay up and running at all times.

There are two workarounds that may help in this case:

Keep Alive - Periodic Check

The *IPsec phase 2 Keep Alive option* to perform a periodic IPsec status check is ideally suited to this case. When enabled, if a given phase 2 is down it will trigger an initiation directly.

This works with VTI because it does not rely on trap policies.

Note: This feature is new in pfSense® Plus software version 22.01 and CE 2.6.0.

Child SA Actions

Another tactic to keep a tunnel up is to set it to initiate immediately at start and automatically reconnect if it gets disconnected. This should only be set on **one** side of a tunnel.

Child SA Start Action

Set the start action to *Initiate at start*. This will trigger a tunnel initiation when the IPsec daemon starts, such as at boot time.

Note: This does not trigger when the IPsec configuration is changed and reloaded, only when the daemon loads the configuration the first time at startup.

Child SA Close Action

Set the close action to *Restart/Reconnect* which will attempt to immediately reconnect the child SA if it gets disconnected.

Depending on the reason the tunnel was disconnected, this may or may not be helpful. For example, if the reason the tunnel disconnected was a local cause, these events may not trigger. The periodic check keep alive method is much more reliable, but only available on current versions of pfSense software.

Tunnel stops attempting connections after timeout

If the remote end of an IPsec tunnel is down when the tunnel attempts to initiate at start, but fails, it may eventually time out and stop trying to connect.

The solution here is similar to the previous scenario above, which is to enable keep alive options for the tunnel which will trigger a fresh initiation periodically if the tunnel is down.

34.34.2 Troubleshooting IPsec Traffic

Tunnel establishes but no traffic passes

The first place to look if a tunnel comes up but will not pass traffic is the IPsec firewall rules tab. If Site A cannot reach Site B, check the Site B firewall log and rules. Conversely, if Site B cannot contact Site A, check the Site A firewall log and rules.

Inspect the firewall logs at **Status > System Logs**, on the **Firewall** tab. Check for log entries indicating traffic is blocked involving the subnets used in the IPsec tunnel. Also check for traffic on the WAN interface used by the tunnel for the protocol **ESP** or UDP port `4500` both of which could be used to carry encapsulated IPsec traffic. If there are log entries which match either case, continue on to check the rules. If there are no log entries indicating blocked packets, revisit the section on IPsec routing considerations in *Client Routing and Gateway Considerations*.

Blocked packets on the IPsec or `enc0` interface indicate that the tunnel itself has established, but traffic is being blocked by firewall rules. Blocked packets on the LAN or other internal interface may indicate that an additional rule may be needed on that interface ruleset to allow traffic from the internal subnet out to the remote end of the IPsec tunnel. Blocked packets on WAN type interfaces would prevent a tunnel from establishing. Typically, this only happens when the automatic VPN rules are disabled. Adding a rule to allow the **ESP** protocol along with UDP ports `500` and `4500` from that remote IP address will allow the tunnel to establish and pass traffic. In the case of mobile tunnels, allow traffic from any source to connect to those ports. See *IPsec and firewall rules* for more details.

Rules for the IPsec interface can be found under **Firewall > Rules**, on the **IPsec** tab. Common mistakes include setting a rule to only allow *TCP* traffic, which means things like ICMP ping and DNS would not work across the tunnel. See *Firewall* for more information on how to properly create and troubleshoot firewall rules.

In some cases it is possible that a setting mismatch can also cause traffic to fail passing the tunnel. In one instance, a subnet defined on a third-party firewall was `192.0.2.1/24`, and on the firewall running pfSense® software it was `192.0.2.0/24`. The tunnel established, but traffic would not pass until the subnet was corrected.

Routing issues are another possibility. Running a `traceroute` (`tracert` on Windows) to an IP address on the opposite side of the tunnel can help track down these types of problems. Repeat the test from both sides of the tunnel. Check *Client Routing and Gateway Considerations* for more information. When using `traceroute`, traffic which enters and leaves the IPsec tunnel will seem to be missing interim hops. This is normal and part of how IPsec works. Traffic which does not properly enter an IPsec tunnel will appear to leave the WAN interface and route outward across the Internet, which would point to either a routing issue such as the firewall running pfSense software not being the gateway (as in *Client Routing and Gateway Considerations*), policy routing rules matching the traffic, an incorrectly specified remote subnet on the tunnel definition, a tunnel which has been disabled.

Some hosts work but not all

If traffic between some hosts over the VPN functions properly, but some hosts do not, this is commonly one of four things:

Missing, incorrect or ignored default gateway
: If the device does not have a default gateway or has one pointing to something other than the firewall running pfSense software, it does not know how to properly get back to the remote network on the VPN (see *Client Routing and Gateway Considerations*).

  Some devices, even with a default gateway specified, do not use that gateway. This has been seen on various embedded devices, including IP cameras and some printers. There isn't anything that can be done about that other than getting the software on the device fixed. This can be verified by running a packet capture on the inside interface of the firewall connected to the network containing the device. Troubleshooting with tcpdump is covered in *Using tcpdump on the command line*, and an IPsec-specific example can be found in *IPsec tunnel will not connect*.

  If traffic is observed leaving the inside interface of the firewall, but no replies return, the device is not properly routing its reply traffic or could potentially be blocking it via a local client firewall.

Incorrect subnet mask
: If the subnet in use on one end is `10.0.0.0/24` and the other is `10.254.0.0/24`, and a host has an incorrect subnet mask of `255.0.0.0` or `/8`, it will never be able to communicate across the VPN because it thinks the remote VPN subnet is part of the local network and hence routing will not function properly. The system with the broken configuration will attempt to contact the remote system via ARP instead of using the gateway.

Host firewall
: If there is a firewall on the target host, it may not be allowing the connections. Check for things like Windows Firewall, iptables, or similar utilities that may be preventing the traffic from being processed by the host.

Firewall rules on pfSense software
: Ensure the rules on both ends allow the desired network traffic.

Connection hangs

IPsec does not gracefully handle fragmented packets. Many of these issues have been resolved over the years, but there may be lingering problems and edge cases. If hangs or packet loss are seen only when using specific protocols (SMB, RDP, etc.), MSS clamping for the VPN may be necessary. MSS clamping can be activated under *Firewall & NAT*. A good starting point for MSS clamping is `1400`. If that works slowly increase the MSS value until the breaking point is hit, then back off a little from there.

Disappearing traffic

If IPsec traffic arrives but never appears on the IPsec interface (`enc0`), check for conflicting routes/interface IP addresses. For example, if an IPsec tunnel is configured with a remote network of `192.0.2.0/24` and there is a local OpenVPN server with a tunnel network of `192.0.2.0/24` then the ESP traffic may arrive, strongSwan may process the packets, but they never show up on `enc0` as arriving to the OS for delivery.

Resolve the duplicate interface/route and the traffic will begin to flow.

34.34.3 Troubleshooting IPsec Logs

Note: Examples presented in this document contain logs edited for brevity but significant messages remain.

Logging for IPsec can provide useful information. To configure IPsec logging for diagnosing tunnel issues with pfSense® software, the following procedure yields the best balance of information:

- Navigate to **VPN > IPsec** on the **Advanced Settings** tab
- Set **IKE SA**, **IKE Child SA**, and **Configuration Backend** to *Diag*
- Set all other log settings to *Control*
- Click Save

Note: Changing logging options is not disruptive to IPsec tunnels.

Tip: Though this section assumes log messages are obtained from the IPsec log, using a manual connection attempt (*Manually connect IPsec from the shell*) can yield more focused results when initiating. When initiating manually using a shell command the messages are printed to the console and not mixed with logs from other connections, making it much simpler to find relevant log messages.

IPsec log interpretation

The IPsec logs available at **Status > System Logs**, on the **IPsec** tab contain a record of the tunnel connection process and some messages from ongoing tunnel maintenance activity. Some typical log entries are listed in this section, both good and bad. The main things to look for are key phrases that indicate which part of a connection worked. If "IKE_SA … established" is present in the log, that means phase 1 was completed successfully, and a Security Association was negotiated. If "CHILD_SA … established" is present, then phase 2 has also been completed and the tunnel is up.

In the following examples, the logs have been configured as listen in *Troubleshooting IPsec Logs* and irrelevant messages may be omitted. Bear in mind that these are samples and the specific ID numbers, IP addresses, and so forth will vary.

Successful connections

When a tunnel has been successfully established both sides will indicate that an **IKE SA** and a **Child SA** have been established. When multiple phase 2 definitions are present with IKEv1, a child SA is negotiated for each phase 2 entry.

Log output from the initiator:

```
charon: 09[IKE] IKE_SA con2[11] established between 192.0.2.90[192.0.2.90]...192.0.2.
↪74[192.0.2.74]
charon: 09[IKE] CHILD_SA con2{2} established with SPIs cf4973bf_i c1cbfdf2_o and TS 192.
↪168.48.0/24|/0 === 10.42.42.0/24|/0
```

Log output from the responder:

```
charon: 03[IKE] IKE_SA con1[19] established between 192.0.2.74[192.0.2.74]...192.0.2.
↪90[192.0.2.90]
```

(continues on next page)

(continued from previous page)

```
charon: 16[IKE] CHILD_SA con1{1} established with SPIs c1cbfdf2_i cf4973bf_o and TS 10.
↪42.42.0/24|/0 === 192.168.48.0/24|/0
```

Failed connection examples

These examples show failed connections for varying reasons. In most cases it is clear from the examples that the initiator does not receive messages about specific items that do not match, so the responder logs are much more informative. This is done to protect the security of the tunnel, it would be insecure to provide messages to a potential attacker that would give them information about how the tunnel is configured.

Phase 1 main / aggressive mismatch

In this example, the initiator is set for Aggressive mode while the responder is set for Main mode.

Log output from the initiator:

```
charon: 15[IKE] initiating Aggressive Mode IKE_SA con2[1] to 203.0.113.5
charon: 15[IKE] received AUTHENTICATION_FAILED error notify
charon: 13[ENC] parsed INFORMATIONAL_V1 request 1215317906 [ N(AUTH_FAILED) ]
charon: 13[IKE] received AUTHENTICATION_FAILED error notify
```

Log output from the responder:

```
charon: 13[IKE] Aggressive Mode PSK disabled for security reasons
charon: 13[ENC] generating INFORMATIONAL_V1 request 2940146627 [ N(AUTH_FAILED) ]
```

Note that the logs on the responder state clearly that Aggressive mode is disabled, which is a good clue that the mode is mismatched.

In the reverse case, if the side set for Main mode initiates, the tunnel to a firewall running pfSense software will establish since Main mode is more secure.

Phase 1 identifier mismatch

When the identifier does not match the initiator only shows that the authentication failed, but does not give a reason. The responder states that it is unable to locate a peer, which indicates that it could not find a matching phase 1, which implies that no matching identifier could be located.

Log output from the initiator:

```
charon: 10[ENC] parsed INFORMATIONAL_V1 request 4216246776 [ HASH N(AUTH_FAILED) ]
charon: 10[IKE] received AUTHENTICATION_FAILED error notify
```

Log output from the responder:

```
charon: 12[CFG] looking for pre-shared key peer configs matching 203.0.113.5...198.51.
↪100.3[someid]
charon: 12[IKE] no peer config found
charon: 12[ENC] generating INFORMATIONAL_V1 request 4216246776 [ HASH N(AUTH_FAILED) ]
```

Phase 1 pre-shared key mismatch

A mismatched pre-shared key can be a tough to diagnose. An error stating the fact that this value is mismatched is not printed in the log, instead this messages is shown:

Log output from the initiator:

```
charon: 09[ENC] invalid HASH_V1 payload length, decryption failed?
charon: 09[ENC] could not decrypt payloads
charon: 09[IKE] message parsing failed
```

Log output from the responder:

```
charon: 09[ENC] invalid ID_V1 payload length, decryption failed?
charon: 09[ENC] could not decrypt payloads
charon: 09[IKE] message parsing failed
```

When the above log messages are present check the Pre-Shared Key value on both sides to ensure they match.

Phase 1 encryption algorithm mismatch

Log output from the initiator:

```
charon: 14[ENC] parsed INFORMATIONAL_V1 request 3851683074 [ N(NO_PROP) ]
charon: 14[IKE] received NO_PROPOSAL_CHOSEN error notify
```

Log output from the responder:

```
charon: 14[CFG] received proposals: IKE:AES_CBC_128/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_1024
charon: 14[CFG] configured proposals: IKE:AES_CBC_256/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_
↪1024
charon: 14[IKE] no proposal found
charon: 14[ENC] generating INFORMATIONAL_V1 request 3851683074 [ N(NO_PROP) ]
```

In this case, the log entry shows the problem exactly: The initiator was set for AES 128 encryption, and the responder is set for AES 256. Set both to matching values and then try again.

Phase 1 hash algorithm mismatch

Log output from the initiator:

```
charon: 10[ENC] parsed INFORMATIONAL_V1 request 2774552374 [ N(NO_PROP) ]
charon: 10[IKE] received NO_PROPOSAL_CHOSEN error notify
```

Log output from the responder:

```
charon: 14[CFG] received proposals: IKE:AES_CBC_256/MODP_1024
charon: 14[CFG] configured proposals: IKE:AES_CBC_256/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_
↪1024
charon: 14[IKE] no proposal found
charon: 14[ENC] generating INFORMATIONAL_V1 request 2774552374 [ N(NO_PROP) ]
```

The hash algorithm is indicated by the `HMAC` portion of the logged proposals. As can be seen above, the received and configured proposals do not have matching HMAC entries. Set them both to match and try again.

Phase 1 DH group mismatch

Log output from the initiator:

```
charon: 11[ENC] parsed INFORMATIONAL_V1 request 316473468 [ N(NO_PROP) ]
charon: 11[IKE] received NO_PROPOSAL_CHOSEN error notify
```

Log output from the responder:

```
charon: 14[CFG] received proposals: IKE:AES_CBC_256/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_8192
charon: 14[CFG] configured proposals: IKE:AES_CBC_256/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_
↪1024
charon: 14[IKE] no proposal found
charon: 14[ENC] generating INFORMATIONAL_V1 request 316473468 [ N(NO_PROP) ]
```

DH group is indicated by the `MODP` portion of the listed proposal. As indicated by the log messages, the initiator was set for 8192 (Group 18) and the responder was set for 1024 (Group 2). This error can be corrected by setting the DH group setting on both ends of the tunnel to a matching value.

Phase 2 network mismatch

In the following example, the phase 2 entry on the initiator side is set for `10.3.0.0/24` to `10.5.0.0/24`. The responder is not set to match as it lists `10.5.1.0/24` instead.

Log output from the initiator:

```
charon: 08[CFG] proposing traffic selectors for us:
charon: 08[CFG] 10.3.0.0/24|/0
charon: 08[CFG] proposing traffic selectors for other:
charon: 08[CFG] 10.5.0.0/24|/0
charon: 08[ENC] generating QUICK_MODE request 316948142 [ HASH SA No ID ID ]
charon: 08[NET] sending packet: from 198.51.100.3[500] to 203.0.113.5[500] (236 bytes)
charon: 08[NET] received packet: from 203.0.113.5[500] to 198.51.100.3[500] (76 bytes)
charon: 08[ENC] parsed INFORMATIONAL_V1 request 460353720 [ HASH N(INVAL_ID) ]
charon: 08[IKE] received INVALID_ID_INFORMATION error notify
```

Log output from the responder:

```
charon: 08[ENC] parsed QUICK_MODE request 2732380262 [ HASH SA No ID ID ]
charon: 08[CFG] looking for a child config for 10.5.0.0/24|/0 === 10.3.0.0/24|/0
charon: 08[CFG] proposing traffic selectors for us:
charon: 08[CFG] 10.5.1.0/24|/0
charon: 08[CFG] proposing traffic selectors for other:
charon: 08[CFG] 10.3.0.0/24|/0
charon: 08[IKE] no matching CHILD_SA config found
charon: 08[IKE] queueing INFORMATIONAL task
charon: 08[IKE] activating new tasks
charon: 08[IKE] activating INFORMATIONAL task
charon: 08[ENC] generating INFORMATIONAL_V1 request 1136605099 [ HASH N(INVAL_ID) ]
```

The responder logs lists both the networks it received (`child config` line in the log) and what it has configured locally (`proposing traffic selectors for...` lines in the log). By comparing the two, a mismatch can be spotted. The `no matching CHILD_SA config found` line in the log will always be present when this mismatch occurs and that directly indicates that it could not find a phase 2 definition to match the values received from the initiator.

Phase 2 encryption algorithm mismatch

Log output from the initiator:

```
charon: 14[CFG] configured proposals: ESP:AES_CBC_128/HMAC_SHA1_96/NO_EXT_SEQ
charon: 14[ENC] generating QUICK_MODE request 759760112 [ HASH SA No ID ID ]
charon: 14[NET] sending packet: from 198.51.100.3[500] to 203.0.113.5[500] (188 bytes)
charon: 14[NET] received packet: from 203.0.113.5[500] to 198.51.100.3[500] (76 bytes)
charon: 14[ENC] parsed INFORMATIONAL_V1 request 1275272345 [ HASH N(NO_PROP) ]
charon: 14[IKE] received NO_PROPOSAL_CHOSEN error notify
```

Log output from the responder:

```
charon: 13[CFG] selecting proposal:
charon: 13[CFG] no acceptable ENCRYPTION_ALGORITHM found
charon: 13[CFG] received proposals: ESP:AES_CBC_128/HMAC_SHA1_96/NO_EXT_SEQ
charon: 13[CFG] configured proposals: ESP:AES_CBC_256/HMAC_SHA1_96/NO_EXT_SEQ
charon: 13[IKE] no matching proposal found, sending NO_PROPOSAL_CHOSEN
charon: 13[IKE] queueing INFORMATIONAL task
charon: 13[IKE] activating new tasks
charon: 13[IKE] activating INFORMATIONAL task
charon: 13[ENC] generating INFORMATIONAL_V1 request 1275272345 [ HASH N(NO_PROP) ]
```

In this case the initiator receives a message that the responder could not find a suitable proposal (`received NO_PROPOSAL_CHOSEN`), and from the responder logs it is obvious this was due to the sites being set for different encryption types, AES 128 on one side and AES 256 on the other.

Phase 2 hash algorithm mismatch

Log output from the initiator:

```
charon: 10[CFG] configured proposals: ESP:AES_CBC_256/HMAC_SHA2_512_256/NO_EXT_SEQ
charon: 10[ENC] generating QUICK_MODE request 2648029707 [ HASH SA No ID ID ]
charon: 10[NET] sending packet: from 198.51.100.3[500] to 203.0.113.5[500] (188 bytes)
charon: 10[NET] received packet: from 203.0.113.5[500] to 198.51.100.3[500] (76 bytes)
charon: 10[ENC] parsed INFORMATIONAL_V1 request 757918402 [ HASH N(NO_PROP) ]
charon: 10[IKE] received NO_PROPOSAL_CHOSEN error notify
```

Log output from the responder:

```
charon: 11[CFG] selecting proposal:
charon: 11[CFG] no acceptable INTEGRITY_ALGORITHM found
charon: 11[CFG] received proposals: ESP:AES_CBC_256/HMAC_SHA2_512_256/NO_EXT_SEQ
charon: 11[CFG] configured proposals: ESP:AES_CBC_256/HMAC_SHA1_96/NO_EXT_SEQ
charon: 11[IKE] no matching proposal found, sending NO_PROPOSAL_CHOSEN
charon: 11[IKE] queueing INFORMATIONAL task
charon: 11[IKE] activating new tasks
charon: 11[IKE] activating INFORMATIONAL task
charon: 11[ENC] generating INFORMATIONAL_V1 request 757918402 [ HASH N(NO_PROP) ]
```

Similar to a phase 1 hash algorithm mismatch, the `HMAC` values in the log entries do not line up. However, the responder also logs a clearer message `no acceptable INTEGRITY_ALGORITHM found` when this happens in phase 2.

Phase 2 PFS mismatch

Log output from the initiator:

```
charon: 06[ENC] generating QUICK_MODE request 909980434 [ HASH SA No KE ID ID ]
charon: 06[NET] sending packet: from 198.51.100.3[500] to 203.0.113.5[500] (444 bytes)
charon: 06[NET] received packet: from 203.0.113.5[500] to 198.51.100.3[500] (76 bytes)
charon: 06[ENC] parsed INFORMATIONAL_V1 request 3861985833 [ HASH N(NO_PROP) ]
charon: 06[IKE] received NO_PROPOSAL_CHOSEN error notify
```

Log output from the responder:

```
charon: 08[CFG] selecting proposal:
charon: 08[CFG] no acceptable DIFFIE_HELLMAN_GROUP found
charon: 08[CFG] received proposals: ESP:AES_CBC_256/HMAC_SHA1_96/MODP_2048/NO_EXT_SEQ
charon: 08[CFG] configured proposals: ESP:AES_CBC_256/HMAC_SHA1_96/NO_EXT_SEQ
charon: 08[IKE] no matching proposal found, sending NO_PROPOSAL_CHOSEN
charon: 08[ENC] generating INFORMATIONAL_V1 request 3861985833 [ HASH N(NO_PROP) ]
```

Perfect Forward Secrecy (PFS) works like DH groups on phase 1, but is optional. When PFS options do not match a clear message is logged indicating this fact: `no acceptable DIFFIE_HELLMAN_GROUP found`.

Note: In some cases, if one side has PFS set to *off*, and the other side has a value set, the tunnel may still establish and work. The mismatch shown above may only be seen if the values mismatch, for example `1` vs. `5`.

Note: Due to the way IPsec negotiates the first child SA will not use the PFS value from phase 2, but the DH group value from phase 1. Subsequent child SA entries or rekey events will use the value from phase 2. Thus, if a tunnel connects OK at first but fails at rekey, ensure the phase 2 PFS values match.

Mismatched identifier with NAT

In this case, pfSense software is configured for a **Peer Identifier** of *Peer IP address*, but the remote device is actually behind NAT. In this case strongSwan expects the actual private before-NAT IP address as the identifier.

Log output from the responder:

```
charon: 10[IKE] remote host is behind NAT
charon: 10[IKE] IDir '192.0.2.10' does not match to '203.0.113.245'
[...]
charon: 10[CFG] looking for pre-shared key peer configs matching 198.51.100.50...203.0.
↪113.245[192.0.2.10]
```

To correct this condition, change the **Peer Identifier** setting to *IP Address* and then enter the pre-NAT IP address, which in this example is `192.0.2.10`.

34.34.4 Troubleshooting Duplicate IPsec SA Entries

In certain cases an IPsec tunnel may show what appear to be duplicate IKE (phase 1) or Child (phase 2) security association (SA) entries.

Lengthy testing and research uncovered that the main way this starts to happen is when both sides negotiate or renegotiate simultaneously. If both peers initiate, reauthenticate, or rekey phase 1 at the same time, it can result in duplicate IKE SAs. If both peers rekey phase 2 at the same time, it can result in duplicate child SAs.

Mitigating this problem involves ensuring that the chance of simultaneous negotiation is minimized or eliminated. The easiest way to reach that goal is to set higher phase 1 and phase 2 lifetimes on one peer, or at least make sure both sides are **not** set identically.

Specific values vary, but the settings below are the best general advice. Using the suggested values may not be possible due to peer constraints, such as third party vendors which do not implement these settings or insist upon other settings. Variations are mentioned which should accommodate most situations, however, use as many strategies as possible.

Note: The values in this document take precedence over default recommendations in other IPsec examples and recipes.

Current version (2.5.0 and later)

The values in this section can be found in the current GUI. Older versions may have different values or some settings may not be available. Always run the most recent released version to ensure the best experience.

Peer A

Phase 1 (IKE SA)

Key Exchange Version
IKEv2 if both peers are capable.

Life Time
The total IKE SA lifetime as a hard upper limit (e.g. `28800`)

Rekey Time
90% of total IKE SA lifetime (e.g. `25920`).

Reauth Time
`0` to disable reauthentication.

If the peer requires IKEv1 or only allows IKEv2 reauthentication, set this as mentioned in **Rekey Time** above and also enable **Make Before Break** on the **Advanced Settings** tab.

Rand Time
Defaults to 10% of IKE SA **Life Time** (e.g. `2880`). A larger **Rand Time** will decrease the chances of both peers renegotiating simultaneously.

Child SA Close Action
Restart/Reconnect so that this side will reconnect child SA entries when they expire or fail.

Phase 2 (Child SA)

Life Time
Total Child SA lifetime (e.g. `3600` for 1 hour). This peer will attempt to rekey the Child SA before it reaches this limit.

Rekey Time
Leave blank to automatically use 90% of the **Life Time**, or choose a lower amount.

Rand Time
Defaults to 10% of Child SA **Life Time** (e.g. `360`). A larger **Rand Time** will decrease the chances of both peers renegotiating simultaneously.

Peer B

Phase 1 (IKE SA)

Key Exchange Version
IKEv2 if both peers are capable.

Life Time
The total IKE SA lifetime as a hard upper limit, but use a higher lifetime than Peer A by at least 10% (e.g. `31680`). With this peer set higher, Peer A will primarily manage IKE SA renegotiation, reducing the chance of conflicts.

Note: If the remote peer insists their lifetime be set to a specific value, then set peer A lower instead by a similar margin.

Rekey Time
90% of total IKE SA **Life Time**

Reauth Time
Blank (disabled) to disable reauthentication.

If the peer requires IKEv1 or only allows IKEv2 reauthentication, set this as mentioned in **Rekey Time** above and also enable **Make Before Break** on the **Advanced Settings** tab.

Rand Time
Defaults to 10% of IKE SA **Life Time** (e.g. `3168`). A larger **Rand Time** will decrease the chances of both peers renegotiating simultaneously. If using the same **Life Time** as Peer A, then increase this value further. If using a larger **Life Time** for Peer B, then leave this at the default or disable it (`0`).

Responder Only
Checked so that this side **will not** automatically initiate IKE SA negotiation.

Note: This peer can still manually initiate a connection from **Status > IPsec**, but it won't happen automatically.

Child SA Close Action
Close connection and clear SA so that when a Child SA expires, this side will remove the SA and not attempt to renegotiate a new entry.

Phase 2 (Child SA)

Life Time
A larger value than the **Life Time** set on Peer A by at least 10%. For example, if Peer A is set to `3600`, set this to `5400`. That way Peer A will primarily manage Child SA renegotiation.

Note: If the remote peer insists their lifetime be set to a specific value, then set peer A lower instead by a similar margin.

Rekey Time
Leave blank to automatically use 90% of the **Life Time**, or choose a lower amount.

Rand Time
Defaults to 10% of Child SA **Life Time** (e.g. `540`). A larger **Rand Time** will decrease the chances of both peers renegotiating simultaneously. If using the same **Life Time** as Peer A, then increase this value further. If using a larger **Life Time** for Peer B, then leave this at the default or disable it (`0`).

Advanced IPsec Settings (both)

Make Before Break
Checked if phase 1 uses IKEv2 and reauthentication, not relevant otherwise.

Version 2.4.5-p1 and older

The settings are almost all the same as the section above, with a couple changes. The primary difference is in the GUI settings for rekey and reauth. Only the differences are noted below, so follow the previous section except for the values noted here.

Warning: On 2.4.5-p1, setting **Responder Only** in the phase 1 options requires an extra patch, `9a69dd4b8ff6eeeaf5779b7388a10743afae8e91`, which can be applied using the *System Patches Package*.

Peer A

Lifetime
The total time at which this peer will renegotiate the IKE SA (e.g. `28800`)

Margin Time
An amount of time, in seconds, before the **Life Time** is reached when renegotiation begins. Defaults to `540`, but larger values can help reduce the chance of simultaneous renegotiation. Due to the default behavior of the IPsec daemon, this time can be randomly increased up to twice its value to further help avoid both sides choosing the same time. A larger value will help avoid potential collisions.

Disable Rekey
Unchecked

Disable Reauth
Checked

Peer B

Lifetime
A higher time than the same field on Peer A by at least 10%. (e.g. `32400`)

Margin Time
If using the same **Life Time** as Peer A, use a larger value to help avoid simultaneous renegotiation. If using a larger **Life Time** value, then leave this blank or set to the same value as Peer A.

Disable Rekey
Unchecked

Disable Reauth
Checked

Other notes

The strongSwan daemon introduces randomness into the renegotiation process which can help mitigate the problem, but still leaves it up to chance if both peers are using the exact same lifetime values. That is why setting one peer higher, beyond the randomness threshold, is a better practice. The randomness also explains why the problem can take a while to manifest in certain environments as the duplicates do not happen until both peers happen to land on the same random rekey time.

34.35 Troubleshooting L2TP

This section covers troubleshooting steps for the most common problems users encounter with L2TP.

34.35.1 Cannot connect

Check that firewall rules have been added to the external interface where the L2TP traffic enters the firewall. Also make sure the client is connecting to the interface IP address chosen on the L2TP settings.

34.35.2 Connected to L2TP but cannot pass traffic

Ensure firewall rules have been added to the **L2TP VPN** interface as described in *Configure firewall rules for L2TP clients*.

Also ensure the remote subnet across the VPN is different from the local subnet. It is not possible to reach a 192.168.1.0/24 network across the VPN when the local subnet where the client resides is also 192.168.1.0/24, traffic destined for that subnet will never traverse the VPN because it is on the local network. This is why it is important to choose a relatively obscure LAN subnet when using a VPN.

34.35.3 Connection Fails with a Windows Client

If the IPsec layer appears to complete, but no L2TP traffic passes, it is likely a known incompatibility between Windows and the strongSwan daemon used on pfSense® software. There is currently no known workaround except to move the Windows client out from behind NAT, or to use a different style VPN such as IKEv2.

34.35.4 L2TP Traffic Blocked Outbound

In some cases, such as when combined with IPsec, L2TP traffic may also require special handling via floating rules. This appears as blocked traffic in the *outbound* direction in the firewall logs, showing an L2TP server interface.

If this happens, add a floating rule as follows:

- Navigate to **Firewall > Rules**, **Floating** tab
- Click  **Add** to add a new rule to the top of the list

- Set **Action** to *Pass*
- Check **Quick**
- Select *L2TP VPN* for the **Interface**
- Set **Direction** to *Out*
- Set **Protocol** to *TCP*
- Set **Source/Destination** as needed, or set to *any*
- Advanced Features:
  - Set **TCP Flags** to *Any flags*
  - Set **State Type** to *Sloppy State*

34.36 Troubleshooting Access when Locked Out of the Firewall

Under certain circumstances an administrator can be locked out of the GUI. There are a number of ways to regain control, so it is not necessarily a major cause for concern. Some methods are a little tricky, but it is nearly always possible to recover access. The worst-case scenarios require physical access, as anyone with physical access can bypass security measures.

Danger: Let the tactics in this document be a lesson: Physical security of a firewall is **critical**, especially in environments where the firewall is physically located in a common area accessible to people other than authorized administrators.

Before taking any of these steps, try the *Default Username and Password*.

34.36.1 Forgotten Password

The firewall administrator password can easily be reset using the firewall console if it has been lost. Access the physical console (*Connect to the Console*) and use option `3` to change the password for the `admin` account. This option can also reset the `admin` account if it is disabled or expired.

See also:

See *3) Reset admin account and password* for details on how this console menu option works.

34.36.2 Forgotten Password with a Locked Console

If the console is password protected, all is not lost. It takes two reboots to accomplish, but the password can be reset with physical access to the console:

- Connect to the console
- Reboot the firewall
- Choose the *Boot Single User* option (`2`) from the loader menu with the ASCII logo
- Press `Enter` when prompted to start `/bin/sh`
- Remount all partitions as rewritable:

  The specific commands vary based on the filesystem.

  - For devices installed using UFS, see *Re-mount UFS Volumes as Read/Write*.
  - For devices installed using ZFS, see *Re-mount ZFS Volumes as Read/Write*.
- Run the built-in `admin` account reset command:

```
# /etc/rc.initial.password
```

- Follow the prompts to change the password
- Run `/sbin/reboot` to reboot.

See also:

See *3) Reset admin account and password* for details on how this console menu option works.

34.36.3 HTTPS Certificate Problems

If the browser refuses to connect due to the certificate itself, or if the GUI web server refuses to run due to the certificate, then the easiest method to regain access is to generate a new self-signed certificate from the console or SSH:

```
# pfSsh.php playback generateguicert
```

That will create a new self-signed GUI certificate, activate it, and restart the GUI web server.

See also:

SSL/TLS Certificate

34.36.4 HTTP vs HTTPS Confusion

Ensure the client is connecting with the proper protocol, either HTTP or HTTPS. If one doesn't work, try the other. If the GUI has not been configured correctly, the firewall may be running the GUI on an unexpected port and protocol combination, such as:

- `http://<hostname>:443`
- `https://<hostname>:80`

To reset this from the console, reset the LAN interface IP Address, enter the same IP address, and the script will prompt to reset the GUI back to HTTP.

Warning: If the browser previously connected to the GUI using HTTPS, it may refuse to connect using HTTP for security reasons (e.g. HSTS).

34.36.5 Blocked Access with Firewall Rules

If a remote administrator loses access to the GUI due to a firewall rule change, then access can still be obtained from the LAN side. The LAN rules cannot prevent access to the GUI unless the anti-lockout rule is disabled. The anti-lockout rule ensures that hosts on the LAN are able to access the GUI at all times, no matter what the other rules on the LAN interface block.

Having to walk someone on-site through fixing the rule from the LAN is better than losing everything or having to make a trip to the firewall location!

34.36.6 Locked Out by Too Many Failed Login Attempts

Attempting to log in to the GUI or SSH and failing many times will cause the connecting IP address to be added to the lockout table.

To regain access, login successfully from another IP address and then manually remove the entry as follows:

- Navigate to **Diagnostics > Tables**
- Select *sshguard*
- Click  by the entry or entries for workstations to allow again.

The lockout table may also be cleared by the console or ssh in the shell:

```
pfctl -T flush -t sshguard
```

34.36.7 Remotely Circumvent Firewall Lockout with Rules

There are a few ways to manipulate the firewall behavior at the shell to regain access to the firewall GUI. The following tactics are listed in order of how easy they are and how much impact they have on the running system.

Add a rule with EasyRule

The easiest way, assuming the administrator knows the IP address of a remote client PC that needs access, is to use the `easyrule` shell script to add a new firewall rule. In the following example, the `easyrule` script will allow access on the WAN interface, from `x.x.x.x` (the client IP address) to `y.y.y.y` (presumably the WAN IP address) on TCP port `443`:

```
# easyrule pass wan tcp x.x.x.x y.y.y.y 443
```

Once the `easyrule` script adds the rule, the client will be able to access the GUI from the specified source address.

Add an allow all WAN rule from the shell

Another tactic is to temporarily activate an "allow all" rule on the WAN to let a client in.

> **Warning:** An "allow all" style rule is dangerous to have on an interface connected to a public or untrusted network, such as a WAN interface connected to the Internet. **Do not forget to remove the rule added by this script**

To add an "allow all" rule to the WAN interface, run the following command at a shell prompt:

```
# pfSsh.php playback enableallowallwan
```

Once the administrator regains access and fixes the original issue preventing them from reaching the GUI, remove the "allow all" rule from the WAN.

Disable the Firewall

An administrator can (very temporarily) disable firewall rules by using the physical console or SSH.

> **Warning:** This completely disables PF which disables firewall rules and NAT. If the network run by this firewall relies on NAT to function, which most do, then running this command will disrupt connectivity from the LAN to the Internet.

To disable the firewall, connect to the physical console or ssh and use option `8` to start a shell, and then type:

```
# pfctl -d
```

That command will disable the firewall, including all NAT functions. Access to the GUI is now possible from anywhere, at least for a few minutes or until a process on the firewall causes the ruleset to be reloaded (which is almost every page save or **Apply Changes** action). Once the administrator has adjusted the rules and regained the necessary access, turn the firewall back on by typing:

```
# pfctl -e
```

Manual Ruleset Editing

The loaded ruleset is retained in `/tmp/rules.debug`. If the administrator is familiar with PF ruleset syntax, they can edit that file to fix the connectivity issue and reload those rules:

```
# pfctl -f /tmp/rules.debug
```

After getting back into the GUI with that temporary fix, the administrator must perform whatever work is required in the GUI to make the fix permanent. When the rules are saved in the GUI, the temporary edit to `/tmp/rules.debug` will be overwritten.

34.36.8 Remotely Circumvent Firewall Lockout with SSH Tunneling

If remote access to the GUI is blocked by the firewall, but SSH access is allowed, then there is a relatively easy way to get in: SSH Tunneling.

If the GUI is on port `443`, set the SSH client to forward local port `443` (or `4443`, or another port) to remote port `localhost:443`. If the firewall GUI is on another port, use that as the target instead. Then point the browser to `https://localhost`. Add the port to the end of the URL if it differs from the default 443, for example `https://localhost:4443`. If the GUI is using HTTP, change the protocol on the URL to `http://`.

Fill out the options as shown in Figure *Setting Up a Port 443 SSH Tunnel in PuTTY*, then click **Add**.

Once the client connects and authenticates, the GUI is accessible from the redirected local port.

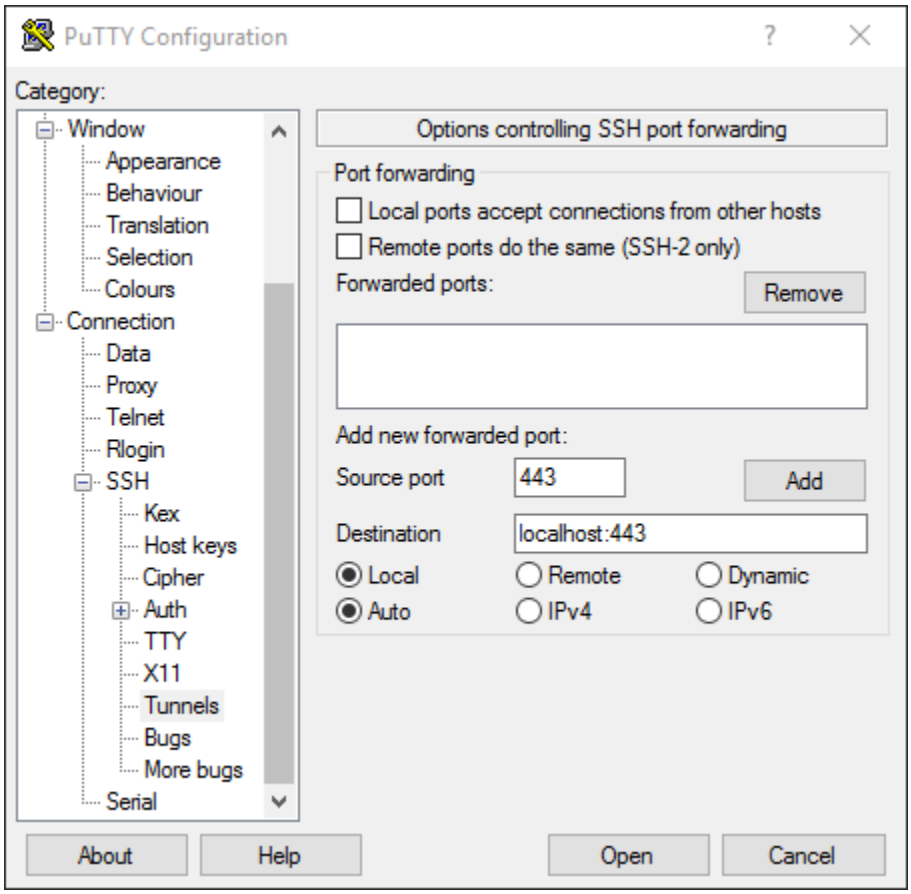


Fig. 4: Setting Up a Port 443 SSH Tunnel in PuTTY

34.36.9 Authentication Server Failure

LDAP and RADIUS authentication for the GUI automatically fall back to the local database if they fail. If the authentication server fails and all local accounts are disabled, locked out, passwords are not known, etc., then to get back in, regain access to the local `admin` account.

Connect to the console (*Connect to the Console*) or ssh and run option `3` to reset the credentials to the *Default Username and Password*.

34.37 Troubleshooting Blocked Log Entries for Legitimate Connection Packets

Sometimes log entries will be present that appear to be blocking legitimate traffic, while labeled with the "Default deny" or even sometimes a pass rule. There are several possible causes for this behavior.

34.37.1 Out-of-State Web Server Packets

The most common example is seeing a connection blocked involving a web server.

Action	Time	Interface	Rule	Source	Destination	Protocol
✕	Apr 17 06:28:59	LAN	Default reject instead of block for multi-WAN (1647375041)	17 :42888	142.251.32.14:443	TCP:PA
✕	Apr 17 06:28:59	LAN	Default reject instead of block for multi-WAN (1647375041)	17 :42888	142.251.32.14:443	TCP:FA
✕	Apr 17 06:34:44	LAN	Default reject instead of block for multi-WAN (1647375041)	17 :55402	142.250.112.95:443	TCP:FA
✕	Apr 17 06:34:50	LAN	Default reject instead of block for multi-WAN (1647375041)	17 :40816	35.244.164.0:443	TCP:FPA
✕	Apr 17 06:41:56	LAN	Default reject instead of block for multi-WAN (1647375041)	17 :48278	216.239.38.57:443	TCP:PA
✕	Apr 17 06:41:56	LAN	Default reject instead of block for multi-WAN (1647375041)	17 :48278	216.239.38.57:443	TCP:FA

Fig. 5: Log entries for blocked out-of-state TCP packets

This is likely due to a TCP FIN packet arriving after firewall has removed the connection state. This happens because on occasion a packet will be lost, and the retransmits will be blocked because the firewall has already closed the connection. Another possible reason for the messages is if a packet arrived too slowly and was outside its expected arrival window. It can also happen when web servers attempt to reuse connections.

In each case, the log entries are harmless and do not indicate a blocked connection. All stateful firewalls do this, though some do not generate log messages for this blocked traffic even if all blocked traffic is logged.

These blocked packets will occur even if rules exist which look as though they should match the traffic, such as an "Allow All" rule. Pass rules for TCP only allow **TCP SYN** packets to create a state. These rules assume TCP traffic with other flags will either be part of an existing state in the state table, or packets with spoofed TCP flags.

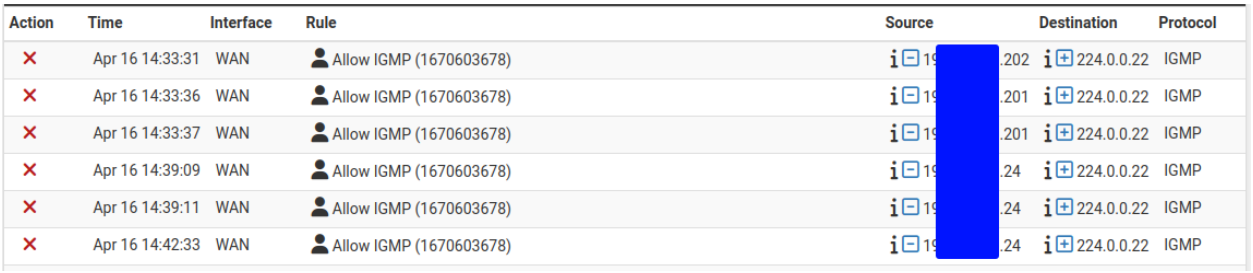
See also:

TCP Flags

34.37.2 Packets with IP Options

PF will drop packets containing IP options unless pass rules specifically allow them using the "Allow IP Options" checkbox in the Advanced section when editing a firewall rule (*IP Options*).

Since this is a special case, the firewall always logs when it drops these packets. However, the rule ID shown in this case is the rule the packet would have otherwise matched.



Action	Time	Interface	Rule	Source	Destination	Protocol
✕	Apr 16 14:33:31	WAN	Allow IGMP (1670603678)	19[redacted].202	224.0.0.22	IGMP
✕	Apr 16 14:33:36	WAN	Allow IGMP (1670603678)	19[redacted].201	224.0.0.22	IGMP
✕	Apr 16 14:33:37	WAN	Allow IGMP (1670603678)	19[redacted].201	224.0.0.22	IGMP
✕	Apr 16 14:39:09	WAN	Allow IGMP (1670603678)	19[redacted].24	224.0.0.22	IGMP
✕	Apr 16 14:39:11	WAN	Allow IGMP (1670603678)	19[redacted].24	224.0.0.22	IGMP
✕	Apr 16 14:42:33	WAN	Allow IGMP (1670603678)	19[redacted].24	224.0.0.22	IGMP

Fig. 6: Log entries for packets blocked due to the presence of IP options

To resolve this, ensure any rules which should pass packets containing IP Options (e.g. IGMP) are configured to allow IP options.

See also:

IP Options

34.37.3 Asymmetric Routing

Blocked packets are also common for legitimate-looking traffic where routed networks and/or Multi-WAN are involved when *Asymmetric Routing* or other related causes are present in the network.

See also:

- *Bypass Firewall Rules for Traffic on Same Interface*
- *Static Route Filtering*
- *Troubleshooting Asymmetric Routing*

34.37.4 Clustering and Load Balancing

In a clustered environment, traffic arriving via the primary and leaving an internal interface can appear to be blocked on the secondary if the destination is a broadcast or multicast address like those used for Microsoft Network Load Balancing. The traffic appears to be blocked on the internal interface of the secondary from a public IP address source. Capturing the traffic on the secondary and inspecting the destination address in Wireshark will reveal the nature of the destination MAC address.

34.38 Troubleshooting ARP Move Log Messages

Log entries on pfSense® software may appear in the system log showing something similar to the following:

```
pfsense kernel: arp: 192.168.1.50 moved from c4:0c:5c:69:6c:05 to 62:1e:3e:43:04:0c on␣
↪em1
```

This indicates that the firewall saw the specified IP address move between the first MAC address and the second. This can happen for several reasons.

IP address conflict
: Two hosts are configured with the same IP address

ARP poisoning
: Someone on the network is ARP poisoning hosts

NIC teaming
: Some NIC teaming or bonding configurations will routinely log messages such as this because of the way they function. In these cases, this message is normal.

IP address moved to a different host or NIC
: If an actively used IP address is reassigned to a different device or different NIC, this message will be logged. This will only occur when an active IP is moved, for instance an expired DHCP lease that later is assigned to a different host will not trigger this as the IP must have an active ARP table entry on the firewall for this to occur.

Apple Bonjour sleep proxy
: Apple's Bonjour sleep proxy will cause these logs to appear because of its network behavior. If both of the listed MAC addresses are Apple vendor MACs, this is likely why and can be disregarded as normal behavior.

This logging can be disabled by setting the tunable `net.link.ether.inet.log_arp_movements` to value `0` under **System > Advanced**, **System Tunables** tab.

34.39 Troubleshooting "login on console as root" Log Messages

Occasionally, the following messages may appear in the system log:

```
login: login on console as root
```

or:

```
login: login on ttyv0 as root
```

This is normal. It means that the console menu stopped and restarted, or someone pressed `Enter` (didn't choose a menu option) at the console menu. To suppress these messages, enable password protection for the console login, and then it will only log in after authentication. If console logins are already enabled, then this means someone logged into the console.

To password protect the console:

- Click **System > Advanced**
- Find the **Console Options** section near the end of the page
- Check **Password protect the console menu**
- Click **Save**

34.40 Troubleshooting "promiscuous mode enabled" Log Messages

The following log messages are recorded when a utility has placed the network card into promiscuous mode:

```
Feb 10 01:41:58    kernel: igc0: promiscuous mode disabled
Feb 10 01:41:57    kernel: igc0: promiscuous mode enabled
Feb 10 01:41:54    kernel: igc0: promiscuous mode disabled
Feb 10 01:41:54    kernel: igc0: promiscuous mode enabled
Feb 10 01:41:50    kernel: igc0: promiscuous mode disabled
```

Promiscuous mode is a mode where the network card will receive every packet on the interface, regardless of the target MAC address, in order to monitor traffic. This is normal for utilities such as `tcpdump`.

34.41 Troubleshooting Low Interface Throughput

In situations where the firewall is not transferring as much data as desired. There are many potential causes for this condition, most of which are listed here along with possible resolutions

34.41.1 Insufficient Hardware

The first thing to check is that the hardware is capable of pushing the expected amount of traffic. In some cases this is more obvious, such as a newer multicore server being unable to transfer small amounts of packets, or an older firewall not being able to transfer high loads. Other cases are more subtle and require some testing and verification.

The most obvious test is to watch the firewall CPU load while transferring data. This can be observed from **Diagnostics > System Activity** or from the shell by running:

```
top -aSH
```

If an IRQ process for a network card is using a significant amount of CPU on a core, then either the hardware is being fully (or over) utilized, or the driver may need adjustments to work as expected. If the firewall is not under any stress whatsoever while transferring data, the problem likely lies elsewhere.

If the amount of "System" CPU is high and the amount of interrupts is low, the problem may be in the amount of packet processing happening in PF or being used for encryption. If PF is pushing the CPU as high as it can, it may require a faster CPU.

If the CPU is being used for encryption, a faster cipher may be chosen, or in some cases a cryptographic accelerator may be utilized.

34.41.2 Hardware/Driver Tuning Required

If a CPU core is fully utilized by interrupts, the network card driver may need tuning. Most of these tweaks are covered on *Tuning and Troubleshooting Network Cards*. Some cards, such as `igb`, are able to use more queues for processing packets which will spread the load across multiple cores and result in higher throughput, but not every workload is helped by these options, so fewer queues may also help.

Another item to check is under **System > Advanced** on the **Networking** tab. Ensure that the boxes are checked for **Disable hardware TCP segmentation offload** and **Disable hardware large receive offload**. If they are already checked, try toggling **Disable hardware checksum offload**. If no difference is observed, toggle it back.

34.41.3 Duplex Mismatch

A duplex mismatch is also possible, though this is more common on circuits 100Mbit/s or less. Some providers are stuck in the Stone Age and still insist on hard-coding ports on CPEs such as fiber converters at *100Mbit/s full-duplex*. If the CPE is hard-coded, but the firewall is not, it would show as using 100Mbit/s half-duplex on **Status > Interfaces**. The duplex mismatch will lead to interface errors, collisions, and low throughput. Setting the speed and duplex is covered on *Forcing Interface Speed or Duplex Settings*.

34.41.4 Traffic Shaping

If the traffic shaping wizard was run previously before an increase in upstream bandwidth, the old limits may still be in effect. Visit **Firewall > Traffic Shaper** and check the root interface queues, and *qInternet* queues to ensure that any listed interface bandwidths are appropriately specified and current.

Also check the **Limiters** tab under the traffic shaper settings, verify that any configured limiters are set for appropriate speeds. Limiters may also need increased queue lengths to handle higher throughput volumes.

34.41.5 MTU Issues

Issues with upload speed frequently end up being issues with the MTU. If the MTU on pfSense® software (default 1500), is higher than the MTU of the upstream link, it can result in packets being fragmented, lost, or otherwise mishandled. Setting MSS clamping on the WANs or changing the MTU of the interface may help.

VPN + MTU Issues

Similar to the above, if large packets or high-throughput seems to break over a VPN, enable MSS Clamping for VPN Networks under **System > Advanced**, **Firewall & NAT** tab in the **VPN Packet Processing** section. The default value for the option is `1400`, but try lower values such as `1350`, `1300`, `1250`, etc.

34.41.6 WAN Connection

There could also be issues between the WAN and the Modem/CPE. It could be a cable, or a quirk in how the two interfaces talk to each other. Place a small switch between the firewall and the Modem/CPE as a test.

34.41.7 Client/Testing Method

The slowness may not be from any cause on the firewall. It could be the client itself or how it connects. Testing a 100Mbit/s WAN over 802.11g wireless, for example, would never show full speed. Testing a 300Mbit/s WAN from a 100Mbit/s LAN connection would likewise not be a valid test.

Ensure the client is connected to the firewall through a connection at least as fast as the WAN.

34.41.8 ISP Issues

If every other factor has been eliminated, test the modem without the firewall involved. If the speed is still low, it may be the ISP to blame, or the Modem/CPE.

34.42 Troubleshooting Multi-WAN

This section describes some of the most common problems with multi-WAN and how to troubleshoot them.

34.42.1 Verify Firewall Rule Configuration

The most common error when configuring multi-WAN is improper firewall rules. Remember, the first matching rule wins and any further rules are ignored. If a policy routing rule is below the default LAN rule in the list, no traffic will ever match that rule because it will match the default LAN rule first. Review *Policy Routing Configuration* and verify the rules are correct.

If the rule ordering and configuration appears correct, it may help to enable logging on the rules. See *Troubleshooting Firewall Rules* for more information. Ensure the appropriate policy routing rule is passing the traffic.

34.42.2 Policy routing does not work for web traffic or all traffic

When a package that can proxy traffic is used, it overrides any policy routes that are defined for client traffic on that port. So no matter which gateway is set in firewall rules, traffic that uses a proxied connection will still go through the proxy.

34.42.3 Failover not working

If problems occur when an Internet connection fails, typically it is because the monitor IP address is still answering, so the firewall thinks the connection is still available. Check **Status > Gateways** to verify. An IP address on the modem may be used as a monitor IP address, which will still be accessible even if the Internet connection is down.

34.42.4 Load balancing not working

- Check that the Gateway Group is properly configured for load balancing, with at least two gateways on the same tier.
- Check that the firewall rules being matched direct traffic to the correct load balancing gateway group.
- Check that all the gateways in the group show as Online under **Status > Gateways**. Connections marked as Offline will not be used.
- Check the testing methodology. Rather than testing with a web browser, try testing with curl or similar utilities which do not retain session data.
- Check that the traffic is not using a proxy or otherwise being initiated from a daemon on the firewall itself.

34.42.5 A gateway is incorrectly marked offline

If a gateway is listed as offline, but the WAN is actually up, several things could be at fault:

- First, test to see if the monitor IP address responds to a ping from a client device on the LAN, and again from **Diagnostics > Ping**.
- If the device with the monitor IP address or other intermediate hop drops ICMP echo request packets without a payload, manual pings would work, but the gateway monitoring would fail. See *Advanced Gateway Settings* and set the payload to a value of `1` or higher.
- If the gateway or monitor IP address does not respond to ICMP echo requests, enter a different monitor IP address to use instead.
- If the monitor IP address is configured as a DNS server for a different WAN, the static routes could be causing a conflict and the echo requests to the gateway may not be following the expected path. Set a non-conflicting monitor IP address on the gateway.
- If there is an outbound NAT rule on the WAN with a **Source** of *any*, it can cause problems with traffic on the firewall, including monitoring traffic, because that will also NAT traffic from the firewall itself. This can be especially problematic if the source address is changed to a CARP VIP. Fix the outbound NAT.

If all else fails, it's possible the circuit really is down, but the testing methodology appears to show it up. Verify the Interface and Gateway settings and run the test again, and try `traceroute` to make sure the traffic is leaving using the expected path.

34.42.6 Ping works by IP address, but web browsing fails

In this case, the most likely cause is DNS. If the firewall DNS settings do not match those in *Interface and DNS Configuration*, clients may not be able to resolve DNS when a WAN is down. Review the settings and fix any problems that are found.

34.42.7 Services on the firewall do not use multiple connections

Services on the firewall itself do not understand load balancing; They will use only the WAN connection with the default gateway. Configuring default gateway switching (*Managing the Default Gateway*) can allow such services to use failover.

Check the Netgate Forum for package-specific alternate techniques.

34.43 Troubleshooting NAT

NAT can be complex, and in all but the most basic environments there are bound to be issues obtaining an ideal configuration. This section covers a few common problems and potential resolutions.

See also:

Hangouts Archive to view the May 2016 hangout for NAT on pfSense® software version 2.3, The June 2016 hangout on Connectivity Troubleshooting, and the December 2013 Hangout on Port Forward Troubleshooting, among others.

34.43.1 Port Forward Troubleshooting

Port Forwards in particular can be tricky, since there are many things to go wrong, most of which could be in the client configuration and not pfSense software.

For information on diagnosing these problems, see *Troubleshooting NAT Port Forwards*,

34.43.2 NAT Reflection Troubleshooting

For detailed information about troubleshooting NAT reflection, see *Troubleshooting NAT Reflection*.

34.43.3 Outbound NAT Troubleshooting

When manual outbound NAT is enabled and there are multiple local subnets, an outbound NAT entry is required for **each** subnet on each outgoing WAN interface. This includes traffic for subnets that must exit with NAT after coming into pfSense software through a VPN connection.

One indication of a missing outbound NAT rule is packets leaving the WAN interface with a source address of a private network. See *Packet Capturing* for more details on obtaining and interpreting packet captures.

34.43.4 1:1 NAT Troubleshooting

For information about troubleshooting 1:1 NAT, see *Troubleshooting 1:1 NAT*.

34.44 Troubleshooting 1:1 NAT

If 1:1 NAT (or even Outbound NAT) is properly configured, but the system still appears to access sites like https://www.pfsense.org/ip and https://ifconfig.me/ from the main WAN IP Address on the firewall, then something may be proxying or redirecting those requests.

With a proxy involved, even though 1:1 NAT is in place, web requests can still be proxied, and thus originate from the proxy.

To verify the 1:1 mapping is working properly in an environment where web requests are proxied, find a different method to verify, such as:

- Login to a remote system and watch the firewall logs or packet capture.
- Initiate some traffic from the system and verify the traffic is originating from the proper IP Address by checking the states or a packet capture.
- Access an HTTPS site that does not flow through the proxy.

34.45 Troubleshooting NAT Port Forwards

This document covers potential problems users may encounter with port forwarding on pfSense® software.

34.45.1 Port Forward Testing Procedures

Follow the Guide

If the *Port Forwarding* guide was not followed exactly, remove the problematic port forward entries and start from scratch by closely following the instructions in the guide.

NAT Reflection

Port forwards on external interfaces do not work from local clients without *NAT reflection*.

Always test port forwards from outside the network, such as from a client in another location, using a VPN to another region, or from a cellular device.

Setup Logging

Edit the firewall rule that passes traffic for the NAT entry and enable logging. Save and click **Apply Changes**. Then, try to access it again from the outside. Check the firewall logs (**Status > System Logs**, **Firewall** tab) to see if the traffic shows as being permitted or denied.

Check States

Check the states table under **Diagnostics > States**, filter on the source, destination, or port number to see if any entries are present. If entries are present that appear to match the NAT performed by the port forward, then the firewall is accepting and translating the traffic properly, so look at internal issues (e.g. client firewalls).

Check Packet Capture

Use a **Packet Capture** or `tcpdump` to see what is happening on the wire.

See also:

Packet Capturing

This is the best means of finding the problem, but requires the most networking expertise. Navigate to **Diagnostics > Packet Capture** to capture traffic, or use `tcpdump` from the shell.

Start with the WAN interface, and use a filter for the appropriate protocol and port. Attempt to access from outside the network and see if the connection attempt shows in the packet capture. If not, the ISP may be blocking the traffic, or if Virtual IPs are involved they may have an incorrect configuration.

If the traffic is seen on the WAN interface, switch to the inside interface and perform a similar capture. If the traffic is not leaving the inside interface, there is a NAT or firewall rule configuration problem.

If packets leave the internal interface, but no traffic is coming back from the destination host, the target host default gateway may be missing or incorrect, it may not be listening on that port, or it may have a local firewall (Windows Firewall, iptables) blocking the traffic.

With certain types of packets return traffic may indicate the host is not listening on that port. For TCP, this is a TCP RST. For UDP, it can be an ICMP Unreachable message.

34.45.2 Check for Common Problems

This section contains common problems with port forwards and potential resolutions.

Missing or incorrect firewall rule

After checking the port forward settings, double check that the firewall rule has the proper settings. An incorrect firewall rule would also be apparent by viewing the firewall logs (*Viewing the Firewall Log*).

Remember, the destination for the firewall rule is the **internal** IP address of the target system and not the address of the interface containing the port forward. See *Rules for NAT* for more details.

Firewall is enabled on the target machine

pfSense software may be forwarding the port properly, but a firewall on the target host may be blocking the traffic. If there is a firewall on the target host, check its logs and settings to confirm whether the traffic is being blocked at that point.

Incorrect Gateway on Target

For pfSense software to properly forward a port for a local system, pfSense software **must** be the default gateway for the target host.

If pfSense software is not the gateway, the target host will attempt to send replies to forwarded requests out whatever router the target has for its gateway. At that point, one of two things will happen: It will be dropped there since there would be no matching connection state on that router, or it would have NAT applied by that router and then be dropped by the system originating the request since the reply is from a different IP address than the one to which the request was initially sent.

Target system has no gateway or cannot use pfSense software as its gateway

A subset of the larger problem of the target machine gateway is when the device has no gateway, or is incapable of having a gateway. In these cases, work around that problem by switching to Hybrid or Manual Outbound NAT and crafting a rule on the LAN or other internal interface facing the local device. This rule should translate traffic from any source going to the target host on the target port.

For example, if there is a file server without a gateway located at `10.3.0.6`, switch to Hybrid Outbound NAT and create a rule like Figure *Manual Outbound NAT Rule for LAN Device with Missing Gateway* to reach it from outside the network. The file server will see the LAN IP address of the firewall as the source of the traffic, and since that is "local" to the server, it will respond properly.

Target machine is not listening on the forwarded port

If the request is rejected instead of timing out when the connection is tested, in all likelihood pfSense software is forwarding the connection properly and the connection is rejected by the target host. This can happen when the target host has no service listening on the port in question, or if the port being forwarded does not match the port on which the target host is listening.

For example, if the target host is supposed to listen for SSH connections, but the port forward was entered for port `23` instead of `22`, the server would most likely reject the request. The difference can typically be detected by trying to connect to the port in question using `nc` or `telnet`. A message such as "Connection refused" indicates something, frequently the inside host, is actively rejecting the connection.

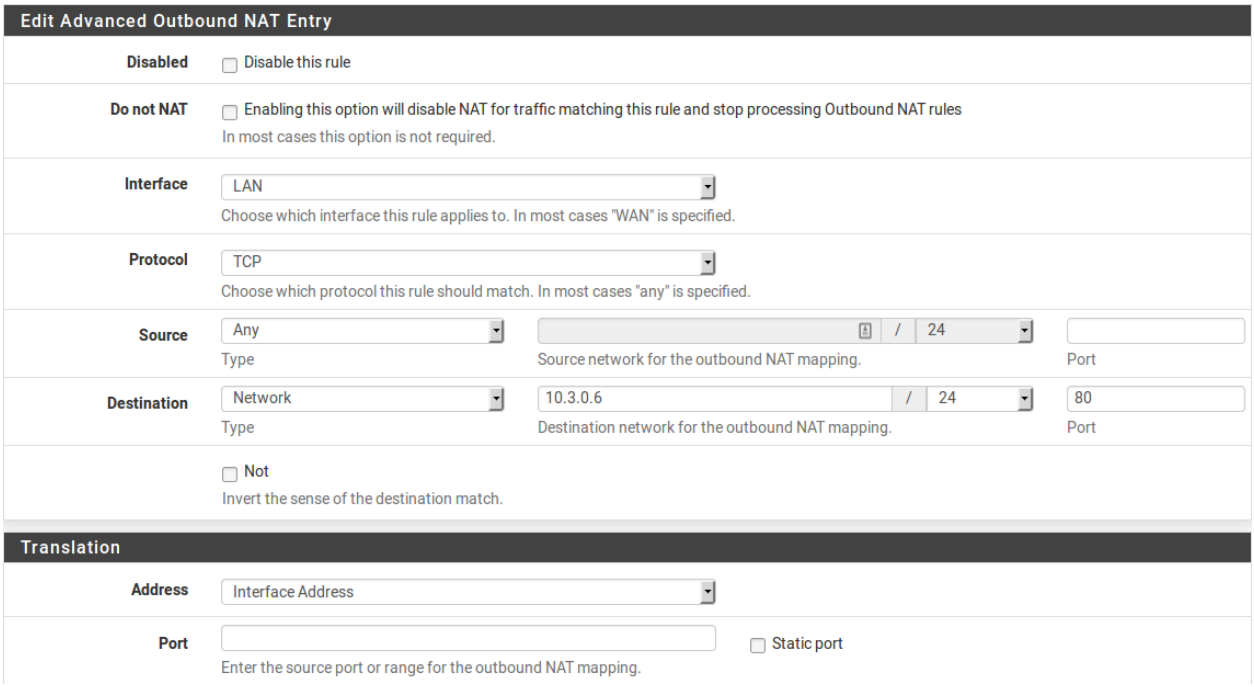


Fig. 7: Manual Outbound NAT Rule for LAN Device with Missing Gateway

See also:

Using **Diagnostics > Test Port** can also help, see *Testing a TCP Port*.

ISP is blocking the port

Some ISPs filter incoming traffic to well-known ports. Check the Terms of Service (ToS) from the ISP to see if there is a clause about running servers. Such restrictions are more common on residential connections than commercial connections. When in doubt, a call to the ISP may clear up the matter.

If the ISP is filtering ports, moving the services to a different port may work around the restriction. For example, if the ISP disallows servers on port `80`, try `8080` or `18080`.

Before attempting to work around a filter, consult the ISP ToS to ensure that running a server is not a violation of their rules.

Testing from inside the network instead of outside

By default, port forwards only work when connections are made from outside the local network. This is a frequent mistake when testing port forwards.

If port forwards must work internally, see *NAT Reflection*. However, Split DNS (*Split DNS*) is the best practice solution to this problem. Split DNS does not rely on NAT reflection or port forwards, and it is worth the effort to implement that instead.

Even with NAT reflection, testing from inside the network isn't necessarily indicative of whether it will work from the Internet. ISP restrictions, restrictions on devices upstream from the firewall, amongst other possibilities, are not possible to see when testing from within the network.

Incorrect or missing Virtual IP address

When using IP addresses that are not the actual IP addresses assigned to an interface, a Virtual IP address is necessary (VIPs, see *Virtual IP Addresses*). If a port forward on an alternate IP address is not working, it may require a different type of VIP. For example, a Proxy ARP type may be necessary instead of an "Other" type VIP.

When testing, also make sure that the client is connecting to the proper VIP.

pfSense software is not the border/edge router

In some scenarios pfSense software is acting as an internal router and there are other routers between it and the Internet also performing NAT. In such a case, a port forward must also be entered on the edge router forwarding the port to pfSense software, which will then use another port forward to get it to the local target host.

Note: Some upstream gear may also be able to change to a bridge mode to eliminate double NAT, or use a half bridge, 1:1 NAT, or DMZ mode to forward all traffic to the firewall running pfSense software.

Forwarding ports to a host behind Captive Portal

Forwarding ports to a host behind a captive portal needs special consideration. See *Port Forwards Behind Portal Only Work When Target Logs In* for details.

Return Routing

There are a few possible issues with return routing on WANs, particularly with multiple WANs.

- If the port forward is on a WAN that is not the default gateway, make sure there is a gateway chosen on the corresponding WAN interface. Otherwise, the firewall rules for the port forward will not send reply packets via the correct gateway.

  The firewall will not add `reply-to` on the rules unless the interface is configured with a gateway.

- If this is on a WAN that is not the default gateway, ensure the traffic for the port forward is **NOT** passed in via Floating Rules or an Interface Group.

  Only rules present on interface tab under Firewall Rules get the `reply-to` keyword to ensure the traffic responds properly via the expected gateway.

- If this is on a WAN that is not the default gateway, make sure the firewall rule allowing the inbound traffic **does not** have the box checked to disable `reply-to`.

- If this is on a WAN that is not the default gateway, make sure the master `reply-to` disable switch is not checked under **System > Advanced**, **Firewall/NAT** tab.

Gateways on Firewall Rules

WAN rules **should NOT** have a gateway set. Make sure that the rules passing traffic for port forwards **do not** have a gateway configured.

Check UPnP IGD & PCP

If the traffic appears to be forwarding in to an unexpected host, it may be happening due to UPnP IGD & PCP. Check **Status > UPnP IGD & PCP** to see if an internal host has configured a port forward unexpectedly. If so, disable UPnP IGD & PCP on either that host or on the firewall.

34.46 Troubleshooting NAT Reflection

NAT Reflection (*NAT Reflection*) is complex, and as such may not work in some advanced scenarios. The best practice is to use Split DNS instead (*Split DNS*). That said, NAT Reflection on current pfSense software releases works reasonably well for nearly all use cases, and problems usually stem from configuration mistakes. Ensure NAT reflection is enabled correctly, and ensure a large range of ports is not being forwarded unnecessarily.

NAT Reflection rules are also duplicated for each interface present in the system, so if a lot of port forwards and interfaces are in use, the number of reflectors can easily surpass the limits of the firewall. If this happens, an entry is printed in the system logs. Check the logs for any errors or information.

34.46.1 Web Access is Broken with NAT Reflection Enabled

If an improperly specified NAT Port Forward is present on the firewall, it can cause problems when NAT Reflection is enabled. The most common way this problem arises is with a local web server, and port `80` is forwarded there with an improperly specified **External Address**.

If NAT Reflection is enabled and the **External Address** is set to *any*, any connection made on the firewall comes up as the local web server. To fix this, edit the Port Forward for the offending port, and change **External Address** to *Interface Address* instead.

If an external address of *any* is required, then NAT Reflection will not work, and Split DNS must be used instead.

34.47 Troubleshooting OpenVPN

This section describes several troubleshooting techniques for OpenVPN, as well as common issues users encounter with OpenVPN along with their solutions.

34.47.1 Check OpenVPN Status

The first place to check is **Status > OpenVPN**, which displays the connection status for each OpenVPN instance.

If a VPN is connected, waiting, reconnecting, etc., it is indicated on that screen.

See also:

For more information, see *OpenVPN Server and Client Status*.

34.47.2 Check Firewall Log

If a VPN connection does not establish, or establishes but does not pass traffic, check the firewall logs under **Status > System Logs** on the **Firewall** tab.

If traffic for the tunnel itself is being blocked, such as traffic to the WAN IP address on port `1194`, then adjust the WAN firewall rules accordingly.

If traffic is blocked on the OpenVPN interface, add rules to the **OpenVPN** tab (or assigned OpenVPN interface tab, if present) to allow traffic there.

34.47.3 Some hosts work, but not all

If traffic between some hosts functions over OpenVPN, but some hosts do not, this is commonly one of four things:

Firewall rules
: Ensure the rules in the firewall GUI on both sides allow the desired network traffic. This may be on the **OpenVPN** tab in **Firewall > Rules** or an assigned OpenVPN interface tab.

Missing, incorrect or ignored default gateway
: If the host that cannot be reached across the VPN does not have a default gateway, or has one pointing to something other than the firewall running OpenVPN, the host does not know how to properly get back to the remote network on the VPN.

  Additionally, some devices, even with a default gateway specified, do not use that gateway. This has been seen on various embedded devices, including IP cameras and printers. In this case, there is no proper firewall-based solution. The software on the device must be fixed.

  This behavior can be verified by running a packet capture on the inside interface of the firewall connected to the network containing the host. If the packet capture contains traffic leaving the inside interface on the firewall, but not replies to that traffic, the device is not properly routing its reply traffic or potentially blocking it via local firewall on the device.

  See also:

  Troubleshooting with tcpdump is covered in *Using tcpdump on the command line*.

Incorrect subnet mask
: If the subnet in use on one end is `10.0.0.0/24` and the other is `10.254.0.0/24`, and a host has an incorrect subnet mask of `255.0.0.0` or `/8`, it will never be able to communicate across the VPN because it thinks the remote VPN subnet is part of the local network and hence routing will not function properly.

  Fix the incorrect subnet mask and then two-way communication is possible.

Host firewall
: If there is a firewall on the target host, it may not be allowing the connections. The Windows firewall, for example, will not allow connections from outside its local subnet in certain configurations.

34.47.4 Check the OpenVPN logs

The OpenVPN logs contain details about the OpenVPN processes, including log messages relating to connections attempts, remote access login records, and other related messages. These logs are in the GUI under **Status > System Logs**, **OpenVPN** tab.

When an OpenVPN instance connects it will log messages similar to the following:

```
openvpn[32194]: UDPv4 link remote: 1.2.3.4:1194
openvpn[32194]: Peer Connection Initiated with 192.168.110.2:1194
openvpn[32194]: Initialization Sequence Completed
```

Note: The number following `openvpn` will differ, it is the process ID of the OpenVPN process handling the connection.

If the `link remote` and `Peer Connection Initialized` messages are not shown when trying to connect, the cause is likely either incorrect client configuration, so the client is not attempting to connect to the correct server, or incorrect firewall rules blocking the client's connection.

34.47.5 Overlapping IPsec connections

The way IPsec configures security policies in the kernel, any enabled IPsec connection matching the local and remote subnets pairs will restrict the firewall to only passing traffic for that pair through IPsec. The traffic will not pass through any other interface, including OpenVPN. Disable any IPsec connections which specify the same local and remote networks as another VPN.

If an IPsec tunnel has been recently disabled or removed, check if the security policy database (SPD) entries are still present at **Status > IPsec** on the **SPD** tab.

If the SPD entries for a disabled tunnel are still active, stop the IPsec daemon and then start it again.

34.47.6 Check the system routing table

Browse to **Diagnostics > Routes** and review the contents of the routing table. For site-to-site VPNs, routes will be present for the remote network(s) to the appropriate *tun* or *tap* interface. If the routes are missing or incorrect, the **Local Network**, **Remote Network**, or custom options are not configured correctly.

If a peer-to-peer setup is in use (shared key or SSL/TLS with `/30` tunnel network) and not client/server, ensure that the custom options do not contain any "push" commands. In a peer-to-peer setup, the peers cannot push settings to each other. Both sides must contain appropriate remote network configurations.

34.47.7 Test from different vantage points

If the connection appears to be up according to the logs, but it does not work from a host on the LAN, try it from the firewall itself. These tests may be easily performed by the **Diagnostics > Ping** page in the GUI.

See also:

- *Ping Host*
- *OpenVPN Adapter Address ICMP Behavior*

First test using the inside interface involved in handling OpenVPN internal traffic as the ping source. This is typically the LAN interface. If that does not work, try again using the *default* source address so that the firewall will source the ping from the OpenVPN interface itself.

If the *default* source ping works, but the internal network ping does not, check the firewall rules and routes on the far side.

34.47.8 Trace the traffic with packet captures

Using packet captures to determine traffic flow is one of the most helpful troubleshooting techniques.

Start with the internal interface (commonly LAN) on the side where hosts initiate the traffic. Next, progress to the *tun* interface on that firewall, then the *tun* interface on the remote firewall. Finally, check the inside interface on the remote firewall. Determining which interfaces contain the traffic which do not can help greatly in narrowing down the location of the problem.

See also:

Packet Capturing

34.47.9 Routes will not push to a client

If a client does not receive routes for networks from the **Local Network** settings or a `push` statement, a couple of things could be happening:

- Check that an SSL/TLS server setup is used with a **Tunnel Network** larger than a `/30`. The `server` mode in OpenVPN only takes effect when using a subnet large enough to contain multiple clients, such as a `/24`.
- If the client is running on Windows 10 or similar, try running the client as Administrator. Some versions of the OpenVPN client require Administrator mode to apply routes to the client PC routing table.

  Note: This should not be a factor on modern versions of the OpenVPN client for Windows, so another step is to uninstall the OpenVPN client on the end user device and install the most current available client.

- Peer-to-peer setups (shared key, SSL/TLS with a `/30` tunnel network) are not capable of pushing routes. Use the **Remote Network** boxes or `route` statements on each side (both client and server) to direct traffic to subnets on the far end of the tunnel.

34.47.10 Why do multiple OpenVPN clients get the same IP address?

If all clients use the same certificate then OpenVPN will assign all clients the same IP address when they connect. To work around this, create separate certificates for each client, which is the best practice.

Another workaround is to check **Duplicate Connections** on the server configuration. This is a bad practice, however, as it is less secure than using unique certificates for each client.

34.47.11 Importing OpenVPN DH Parameters

When importing an existing OpenVPN setup into pfSense, there is no need to import DH Parameters. DH parameters are not specific to a given setup in the way that certificates or keys are.

See also:

DH Parameters Length

34.48 Troubleshooting Windows OpenVPN Client Connectivity

Historically, OpenVPN client software on Windows had issues with routing due to a lack of privileges. Current versions of the OpenVPN client software for Windows run as a service which only requires administrative privileges during the installation process and **not** when the client software runs afterward.

Older legacy client software versions may still need such privileges, but in nearly all cases the correct course of action is to upgrade the software.

See also:

- *Installing the OpenVPN Client on Windows*
- *Routes will not push to a client*
- *OpenVPN Client Export Package*

When installing OpenVPN client software using the *OpenVPN Client Export Package*, the installer will not upgrade an existing installation of OpenVPN client software. To ensure the client software is updated properly, uninstall the old client software first, reboot, then install the new copy.

34.49 Troubleshooting OpenVPN Internal Routing (iroute)

For a site-to-site PKI (SSL) OpenVPN setup with a tunnel network larger than `/30`, OpenVPN must have an internal route for the client subnet. Without the internal route, the firewall will forward traffic into OpenVPN but OpenVPN will drop the traffic as it has no way to determine the proper destination. There are a couple common scenarios where this may have difficulties.

The **Remote Network** (`route`) definitions on the server settings inform the firewall operating system that the networks must be routed to an OpenVPN instance. The **Remote Network** (`iroute`) options on the **Client Specific Override** inform OpenVPN internally which networks are associated with a specific client certificate.

See also:

- *Client Specific Overrides*
- *OpenVPN Site-to-Site Configuration Example with SSL/TLS*

34.49.1 Check Internal Route Configuration

Internal routes are set by a **Client Specific Overrides** entry which matches the client certificate **common name**. In the override, the **IPv4/IPv6 Remote Network/s** boxes can set up this internal routing, or it can be performed manually using an `iroute` statement in the advanced settings.

The firewall creates `iroute` statements automatically for each network listed in the **IPv4** and **IPv6 Remote Network/s** fields of an override.

Next, ensure that the **common name** matches and that OpenVPN is learning the internal route as it should be. Log verbosity in OpenVPN may need increased to see if this is working. On **Status > OpenVPN** the internal routing for the OpenVPN server may also be viewed while the client is connected.

For each network that needs an `iroute` statement, the **server** definition must also have the same network(s) listed as **IPv4/IPv6 Remote Networks** or as `route` statements in the **Custom options** box.

34.49.2 Example Configuration

This is a basic example which demonstrates routing a single subnet to a specific client.

- Server1 custom options:

```
push "route a.a.a.0 255.255.255.0";
route b.b.b.0 255.255.255.0;
```

- Client Specific Overrides for **Common Name** `client1`:
  - **IPv4 Remote Network/s** set to `b.b.b.0/24` **OR**
  - Advanced options:

```
iroute b.b.b.0 255.255.255.0;
```

client1 custom options:

```
(blank -- no route statements needed)
```

34.49.3 Single Client Strategy Without Internal Routing

For a site-to-site setup between only two locations, the tunnel network can be a `/30` so that OpenVPN uses peer-to-peer mode and does not require `iroute` statements to reach client networks. In this case, use the **IPv4/6 Remote Network/s** on both sides to set up the routes and there is no need for an override.

See also:

See the note at *IPv4/IPv6 Tunnel Network* for more information.

34.50 Troubleshooting Lost Traffic or Disappearing Packets

If there are issues with traffic being lost, or packets that seem to disappear or never show up (or leave) an interface, there are a few potential causes to consider.

See also:

Troubleshooting Network Connectivity contains procedures to diagnose and analyze these types of problems in great detail.

34.50.1 IPsec (Tunnel Mode)

If a packet matches the traffic selectors set in tunnel mode IPsec phase 2 entries exactly (matches both source and destination), then the kernel will attempt to inject that packet into IPsec even if the IPsec tunnel is down.

This is commonly observed by users who set up an alternate VPN (e.g. WireGuard or OpenVPN) for the same networks previously handled by IPsec without fully removing the IPsec configuration.

Check the contents of the Security Policy Database (SPD) at **Status > IPsec** on the **SPDs** tab to see if one of the policies listed there overlaps the network that isn't working properly. If it does, remove or disable the IPsec configuration and stop the IPsec daemon, then start it again. In some rare cases it may require a reboot to fully clear the old policy.

34.50.2 Captive Portal

If a host is on a network segment in an active Captive Portal zone, then it must be authenticated to the portal or setup as a bypass (e.g. Allowed IP/Host/MAC address) to pass traffic through the tunnel. This includes traffic from inbound NAT such as port forwards or 1:1 NAT.

34.50.3 Firewall Rules

Firewall rules, including floating rules, could be blocking and not logging the block action. Check the floating tab for matching rules (e.g. from packages such as pfBlockerNG) and also look at the content of block tables maintained by packages such as Snort and Suricata.

34.50.4 Routing Problems

If the firewall has a packet but nowhere to deliver the packet, then the firewall can drop that packet. The most common way this happens is from a lack of default route on the firewall itself.

If a packet arrives for a network that is not on a directly connected interface and the firewall has no default route, then the firewall has no idea where it should send that packet. This, it has no choice but to drop.

This can be sidestepped in certain cases by policy routing, which can make it tricky to diagnose at first. For example if the pass rules on LAN all have a gateway set then traffic from LAN might work, but traffic from the firewall itself (e.g. DNS resolution) would fail.

34.50.5 Hardware Checksum Offloading

It's possible that a problem in hardware checksum offloading is leading to the packets being rejected by various parts of the network (e.g. OS, NIC, switch, peers, etc.)

Try disabling Checksum Offloading as follows:

- Navigate to **System > Advanced**, **Networking** tab
- Check **Disable hardware checksum offload** under **Network Interfaces**
- Click **Save**

Then try to reproduce the problem.

Note: A reboot may be desired after applying this option, but it should not be necessary.

This has historically been an issue with Realtek NICs, some Via Rhine NICs, and some specific Intel `fxp` chips, as well as virtualized/emulated NICs in some hypervisors, especially those that use VirtIO.

The problem may also manifest as a PPPoE connection that will establish a login and bring up the interface, but will not pass traffic.

34.51 Troubleshooting a Broken pkg Database

In rare edge cases it is possible for the `pkg` database in `/var/db/pkg/` to become corrupted. In the unlikely event this happens to a firewall, it can usually be corrected by running a few commands to re-create the database.

Note: The following commands only account for the base system of a typical CE installation.

- Ensure that the package database directory exists:

```
# /bin/mkdir -p /var/db/pkg/ /root/var/db/pkg/
```

- Force an update of the package repository data:

```
# /usr/local/sbin/pkg-static update -f
```

- Force a reinstallation of the pfSense® software base package and kernel:

```
# /usr/local/sbin/pkg-static install -yf pkg pfSense pfSense-kernel-pfSense
```

- Refresh the `php.ini` and other files to ensure they are loading the correct modules:

```
# /etc/rc.php_ini_setup
```

- If any additional packages were installed, reinstall them manually using the GUI if possible, or by using `pkg-static install` as above:

```
# /usr/local/sbin/pkg-static install -yf <additional-package> <another-additional-
↪package>
```

34.52 Troubleshooting Routes

When diagnosing traffic flow issues, one of the first things to check is the routes known to the firewall.

34.52.1 Viewing Routes

Viewing the routing table is described in detail at *Route Table Contents*. For routed destinations, check the contents of the table and see if the destination is listed as expected.

It is also possible to check how the operating system will route a specific destination at the CLI. For example, to see how the firewall will reach `8.8.8.8`:

```
$ route -n get 8.8.8.8
   route to: 8.8.8.8
destination: 0.0.0.0
       mask: 0.0.0.0
    gateway: 198.51.100.1
        fib: 0
  interface: igb0
      flags: <UP,GATEWAY,DONE,STATIC>
 recvpipe  sendpipe  ssthresh  rtt,msec    mtu        weight    expire
       0         0         0         0      1500         1         0
```

In that example, there isn't a specific static route in place for the destination, so it uses the default route.

In the next example, there is a specific route for `8.8.4.4` using an alternate gateway:

```
: route -n get 8.8.4.4
   route to: 8.8.4.4
destination: 8.8.4.4
    gateway: 10.0.14.1
        fib: 0
  interface: igb5
      flags: <UP,GATEWAY,HOST,DONE,STATIC>
 recvpipe  sendpipe  ssthresh  rtt,msec    mtu        weight    expire
       0         0         0         0      1500         1         0
```

While viewing the routing table as a whole is helpful, sometimes querying the OS in this way is faster and easier when a specific destination is known.

34.52.2 Using traceroute

Traceroute is a useful tool for testing and verifying routes and multi-WAN functionality, among other uses. It shows each "hop" along the path a packet travels from one end to the other, along with the latency encountered in reaching that intermediate point. On pfSense® software, a traceroute can be performed by navigating to **Diagnostics > Traceroute**, or by using `traceroute` at the command line. From clients running Windows, the program is available under the name `tracert`.

See also:

For details on how traceroute works and how to perform a traceroute using the GUI, see *Traceroute*.

34.52.3 Routes and VPNs

Depending on the VPN being used, the firewall may not have routes in the table for the remote side. IPsec in tunnel mode does not use the routing table, it is instead handled internally in the kernel using IPsec security policy database (SPD) entries. Static routes will never cause traffic to be directed across a tunnel mode IPsec connection. VTI mode IPsec, OpenVPN, and WireGuard use the system routing table and as such entries are present for networks reachable via those types of VPNs.

For example, the following output includes an OpenVPN tunnel:

```
#netstat -rWn
Routing tables

Internet:
Destination        Gateway            Flags       Use    Mtu      Netif Expire
default            198.51.100.1       UGS       92421   1500        em0
10.6.0.0/16        10.6.203.1         UGS           0   1500     ovpnc2
10.6.203.0/24      10.6.203.2         UGS           0   1500     ovpnc2
10.6.203.1         link#9             UH            0   1500     ovpnc2
10.6.203.2         link#9             UHS           0  16384        lo0
10.7.0.0/24        link#2             U       1260771   1500        em1
10.7.0.1           link#2             UHS           0  16384        lo0
127.0.0.1          link#7             UH          866  16384        lo0
198.51.100.0/24    link#1             U       1251477   1500        em0
198.51.100.7       link#1             UHS           0  16384        lo0
```

The OpenVPN interface is `10.6.203.2`, with a gateway of `10.6.203.1` and the interface is `ovpnc2`. The network reachable using OpenVPN in this example is `10.6.0.0/16`.

With tunnel mode IPsec, traceroute is not as useful as with routed setups, because a tunnel mode IPsec connection does not have an interface or IP addresses. When running traceroute to a destination across tunnel mode IPsec, the client will experience a timeout when crossing the IPsec tunnel.

34.53 Troubleshooting in Single User Mode

Warning: This option must only be used under the guidance of a Netgate TAC representative or a FreeBSD user with advanced knowledge.

Single user mode is a special boot mode in which only the operating system kernel is loaded and a bare minimum of functions are available. For example, networking is not configured, the GUI is not available, and no system services are started. The root filesystem defaults to read-only and other filesystems are not mounted. In single user mode, commands must be entered into a shell prompt on the primary system console.

Warning: This action **requires** access to the primary system console, which depending on hardware and configuration may be a video console or serial console. Ensure console access is available and working before attempting to enter single user mode.

34.53.1 Entering Single User Mode

There are multiple ways to enter single user mode:

Next Boot Configuration
From the console or SSH menu, choose option `5` *Reboot System* and then enter `S` (capital `s`) to reboot and automatically enter single user mode.

> **Warning:** ZFS systems will not automatically clear this configuration when exiting single user mode. See *Re-mount ZFS Volumes as Read/Write*.

Choose Single User Mode on Boot Menu
From the loader menu with the ASCII logo, choose the menu option to enter single user mode, typically option `2`

Use the Loader Prompt
Enter the loader directly, either by using option `3` from the ASCII loader menu or if the boot menu is unavailable, press the `Space` bar while the kernel is loading at boot time, which will start the bootloader with an `OK` or `loader>` prompt.

Note: The correct moment to hit `Space` varies by hardware.

The OS may display prompt such as "Hit [Enter] to boot immediately, or any other key for command prompt." at the proper moment to press `Space`. If that message is not present, press `Space` when "Loading kernel…" and/or the kernel loading spinner are visible.

Do not press a key at the `Hit any key to stop autoboot` prompt as that is too early. That prompt is for U-Boot or similar and is before the OS boot process starts.

Enter the following command at the loader prompt:

```
boot -s
```

Warning: Check the prompt again before typing this command. It must say either `loader>` or `OK`. If the prompt shows a different string such as `Marvel>>` that is not the correct prompt for this action. Reboot and try again.

After invoking one of those options, the kernel will load and then immediately after, the operating system will prompt for a shell:

```
Enter full pathname of the shell or RETURN for /bin/sh:
```

At that point, press `Enter` and the resulting shell prompt is now running in single user mode.

Tip: Reminder: At this point, all filesystems are still read-only.

34.53.2 Exiting Single User Mode

There are two ways to exit single user mode, and the method to use depends on the changes made.

Reboot
Run `/sbin/reboot` or an equivalent command to force an operating system reboot. This is the safest choice as it will ensure the system is fully reinitialized. This method is required when making boot-time changes such as those in *Managing Loader Tunables*.

Continue Boot
Exiting the single user mode shell via `exit`, `logout`, or `^D` (Ctrl-D) will terminate the single user mode shell and then continue to boot the system into its regular multi-user mode. This may be OK for simple changes such as `config.xml` alterations, but will not activate changes which must be present before the kernel loads.

34.53.3 Working in Single User Mode

Mounting File Systems

In single user mode the root file system defaults to read-only and other filesystems are not mounted.

Re-mount UFS Volumes as Read/Write

To mount all UFS type filesystems in read-write mode, run the following command:

```
# /sbin/mount -a -t ufs
```

Re-mount ZFS Volumes as Read/Write

For ZFS the procedure to mount the all filesystems as read-write in single user mode requires several commands.

First, remount the root slice as read-write:

```
# /sbin/mount -u /
```

If that produces an error, try:

```
# /sbin/zfs set readonly=off pfSense/ROOT/default
```

Next, mount all ZFS filesystems, datasets, etc.:

```
# /sbin/zfs mount -a
```

On newer versions, such as pfSense® plus software installations with ZFS boot environments, this may not mount all the ZFS datasets as a few are marked not to mount automatically. In this case, mount them manually:

```
# /sbin/zfs mount pfSense/ROOT/default/cf
# /sbin/zfs mount pfSense/ROOT/default/var_db_pkg
```

Finally, if the system was put into single user mode using the reboot menu `S` option, the `nextboot` configuration **must** be cleared manually:

```
# /sbin/nextboot -D
```

Warning: Failing to clear the `nextboot` configuration will result in the system booting back into single user mode on every boot.

Tip: It is safe to run this command even if the system was not booted in that way, so always run this command on ZFS systems before exiting single user mode.

Running Commands

Be aware that when running commands single user mode, their behavior may be unexpected as the system does not have a full traditional terminal setup, and some default shell environment variables and console/terminal settings are not available. For example, using `vi`, the arrow keys on the keyboard may not work or the dimensions of the terminal may be incorrect.

Some commands may be run without typing the full path name to the binary, but it is safest to use the full path where possible.

34.54 Troubleshooting Snort Rule Updates

34.54.1 MD5 Signature Mismatch

Periodically, Sourcefire redesigns their site or updates the engine and rules, and the snort package needs an update to accommodate this change. Removing and then installing the snort package again is required to restore proper functionality, assuming the package has been updated to match the upstream rule format.

34.54.2 Upstream Issues

Rule problems can almost always be solved by waiting 20-30 minutes and then trying the download again. Failing that, uninstall the package completely and then reinstall the package to ensure the snort binaries are the latest/correct ones.

34.54.3 Space Issues

If the `/tmp` slice is small, because the firewall is running with `/tmp` on a RAM disk, current rulesets can easily fill the slice up and cause numerous rule-related errors.

If there is sufficient RAM, increase the size of `/tmp` using the options on **System > Advanced**, **Miscellaneous** tab.

34.55 Troubleshooting the Squid Package

Danger: The add-on packages Squid, SquidGuard and Lightsquid are deprecated in pfSense Plus and pfSense CE software due to many unfixed upstream security vulnerabilities. Netgate **STRONGLY** recommends that users uninstall these packages. The packages will no longer function in future major releases of pfSense Plus and pfSense CE software.

34.55.1 Disk Usage Issues

The `swap.state` from Squid file can grow large and consume all available drive space. See *Tuning the Squid Package* for more details.

34.55.2 Sites not loading with splice / Error 409 in access log

As a security measure, squid will not allow a user to connect to a site that has a hostname that does not match its IP address. This prevents clients from hardcoding or altering DNS responses to evade access controls. The side effect of this, however, is that sites which employ round-robin DNS or other DNS optimizations can cause squid to block or drop connections those sites unintentionally. The squid access log will have a `409 (Conflict)` error code when a connection is dropped for this reason.

This happens with sites such as Google or Facebook when the client and squid use different sources for DNS, and thus get different DNS results for the same query because the results are randomized. Even though the address for the server is valid, the disparity causes squid to drop the connection.

The solution is to have the clients use the firewall as their DNS server, so that both squid and clients use the same DNS source and the results will match.

34.55.3 Clear Cache

Resetting the cache in squid can often clear up issues without performing a more complicated procedure. Before performing a full reset, try clearing and resetting the cache:

```
# mv /var/squid/cache /var/squid/cache.old
# squid -z
# rm -rf /var/squid/cache.old
```

The old cache should be moved, then reset, and then the old cache should be removed, as above, because removing the cache directory can be time-consuming, and if it is moved first, then its removal will not prevent squid from being run while it is happening.

34.55.4 Complete Reset

When troubleshooting squid/squidGuard there are some procedures that may be followed to ensure things are completely reset.

- **Remove** the packages from **System > Packages** on the **Installed Packages** tab in the proper order:
  1. Lightsquid
  2. squidGuard
  3. Squid
- Remove the contents of the squid directory and cache:

```
# rm -rf /var/squid
```

- Remove the Squid and related package include files:

```
# rm /usr/local/pkg/*squid*
# rm -rf /usr/local/etc/squid/
```

- Ensure any leftover PBI symlinks are removed:

```
# find / -type l -lname '/usr/pbi/*' -delete
```

- [Optional] Remove the settings from inside `config.xml` using one of the following methods:
  - From **Diagnostics > Command Prompt** in the **PHP Execute** box:

```
$squid_sections = array("squid", "squidnac", "squidcache", "squidauth",
 "squidextauth", "squidtraffic", "squidupstream", "squidusers");
foreach ($squid_sections as $sec) {
 if (is_array($config['installedpackages'][$sec]))
        unset($config['installedpackages'][$sec]);
}
write_config("Removed Squid");
```

  - Or to remove Squid, squidGuard, Lightsquid, and anything else with `squid` in its package name from **Diagnostics > Command Prompt** in the **PHP Execute** box:

```
foreach (array_keys($config['installedpackages']) as $sec) {
 if (strpos($sec, "squid") !== false)
```

(continues on next page)

(continued from previous page)

```
        unset($config['installedpackages'][$sec]);
}
write_config("Removed all squid-related settings");
```

  - Or backup `config.xml`, edit the settings out, then restore.

- Navigate to **System > Packages** and on the **Available Packages** tab, reinstall the following packages in order:
  1. Squid
  2. squidGuard
  3. Lightsquid

See also:

For assistance in solving problems, post on the Cache/Proxy category of Netgate Forum.

34.56 Troubleshooting Hardware Shutdown and Power Off

If a firewall device does not automatically power itself off, this is typically a case of FreeBSD and ACPI not working well together on a particular hardware combination.

As a test, enter this at the CLI then attempt a power-down:

```
# sysctl hw.acpi.disable_on_reboot=1
```

For a more permanent solution, add an entry under **System > Advanced** on the **Tunables** tab to set:

```
hw.acpi.disable_on_reboot=1
```

34.57 Troubleshooting Clock Issues

Time and clock issues are relatively common on hardware, but on firewalls they are critical, especially if the firewall is performing tasks involving validating certificates as part of a PKI infrastructure.

Proper time synchronization is an absolute necessity on systems which do not have a battery onboard to preserve their date and time settings when power is removed.

Not only will getting this all in line help with critical system tasks, but it also ensures that the log files on the firewall are properly timestamped, which aids with troubleshooting, record keeping, and general system management.

34.57.1 Time Keeping Problems

Hardware can have significant problems keeping time, though such problems are generally isolated to older, low-quality hardware. All PC clocks will drift to some extent, but some hardware can drift as much as one minute for every couple minutes that pass and rapidly get out of sync. NTP is designed to periodically update the system time to account for normal drift. It cannot reasonably correct clocks that drift significantly. This is very uncommon, but there are a few methods that can potentially work around these problems.

The best way to avoid time keeping problems is to use quality hardware that has been tested and does not experience these issues, such as official appliances found in the Netgate Store.

There are several items to check if hardware has significant time keeping problems.

Network Time Protocol

By default, pfSense software attempts to synchronize its time using the ntp.org Network Time Protocol (NTP) server pool. This ensures an accurate date and time on the firewall, and will accommodate normal clock drift. If the firewall date and time are incorrect, ensure NTP synchronization is functioning. The most common problem preventing synchronization is the lack of proper DNS configuration on the firewall. If the firewall cannot resolve hostnames, NTP synchronization will fail. The results of synchronization are shown at boot time in the system log, and the status of the NTP clock synchronization can be viewed at **Status > NTP**. The **NTP Status** widget for the Dashboard also offers basic information about the NTP server selected for use by the firewall.

BIOS Updates

Netgate has observed older hardware that ran fine for years on Windows encounter major timekeeping problems once redeployed on FreeBSD (and by consequence, pfSense software). The systems were running a BIOS version several revisions out of date. One of the revisions addressed a timekeeping issue that apparently never affected Windows. Applying the BIOS update fixed the problem. The first thing to check is to make sure the hardware in question has the latest available BIOS.

PNP OS settings in BIOS

Other hardware might have time keeping difficulties in FreeBSD and pfSense software unless **PNP OS** in the BIOS was set to *No*. If the BIOS does not have a **PNP OS** configuration option, look for an **OS** setting and set it to *Other*.

Disable ACPI

A few BIOS vendors have produced ACPI (Advanced Configuration and Power Interface) implementations which are buggy at best and dangerous at worst. On more than one occasion Netgate has encountered hardware that would not boot or run properly while ACPI is active.

While ACPI can be disabled in the BIOS, and in the OS, this is not a best practice and as such hardware that requires such changes should be replaced.

Adjust Time Counter Hardware Setting

On rare systems, the `kern.timecounter.hardware` sysctl value may need to be changed to correct an inaccurate clock. This is known to be an issue with older versions of VMware such as ESX 5.0 in combination with an amd64-based pfSense software or FreeBSD image. That special case was a bug in the hypervisor that is fixed in ESX 5.1 and later.

On these systems the default timecounter will eventually stop the clock from ticking, causing problems with encryption, VPNs, and services in general. On other systems, the clock may skew by large amounts with the wrong timecounter.

To change the timecounter, browse to **System > Advanced**, on the **System Tunables** tab and add an entry to set `kern.timecounter.hardware` to `i8254`

This will make the system use the i8254 timecounter chip, which typically keeps good time but may not be as fast as other methods. The other timecounter choices will be explained later in this section.

If the system keeps time properly after making this change, leave the tunable entry in place to make this change permanent. If the change did not help, remove the tunable or try another value.

Depending on the platform and hardware, there may also be other time counters to try. For a list of available time counters found on a firewall, execute the following command:

```
# sysctl kern.timecounter.choice
```

The firewall will print a list of available time counters and their "quality" as reported by FreeBSD:

```
kern.timecounter.choice: TSC-low(1000) ACPI-safe(850) i8254(0) dummy(-1000000)
```

Try any of those four values for the sysctl `kern.timecounter.hardware` setting. In terms of "quality" in this listing, the larger the number the better, but the actual usability varies from system to system.

TSC
: A counter on the CPU, but is tied to the clock rate and is not readable by other CPUs. It can be used in bare metal SMP systems, but it requires that TSCs on all CPUs be synchronized. It cannot be used reliably on systems with variable-speed CPUs or virtualized system with multiple CPUs.

i8254
: A clock chip found in most hardware, which tends to be safe but can have performance drawbacks.

ACPI-safe
: If this option is properly implemented in the hardware, this is a good choice because it does not suffer from the performance limitations of i8254, but in practice its accuracy and speed vary widely depending on the implementation.

ACPI-fast
: A faster implementation of the ACPI timecounter available on hardware that does not suffer from known ACPI issues.

HPET
: High Precision Event Timer available in some hardware. When available, it is generally considered a good source of accurate time counting.

This and more information on FreeBSD Timecounters can be found in the paper Timecounters: Efficient and precise timekeeping in SMP kernels by Poul-Henning Kamp of the FreeBSD Project, and in the FreeBSD source code.

Adjust the Kernel Timer Frequency

In rare cases adjusting the kernel timer frequency, or `kern.hz` kernel tunable, can help performance or stability. This is especially true on virtualized environments. The default is `1000`, but in some cases `100`, `50`, or even `10` will be a better value depending on the system. When pfSense software is installed in VMware, it detects it and automatically sets this tunable to `100`, which should work fine in nearly all cases with VMware products.

To adjust this setting, add the following line, with the new value, as a *Loader Tunable*:

```
kern.hz=100
```

34.57.2 GPS Time Synchronization

To aid in maintaining an accurate clock, GPS time synchronization is also provided by pfSense software on compatible hardware. In combination with external time servers, compatible GPS devices can help keep the clock accurate. For more detail, see *NTPD*.

34.58 Troubleshooting Time Zone Configuration

34.58.1 Processes use different time zones

Processes on FreeBSD (and thus pfSense® software) only pick up time changes when they are started. If the firewall has not been rebooted since the last time zone change, doing so will ensure that all running processes are using the correct time zone.

34.58.2 Clock does not use the expected zone offset

If the clock is several hours off, but accurate to the minute, it is most likely a time zone setting issue. If using a GMT offset time (e.g. *-0500*), change to a more specific **geographic** time zone such as *America/New_York* instead.

Using geographic zones is the best practice as they use an accurate offset, include local Daylight Saving Time behavior, and also consider historical changes in time zones.

The following text from the time zone database explains the behavior of the GMT zones further:

```
# We use POSIX-style signs in the Zone names and the output abbreviations,
# even though this is the opposite of what many people expect.
# POSIX has positive signs west of Greenwich, but many people expect
# positive signs east of Greenwich.  For example, TZ='Etc/GMT+4' uses
# the abbreviation "GMT+4" and corresponds to 4 hours behind UTC
# (i.e. west of Greenwich) even though many people would expect it to
# mean 4 hours ahead of UTC (i.e. east of Greenwich).
```

This behavior is also noted on the Wikipedia page for the time zone database.

34.59 Troubleshooting Traceroute Output

When `traceroute` is run from LAN to a destination on the Internet, the firewall itself may be missing from the `traceroute` output depending on the configuration.

This happens on Multi-WAN due to the way that `route-to` and `reply-to` work. Policy routing (`route-to`/`reply-to`) does not decrease the IP TTL when forwarding packets, so the firewall does not appear as a hop.

This may also happen with IPsec due to the way IPsec traffic is handled in the kernel. The traffic is not "routed" in a traditional sense.

See also:

This behavior may change in future versions of pfSense® software, see Redmine issue #932.

34.60 Troubleshooting Traffic Shaping

Traffic Shaping/QoS is a tricky topic, and can prove difficult to get right the first time. This section covers several common pitfalls.

34.60.1 BitTorrent traffic not using the P2P queue

BitTorrent is known for not using standard ports. Clients are allowed to declare which port other clients use to reach them, which means chaos for network administrators trying to track the traffic based on port alone. Clients can also choose to encrypt their traffic. Regular shaper rules don't have any way to examine the packets to tell what program the traffic appears to be, so it is forced to rely on ports. This is why it may be a good idea to use the P2P Catch all rule, and/or make rules for each type of desirable traffic and treat the default queue as low priority.

34.60.2 UPnP IGD & PCP traffic shaping

Out of the box, traffic allowed in by the UPnP IGD & PCP daemon will end up in the default queue. This happens because the rules generated dynamically by the UPnP IGD & PCP daemon do not have any knowledge of queues unless UPnP IGD & PCP is configured to send traffic into a specific queue.

The nature of the traffic depends on what the client devices utilizing UPnP IGD & PCP on a network are doing. This could be low priority traffic like BitTorrent, or high priority traffic such as videoconferencing or online gaming.

To configure UPnP IGD & PCP to use a specific ALTQ queue:

- Setup ALTQ shaping and decide which queue to use for UPnP IGD & PCP
- Navigate to **Services > UPnP IGD & PCP**
- Enter the chosen ALTQ queue name into the **Traffic Shaping** field
- Click **Save**

This trick only works with the ALTQ shaper. At this time, the firewall is not capable of assigning UPnP IGD & PCP traffic to a limiter.

34.60.3 ACK queue bandwidth calculations

This is a complex topic and most users gloss over it and guess a sufficiently high value. For more detailed explanations with mathematical formulas, check the Traffic Shaping section of the Netgate forum. There is a sticky post in that board which describes the process in great detail, and there is also a downloadable spreadsheet which can be used to help ease the process.

34.60.4 Why is <x> not properly shaped?

The reason is nearly always one of these choices:

- The traffic matched a different rule than expected
- The traffic did not match any rule

As with other questions in this section, this tends to happen because of rules entered either internally or by other packages that do not have knowledge of queues. Since no queue is specified for a rule, it ends up in the default or root queue, and not shaped.

Working around the limitation may require altering the rules to better match the traffic, or disabling internal rules that are matching the traffic in unexpected ways. Another tactic is to identify all other traffic and then use different shaping options on the default queue.

In rare cases, such as BitTorrent, it may be impossible to accurately identify all traffic of a given type. One workaround is to isolate the traffic to one specific device on the network and then match based on that client device address.

34.60.5 WAN connection speed changes

To update the speed of a WAN if it changes, edit the appropriate queues under **Firewall > Traffic Shaper** to reflect the new speed.

The queues that need updating are:

- The root queue for each WAN interface for the upload speed
- The root queue for each LAN interface for the download speed
- `qInternet` queue for each LAN interface for the download speed

If this firewall has multiple WANs, the LAN root and `qInternet` queue must use the total download speed of **all** WANs.

Alternately, if the wizard created all the queues and rules and these have not been changed, then complete the wizard again and update the speed using the wizard.

34.61 Troubleshooting Traffic Shaping Graphs

The RRD for traffic shaping graphs must be reset when a change is made to the traffic shaper settings. The RRD files are in a very specific format and refer to the number and name of the queues as they exist in the shaper configuration. Should this data change, the RRD file data becomes invalid and must be reset.

Therefore, any time a traffic shaper setting is changed, the *queue* and *queuedrops* graphs are reset in order to ensure that the RRD schema matches up with the current shaper configuration.

34.62 Troubleshooting Unexpected Reboots

Unexpected reboots are caused by one of two things – hardware problems, or FreeBSD kernel panics. The vast majority of the time it is hardware problems. Hardware diagnostics should be run before trying anything else.

If the reboot was caused by a kernel panic and the firewall has swap space available, the GUI will display a prompt asking to *view the crash report*. If there is no kernel panic, then the cause is most likely a hardware problem.

For hardware issues, check on:

- Failing power supply
- Flaky electricity in general
- Overheating CPU
- Overheating or faulty RAM
- Faulty hard drive/SSD/other storage
- Faulty drive cables
- and many others…

If the firewall does panic, and the panic message contains a backtrace that mentions things like memory allocation, `mbuf`, `uma_zalloc_arg`, or similar, then it may be crashing due to mbuf exhaustion. See *Tuning and Troubleshooting Network Cards* for information on how to overcome that problem.

Devices with eMMC or small disks may not have any swap space and thus may not be able to recover panic information. In these cases they will typically print the panic information on the console when it occurs. If the device has a serial console, record the output there.

See also:

See *Obtaining Panic Information for Developers* for details.

34.63 Troubleshooting Upgrades

This document describes methods of troubleshooting problems firewalls may encounter when attempting to run a pfSense® software upgrade.

34.63.1 Check Disk Space and ZFS Boot Environments

Running out of space during an upgrade can lead to a variety of failure states that are difficult to fully document. The upgrade process attempts to estimate the amount of disk space it will need, but the calculations are not always accurate enough to prevent problems during the upgrade process itself.

If there are any problems with an upgrade one of the best places to start is by cleaning up disk usage where possible. On systems running pfSense Plus software on ZFS, ZFS Boot Environments can consume significant disk space over time if they are not cleaned up. See *Check and Clean Up ZFS Boot Environments* for details. Additionally, if there are packages installed, package logs and data in particular may be consuming large amounts of disk space.

34.63.2 Low Memory Hardware and AWS/Azure Instances

Hardware with **1 GiB or less** available memory may have issues upgrading depending on which features, services, or packages are running. This includes some Netgate hardware such as the Netgate 1100 when running with ZFS and/or certain services/packages. For the best chance of success in these cases, temporarily disable any non-critical services before starting the upgrade.

Tip: A *Pre-Upgrade Reboot* can also temporarily reduce memory used for ZFS caching, which can help in this situation as well.

---

pfSense Plus software can no longer run on AWS ".nano" size instances as they lack sufficient RAM to upgrade properly. Attempting to upgrade a ".nano" instance to pfSense Plus software version 24.03 will fail before the upgrade is performed. Migrate the instance to a ".micro" or larger size **before** attempting to upgrade, or redeploy instead.

Similar to the above, pfSense Plus software can no longer run on Azure A0 instances. Migrate to instances with more memory.

34.63.3 Remove / Dismount any Installation Media

Some users leave an installation disk plugged in or mounted for various reasons, but this can interfere with the upgrade process. Be sure to dismount and remove any installation media such as a USB thumb drive, optical disk, or ISO in a virtual drive.

A particularly problematic case is when a non-UEFI system has UEFI boot media inserted, the upgrade process may attempt to upgrade the bootloader on the installer and fail:

```
Number of packages to be reinstalled: 1
[1/1] Reinstalling pfSense-boot-23.09...
[1/1] Extracting pfSense-boot-23.09: .......... done
mount_msdosfs: /dev/msdosfs/EFISYS: Read-only file system
pkg-static: POST-INSTALL script failed
 failed.
__RC=1 __REBOOT_AFTER=10
```

34.63.4 Cosmetic Problems Post-Upgrade

If cosmetic problems occur after performing an upgrade, this is nearly always due to stale browser cache entries for CSS, JavaScript, or other files where the browser does not pull the updated version. Force a refresh of the page in the browser (e.g. Shift+Reload or Ctrl+F5) or clear the browser cache to resolve the issues.

34.63.5 Upgrade Log

`pfSense-upgrade` keeps a log of the last update attempt, which may contain additional useful information. This log is located at `/conf/upgrade_log.latest.txt`. Please include the contents of this log with any forum post or Netgate TAC request when requesting assistance with upgrade problems.

34.63.6 Upgrade not Offered / Library Errors

If the update system does not offer an upgrade to the most recent version, the upgrade will not proceed, or the upgrade process encounters errors with shared libraries, take the following steps:

- Navigate to **System > Updates**
- Set **Branch** to *Previous stable version*
- Wait a few moments for the upgrade check to complete
- Optional: Confirm that the latest version of *pfSense-upgrade* is present using `pkg-static info -x pfSense-upgrade`.

  If the correct version is not present, wait a bit longer and check again as that package may be updating in the background.
- Set **Branch** to *Latest stable version*
- Wait a few moments for the upgrade check to complete

If the upgrade is still not offered, refresh the repository configuration and upgrade script by running the following commands from the console or shell:

```
# pkg-static clean -ay; pkg-static install -fy pkg pfSense-repo pfSense-upgrade
```

If that procedure results in an error, or the upgrade is still not offered, then attempt to update to an intermediate version. For example, to get from Plus 2.4.5-p1 to Plus 21.05 or later, upgrade to Plus 21.02.x before proceeding to later versions if a direct upgrade does not succeed. In these cases, the appropriate version will be visible as a branch to select.

34.63.7 Repository Metadata Version Errors

If `pkg` is unable to update and complains about the repository metadata version, the `pkg` utility may need to be updated manually to version 1.13.x or later.

Example metadata version error:

```
>>> Updating repositories metadata...
Updating pfSense-core repository catalogue...
pkg-static: repository meta /var/db/pkg/pfSense-core.meta has wrong version 2
pkg-static: Repository pfSense-core load error: meta cannot be loaded No error: 0
```

To correct the problem, manually bootstrap `pkg` from an ssh or console shell:

```
# pkg-static bootstrap -f
```

Warning: Do not run the above command from **Diagnostics > Command** as it requires interactive input. If ssh and console shells are unavailable, use this variation instead:

```
env ASSUME_ALWAYS_YES=yes pkg-static bootstrap -f
```

34.63.8 Rewrite Repository Information

In some cases the repository information may need to be rewritten:

- Navigate to **System > Updates**
- Set the **Branch** to *Latest Development Snapshots*
- Wait for the page to refresh
- Set the **Branch** to *Latest stable version*

If the update still does not appear, run the commands above from the console or shell.

34.63.9 CLI Troubleshooting

If the GUI update is not functioning as expected, there are a handful of shell commands that can help gather information or resolve problems.

Force pkg Metadata Update

Often times DNS or connectivity problems will prevent the firewall from finding updates. A quick way to verify this is to force a *pkg* metadata update:

```
# pkg-static update -f
```

This command forces an update and will print errors if problems are found, a few potential errors include:

No address record
: The firewall cannot resolve the update server hostname. This could be a problem with DNS from the firewall itself, or connectivity from the firewall to the Internet in general, such as a missing or incorrect default route.

No route to host
: The firewall cannot reach the update server because it cannot find a route there. Most likely, the firewall is missing its default route or the WAN with the default route is down.

Operation timed out
: The firewall was unable to download a file in a timely manner. This is most likely due to degraded connectivity between the firewall and the update servers. It could also be a routing issue, or a problem with IPv6 on the firewall (See *IPv6 Connectivity Problems*).

An error occurred while fetching package
: A general error that could have a few different causes. It may indicate that `pkg` does not trust the package servers. Try running `certctl rehash` from the console, a root shell prompt, or via **Diagnostics > Command Prompt**. This will allow `pkg` to utilize the system certificates until the next reboot.

Authentication error
: There is a proxy between the firewall and the update servers, and it requires authentication. Move the firewall so it is not behind a proxy, or configure the proxy under **System > Advanced**, **Miscellaneous** tab.

No trusted public keys found
: The firewall is attempting to update from the wrong repository. Ensure the correct branch is selected as mentioned in *Rewrite Repository Information*. May require reinstallation to resolve. For CE installations, try the following command:

```
# fetch -qo /usr/local/share/pfSense/keys/pkg/trusted/ \
 https://raw.githubusercontent.com/pfsense/pfsense/RELENG_2_4_5/src/usr/local/share/
↪pfSense/keys/pkg/trusted/pkg.pfsense.org.20160406
```

Debug pkg Metadata Update

If the previous command does not yield any meaningful information, try running the command in debug mode:

```
: pkg-static -d update
```

The output may yield additional useful messages, as in the following example:

```
DBG(1)[69233]> pkg initialized
Updating pfSense-core repository catalogue...
DBG(1)[69233]> PkgRepo: verifying update for pfSense-core
DBG(1)[69233]> PkgRepo: need forced update of pfSense-core
```

(continues on next page)

(continued from previous page)

```
DBG(1)[69233]> Pkgrepo, begin update of '/var/db/pkg/repo-pfSense-core.sqlite'
DBG(1)[69233]> Request to fetch pkg+https://pkg.pfsense.org/pfSense_v2_7_1_amd64-core/
↪meta.conf
DBG(1)[69233]> curl_open
DBG(1)[69233]> Fetch: fetcher used: pkg+https
DBG(1)[69233]> curl> fetching https://pkg.pfsense.org/pfSense_v2_7_1_amd64-core/meta.conf
DBG(1)[69233]> CURL> attempting to fetch from , left retry 3

* Couldn't find host pkg01-atx.netgate.com in the .netrc file; using defaults
*   Trying 208.123.73.209:443...
* Connected to pkg01-atx.netgate.com (208.123.73.209) port 443
* ALPN: curl offers http/1.1
*  CAfile: none
*  CApath: /etc/ssl/certs/
* SSL certificate problem: self-signed certificate in certificate chain
* Closing connection
DBG(1)[69233]> CURL> attempting to fetch from , left retry 2
```

That error suggests a problem with `pkg-static` trusting the package server, so in this case, try running `certctl rehash` from the console, a root shell prompt, or via **Diagnostics > Command Prompt** to allow `pkg` to utilize the system certificates until the next reboot.

Manual Update Check

To run a manual update check from the CLI:

```
# pfSense-upgrade -d -c
```

When run successfully, this command will print a line stating that a new version is available, and the version number of the available update. Errors displayed during that process are likely to be the same as those covered in *Force pkg Metadata Update*.

34.63.10 packages.netgate.com Has no A/AAAA Record

`pkg` does not use A/AAAA records. It uses service (SRV) records. The update server meta names such as `packages.netgate.com` are not meant to be accessed directly using a browser.

To find the actual update servers, lookup the SRV record for the host:

```
$ host -t srv _https._tcp.packages.netgate.com
_https._tcp.packages.netgate.com has SRV record 10 10 443 pkg01-atx.netgate.com.
_https._tcp.packages.netgate.com has SRV record 10 10 443 pkg00-atx.netgate.com.

$ host pkg01-atx.netgate.com
pkg01-atx.netgate.com has address 208.123.73.209
pkg01-atx.netgate.com has IPv6 address 2610:160:11:18::209

$ host pkg00-atx.netgate.com
pkg00-atx.netgate.com has address 208.123.73.207
pkg00-atx.netgate.com has IPv6 address 2610:160:11:18::207
```

Accessing the hosts using their direct hostnames will work with a browser:

```
$ curl --output /dev/null --silent --head --fail \
  "https://pkg00-atx.netgate.com/pfSense_v2_6_0_amd64-core/meta.txz"
$ echo $?
0
```

34.63.11 IPv6 Connectivity Problems

If IPv6 is configured on the firewall, the pfSense software will prefer to use it when performing an update. There are cases when a firewall may have broken IPv6 connectivity, however, that contribute to problems updating. This could manifest as a timeout or routing error when upgrading.

Typically, the operating system will attempt to fall back to IPv4, but the extra time the fallback takes could also lead to a timeout.

The firewall can be configured to prefer IPv4 to eliminate this as a potential cause. See *Controlling IPv6 Preference for traffic from the firewall itself* for details.

Alternately, from ssh or a console shell, force the upgrade to use IPv4 manually:

```
# pfSense-upgrade -4
```

34.63.12 Segmentation Fault in pkg

Certain cryptographic hardware can have a software-induced race condition which leads to a problematic state. In this state, `pkg` will crash with a segmentation fault:

```
1085486128:error:14099044:SSL routines:ssl3_send_client_verify:internal error:
Child process pid=30149 terminated abnormally: Segmentation fault
```

In this case, the device must be powered off and back on to recover. A warm reboot is not sufficient to reset the hardware.

- Navigate to **Diagnostics > Halt System**
- Click  **Halt**
- Wait for the device to shut down. Monitor the console to ensure that the shutdown completes.
- Unplug the power adapter
- Plug the power adapter back in

34.63.13 Forced pkg Reinstall

Forcing a reinstallation of all packages may resolve problems that otherwise may require a full reinstallation. This is not ideal, as a clean installation is more likely to have a positive result, but that is not always an option in every situation (e.g. remote installation with no console access).

To forcefully reinstall all packages, take the following steps:

- Make a backup
- Clean the repository and forcefully reinstall `pkg`, repository data, and the upgrade script:

```
# pkg-static clean -ay; pkg-static install -fy pkg pfSense-repo pfSense-upgrade
```

- Force a reinstallation of everything:

```
# pkg-static upgrade -f
```

- Review the list of changes and enter `y` to proceed
- Manually reboot the firewall

34.63.14 Last Resort

If nothing else works then a reinstallation will eliminate any possibility of problems related to the upgrade itself.

pfSense software offers multiple options to easily restore the configuration. The fastest method is `Recover config.xml` as discussed in *Automatically Restore Configuration During Installation*. Using that method, the pfSense software installation can pick up the existing configuration from the existing install and use it, eliminating the need for any manual restore process. The firewall will boot up after installation with the old settings and reinstall packages as needed.

34.64 Troubleshooting Upgrades on Netgate 1100 and Netgate 2100 Devices

34.64.1 Low Available RAM

Hardware with **1 GiB or less** available memory may have issues upgrading depending on which features, services, or packages are running. This includes some Netgate hardware such as the Netgate 1100 when running with ZFS and/or certain services/packages. For the best chance of success in these cases, temporarily disable any non-critical services before starting the upgrade.

Tip: A *Pre-Upgrade Reboot* can also temporarily reduce memory used for ZFS caching, which can help in this situation as well.

34.64.2 EFI Partition Size

Some older installations of pfSense Plus software on **Netgate 1100**, **Netgate 2100**, and **Netgate 2100 MAX** devices contain an EFI partition which does not have sufficient space to accommodate the new EFI loader for version 23.01 and later. This primarily affects UFS-based systems **initially** installed with pfSense Plus software version 21.02-p1 or before.

Upgrade Notice

Users of affected devices will see a warning about the EFI partition when attempting to upgrade.

When the upgrade check runs, it inspects the system for this problem and files a notice if it identifies a problem:

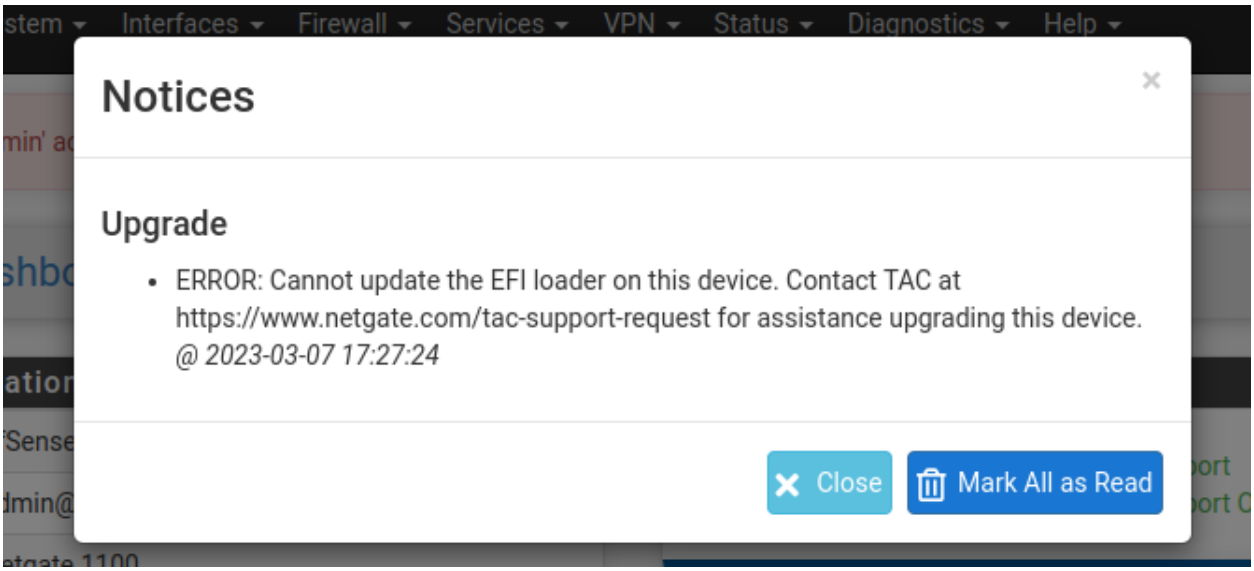


Fig. 8: Notice in the GUI after upgrade check

A similar notice is printed at the command line when checking for updates there:

```
: pfSense-upgrade -c
ERROR: Cannot update the EFI loader on this device. Contact TAC at
https://www.netgate.com/tac-support-request for assistance upgrading this device.
```

Check if the Device is Affected

Users can inspect the EFI partition size by checking the output of `gpart show`.

If the EFI partition size is small (`800K`), then the device must be reinstalled. Larger EFI partition sizes (`64M` or larger) are OK.

This example is a device with an undersized EFI partition:

```
: gpart show
=>       1  15269887  diskid/DISK-B1C82821  MBR  (7.3G)
         1      1600                       1  efi  (800K)
      1601     70012                       2  fat32  (34M)
     71613  15198275                       3  freebsd  (7.2G)
```

Note the size of the `efi` type partition, which is `800K`.

This example is a device with an EFI partition which **can** be upgraded:

```
: gpart show
=>       1  15269887  mmcsd0  MBR  (7.3G)
         1    409600       1  efi  (200M)
    409601     70012       2  fat32  (34M)
    479613  14790275       3  freebsd  [active]  (7.1G)
```

Take a Backup

Before altering the system, *take a local backup*. This backup can be restored at the end of the procedure to retain all current settings.

Tip: Use the *AutoConfigBackup* (ACB) service to store a remote backup, but be sure to note the current device key in ACB as reinstalling will result in the system having a different key unless a backup containing the previous SSH key data is restored.

While AutoConfigBackup is convenient for off-site backups, local file backups can optionally hold and restore much more data including SSH keys, RRD files, and DHCP lease data. Backing up and restoring all the extra data is not strictly necessary, but it makes for a much smoother transition during this kind of reinstallation. Additionally, a local backup can be used with a function such as the External Configuration Locator (ECL) to automatically restore the configuration on the first boot after reinstalling.

Reinstall to Upgrade

Users with affected units **must** reinstall pfSense Plus software to run version 23.01 or later.

To perform this procedure:

- Download the *Netgate Installer* for AARCH64 (*Download Installation Media*).
- Follow the reinstallation instructions in the product manual for the device:
  - Netgate 1100 Reinstallation Procedure
  - Netgate 2100 Reinstallation Procedure

Tip: This is a perfect opportunity to change filesystems from UFS to ZFS!

ZFS is more reliable and has more features than UFS (e.g. ZFS Boot Environments), however ZFS can be memory hungry. Either filesystem will work, but if RAM usage is critical to other tasks that will run on this firewall, UFS can be a more conservative choice. ZFS memory usage can be tuned, however, so that shouldn't be the only deciding factor. See *ZFS Tuning* for details.

Restore the Backup

Restore a backup file after completing the installation using a local file, *Automatic Configuration Backup Service*, or another method such as ECL.

See *Restoring from Backups* for information on the various methods available to restore configuration data.

34.65 Troubleshooting Website Access

If some sites load, but other sites will not, there are a few possible causes.

- Check all the items listed on *Connectivity Troubleshooting* before proceeding
- Ensure the WAN gateway is reachable and set to the proper address
- Ensure the subnet mask on the WAN interface of the firewall is correct
- Ensure the subnet mask on the client stations and on every interface (and VPN) in the pfSense® software configuration is correct
- Ensure the WAN MTU is properly set (See http://www.dslreports.com/faq/695 to determine the MTU), use MSS to lower the MTU if necessary
- Use `traceroute` to determine where the traffic stops. It may be an upstream connectivity issue and not the firewall or ISP.
- *Disable hardware checksums* and see if the problem disappears
- Check **Clear invalid DF bits instead of dropping the packets** on **System > Advanced**, **Firewall/NAT** tab.
- Check **Disable Firewall Scrub** on **System > Advanced**, **Firewall/NAT** tab.

34.66 Troubleshooting Wireless Connections

When it comes to wireless, there are a lot of things that can go wrong. From faulty hardware connections to radio interference to incompatible software or drivers, or simple settings mismatches, anything is possible, and it can be a challenge to make it all work on the first try. This section will cover some of the more common problems that have been encountered by users and developers of pfSense® software.

34.66.1 Check the Antenna

Before spending any time diagnosing an issue, double and triple check the antenna connection. If it is a screw-on type, ensure it is fully tightened. For mini-PCI cards, ensure the pigtail connectors are properly connected and snapped in place. Pigtails on mini-PCI cards are fragile and easy to break. After disconnecting and reconnecting pigtails a few times, they may need to be replaced.

34.66.2 Check Wireless Status

The status of connected wireless clients and nearby access points can be viewed by navigating to **Status > Wireless**. This menu option only appears when a wireless interface is present and enabled.

On this page, click **Rescan** and then refresh the page after 10 seconds to see other nearby access points. If they are on the same or a nearby channel, there could be interference.

In the list of associated clients, several flags are listed that explain the capabilities of the connected client. For example if the client has an "H" flag, this indicates High Throughput used by 802.11n. If a client is connected without that flag, they may be using an older lower standard. A full list of wireless flags is contained in *Wireless Status*, including access point capability descriptions.

34.66.3 Try with multiple clients or wireless cards

To eliminate a possible incompatibility between wireless functions on pfSense software and a wireless client, be sure to try it with multiple devices or cards first. If the same problem is repeatable with several different makes and models, it is more likely to be a problem with the configuration or related hardware than the client device.

34.66.4 Signal Strength is Low

If the signal is weak even when nearby the access point antenna, check the antenna again. For mini-PCI or mini-PCIe cards, if only one pigtail in use and there are two internal connectors, try hooking the pigtail up to the other internal connector on the card. Also try changing the **Channel** or adjusting the **Transmit Power**, or the **Antenna Settings** on the wireless interface configuration. For mini-PCI and mini-PCIe cards, check for broken ends on the fragile pigtail connectors where they plug into the card. If the **Regulatory Domain** settings have not been configured, set them before testing again.

34.66.5 Stuck Beacon Errors

If a "Stuck Beacon" error is found in the system or wireless log, it is usually an indication that the chosen wireless channel is too noisy:

```
kernel: ath0: stuck beacon; resetting (bmiss count 4)
```

The sensitivity of this behavior can be tuned by adding a System Tunable entry for `hw.ath.bstuck` with a value of `8` or higher.

If the errors persist, use a Wi-Fi survey app or program to determine an open or less-used channel to use instead of the current channel.

34.66.6 Interface Unavailable for Assignment

If a wireless interface does not appear in the list of interfaces on the **Interfaces > Assignments** page, there are two possible issues:

If the wireless card is compatible, a wireless instance must first be created as described in *Creating and Managing Wireless Instances*. Once the instance is created, it will be available for assignment.

If the wireless card is not compatible, it will not be available for selection as a parent interface when creating a wireless instance.

34.67 General

- *Troubleshooting Captive Portal*
- *Troubleshooting Offline DHCP Leases*
- *Troubleshooting DHCPv6 Client XID Mismatches*
- *Troubleshooting FTP Connections*
- *Troubleshooting ARP Move Log Messages*
- *Troubleshooting Clock Issues*
- *Troubleshooting Time Zone Configuration*
- *Troubleshooting OS Issues with a Debug Kernel*

34.68 Authentication / User Manager

- *Troubleshooting Authentication*
- *Troubleshooting Access when Locked Out of the Firewall*
- *Troubleshooting "login on console as root" Log Messages*

34.69 Connectivity / Networking

- *Troubleshooting Network Connectivity*
- *Troubleshooting GUI Connectivity*
- *Troubleshooting "No buffer space available" Errors*
- *Troubleshooting Website Access*
- *Troubleshooting Low Interface Throughput*
- *Troubleshooting Lost Traffic or Disappearing Packets*
- *Troubleshooting Traceroute Output*

34.70 DNS

- *Troubleshooting DNS Resolution Issues*
- *Troubleshooting the DNS Cache*
- *Troubleshooting DNS Queries*
- *Troubleshooting Thread Errors with Hostnames in Aliases*

34.71 Hardware

- *Troubleshooting Boot Issues*
- *Troubleshooting Multiple Disks*
- *Troubleshooting in Single User Mode*
- *Troubleshooting DMA and LBA Errors*
- *Troubleshooting Disk and File System Issues*
- *Troubleshooting Full File System or Inode Errors*
- *Troubleshooting File System Capacity Shrinking*
- *Troubleshooting Disk Lifetime*
- *Troubleshooting Disk Writes*
- *Troubleshooting High CPU Load*
- *Troubleshooting "promiscuous mode enabled" Log Messages*
- *Troubleshooting Unexpected Reboots*
- *Troubleshooting Wireless Connections*
- *Troubleshooting Hardware Shutdown and Power Off*

34.72 High Availability

- *Troubleshooting High Availability*
- *Troubleshooting High Availability Clusters in Virtual Environments*
- *Troubleshooting High Availability DHCP Failover*
- *Troubleshooting VPN Connectivity to a High Availability Secondary Node*

34.73 Installation / Upgrades

- *Troubleshooting Installation Issues*
- *Troubleshooting Upgrades*
- *Troubleshooting Upgrades on Netgate 1100 and Netgate 2100 Devices*
- *Troubleshooting a Broken pkg Database*

34.74 Rules/NAT

- *Troubleshooting Firewall Rules*
- *Troubleshooting Bogon Network List Updates*
- *Troubleshooting Blocked Log Entries for Legitimate Connection Packets*
- *Troubleshooting NAT*
- *Troubleshooting 1:1 NAT*
- *Troubleshooting NAT Port Forwards*
- *Troubleshooting NAT Reflection*
- *Troubleshooting Traffic Shaping*
- *Troubleshooting Traffic Shaping Graphs*

34.75 Routing / Multi-WAN

- *Troubleshooting Routes*
- *Troubleshooting Asymmetric Routing*
- *Troubleshooting Multi-WAN*
- *Troubleshooting Gateway Monitoring*

34.76 VPN

34.76.1 IPsec

- *Troubleshooting IPsec VPNs*
- *Troubleshooting IPsec Connections*
- *Troubleshooting IPsec Traffic*
- *Troubleshooting IPsec Logs*
- *Troubleshooting Duplicate IPsec SA Entries*

34.76.2 OpenVPN

- *Troubleshooting OpenVPN*
- *Troubleshooting Windows OpenVPN Client Connectivity*
- *Troubleshooting OpenVPN Internal Routing (iroute)*

34.76.3 Other

- *Troubleshooting L2TP*
- *Troubleshooting Cisco VPN Pass Through*

34.77 Packages

- *Troubleshooting the HAProxy Package*
- *Troubleshooting Snort Rule Updates*
- *Troubleshooting the Squid Package*

CHAPTER

THIRTYFIVE

PFSENSE® SOFTWARE CONFIGURATION RECIPES

The recipes in this section walk administrators through configuring various aspects of pfSense software.

35.1 WAN Connectivity with 802.1X Authentication Bridging and VLAN 0 PCP Tagging

Some Internet Service Providers require their customers to utilize the ISP modem in conjunction with an Optical Network Terminal (ONT) to be granted access to their fiber network. AT&T is one major example of such a provider. However, in some cases it is possible to bypass the modem and connect a firewall directly.

This guide covers the process of configuring a firewall to accommodate this type of authentication.

Note: This guide primarily applies to the AT&T Residential Fiber Network in North America, but can be adapted to any ISP utilizing a similar configuration.

Warning: The configuration options used in this guide are only present on pfSense® Plus software version 23.05-RELEASE and later.

35.1.1 Use Case

The purpose of this configuration is to provide authentication for access to the fiber network. Some ISP modems offer an "IP-Passthrough" mode which enables end users to have their public IPv4 and IPv6 addresses/blocks assigned directly to the equipment behind it (i.e. the firewall). However, this comes with a few drawbacks:

Modem Memory Limitations
: The fiber modem may still track states even in IP-Passthrough mode. Some modems have a hard limit on the number of states that they can handle at a time, becoming unstable under significant load.

Limitations in IPv6 Implementation
: In IP-Passthrough mode the modem is usually provisioned with an IPv6 prefix (`/60` for AT&T, for example), but will only hand out a single `/64` prefix out of the larger allocation via DHCP-PD to the firewall. This means that only a single LAN on the firewall can be provisioned with IPv6 by default. It is possible to request multiple `/64` networks out of the IPv6 prefix block, but that is an ugly workaround.

Multiple Points of Failure
: Having an ONT, a modem, and a firewall all needing to be powered at all times and available at all times introduces unnecessary additional points of potential hardware failure that can bring down connectivity even if the physical fiber link is in working order.

Bypassing the ISP equipment and attaching directly to the ONT with a pfSense Plus firewall eliminates or reduces the above limitations, allowing for greater control and flexibility.

Warning: The best practice when bypassing the ISP modem is to disable the Wi-Fi Access Point in the ISP equipment. This scenario requires an alternate means of Wi-Fi and switch connectivity behind the firewall to ensure connectivity parity with the ISP-provided all-in-one solution for Wi-Fi connectivity.

35.1.2 Requirements

Authenticating the firewall and allowing it to connect to the provider requires the following:

1. A firewall with at least **three** unique, discrete interfaces: One for the modem, one for the WAN/ONT connection, and one for the inside network(s).
2. The modem must be able to authenticate access using 802.1X EAP-TLS authentication. ISP modems using this type of 802.1X authentication have a "burned in" certificate and will initiate authentication when attached to a physical network on the red "ONT" port. This is handled on boot-up of the modem normally when it is in-line between the ONT and the local equipment, and it will periodically retry authentication.
3. All traffic after authentication must be 802.1Q tagged on VLAN `0` with a Priority Code Point (PCP) of `1`. PCP is a means of defining traffic priority. A PCP of `1` is "Best Effort" and is how most ISPs, including AT&T, expect traffic to be marked. Configuring a PCP on a non-VLAN interface in pfSense Plus will tag the traffic on VLAN `0` and include the PCP value.
4. The WAN interface on pfSense Plus software must have the MAC address spoofed to match the WAN interface of the fiber modem. This MAC address may be printed on sticker attached to the modem, or it may be visible in the web interface on the modem.
5. The pfSense Plus software interface attached to the modem must be set to operate in promiscuous mode.
6. The firewall must send all IPv6 DHCP requests with a defined and expected DUID. A DUID is a unique identifier a device uses when requesting a DHCPv6 lease. Normally pfSense software will use an automatically generated random identifier, but ISPs such as AT&T expect a DUID-EN (DUID Enterprise Number) of `3561` and an identifier tied to the serial number of the modem. The identifier for a modem can be generated using an open source script.

   See also:

   To learn more about DUIDs, see *DHCP6 DUID*.
7. The firewall must send a prefix hint when requesting a DHCPv6 Prefix Delegation. Typically, this is `/60` for AT&T. A `/60` prefix allows for 16 interfaces to each have a unique /64 subnet assigned from this block.

35.1.3 Modem Bypass Configuration

Physical Connections

Set up the physical connections as shown in *Diagram of Auth Bridge Wiring Layout*:

- Connect the ONT device LAN/Modem port to the NIC on the firewall which will be the **WAN** interface
- Connect the ISP modem ONT/WAN port (may be marked in red) to the NIC on the firewall which will be the **MODEM** interface
- Connect the NIC on the firewall which will be the **LAN** interface to a switch or other means of local connectivity

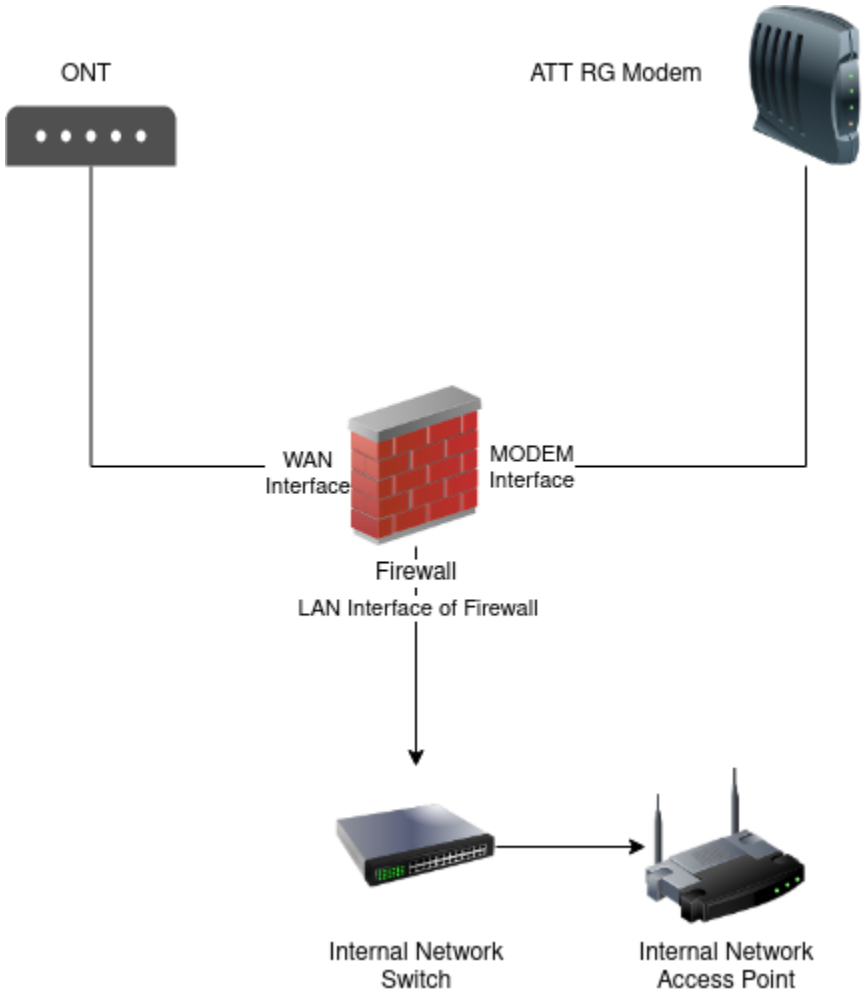


Fig. 1: Diagram of Auth Bridge Wiring Layout

Configure Firewall Interfaces

The next step is to configure the interfaces in the pfSense Plus software GUI.

Modem Interface

Assign and configure a new interface for the ISP Modem:

- Navigate to **Interfaces > Assignments**
- Set **Available network ports** to the physical interface attached to the ISP Modem
- Click  **Add**
- Note the name of the new interface (e.g. **OPT1**)
- Navigate to the newly added OPT interface using the **Interfaces** menu (e.g. **Interfaces > OPT1**)
- Configure the interface options as follows:

  Enable interface
  Checked

  Description
  `MODEM`

  IPv4 Configuration Type
  None

  IPv6 Configuration Type
  None

  Enable Promiscuous Mode
  Checked

- Click **Save**
- Click **Apply Changes**

The interface will now be available at **Interfaces > MODEM** and will appear as a choice with that name in various interface fields in the GUI.

WAN/ONT Interface

Now configure the WAN interface to send traffic that the ONT and ISP will accept:

- Navigate to the interface attached to the ONT (e.g. **Interfaces > WAN**)
- Configure the options as follows:

  Enable interface
  Checked

  Description
  `WAN` or another similarly descriptive name

  IPv4 Configuration Type
  DHCP

  IPv6 Configuration Type
  DHCP6

MAC Address
Enter the MAC address of the WAN interface on the **ISP Modem**

Priority Tag
1

DHCPv6 Prefix Delegation size
Set to match the value provided by the ISP, e.g. `60` for AT&T

Send IPv6 prefix hint
Checked

Do not wait for a RA
Checked

- Click **Save**
- Click **Apply Changes**

LAN Interface

Finally, configure the LAN and other local interfaces:

- Navigate to **Interfaces > LAN** or its equivalent
- Configure the options as follows:

  Enable interface
  Checked

  Description
  `LAN` or another similarly descriptive name

  IPv4 Configuration Type
  Static IPv4 using whichever private LAN subnet is already in place.

  IPv6 Configuration Type
  Track Interface

  Track IPv6 Interface

    IPv6 Interface
    WAN or whichever interface is attached to the ONT

    IPv6 Prefix ID
    1

- Click **Save**
- Click **Apply Changes**

Repeat this for any remaining inside interfaces. For each additional interface, increment the **IPv6 Prefix ID** value by `1` in **hexadecimal**. On AT&T or other providers who delegate a `/60` prefix size, the maximum ID value is `f`. The help text under the **IPv6 Prefix ID** field is automatically adjusted to show the minimum and maximum allowed values for the ID.

Configure IPv6 DUID

Set the custom DUID to send to the ISP:

- Navigate to **System > Advanced**, **Networking** tab
- Configure the options on the page as follows, leaving other unmentioned options at their current values:

  DHCP6 DUID
  DUID-EN: Assigned by Vendor based on Enterprise Number

  DUID-EN

  IANA Private Enterprise Number
  `3561`

  Identifier
  Enter the DUID value generated by the gen-duid.sh script

- Click **Save**

Configure Authentication Passthrough

Passing through the authentication traffic between the modem and ISP requires two Ethernet rules to bridge the appropriate traffic.

Enable Ethernet Rules

The Ethernet rules feature is disabled by default and must be manually enabled before use:

- Navigate to **System > Advanced**, **Firewall & NAT** tab
- Check **Enable Ethernet Filtering** in the **Advanced Options** section
- Click **Save**

See also:

Ethernet (Layer 2) Rules

Add WAN-Modem Bridge Rule

Add a rule to bridge 802.1X authentication traffic from the WAN/ONT to the MODEM interface:

- Navigate to **Firewall > Rules**, **Ethernet** tab
- Click  **Add**
- Configure the rule as follows:

  Action
  Pass

  Quick
  Checked

  Interface
  WAN or whichever interface is attached to the ONT

Direction
in

Protocol
IEEE 802.1X

Source
any

Destination
any

- Click **Display Advanced** and set:

Bridge To
MODEM

- Click **Save**

Add Modem-WAN Bridge Rule

Add a rule to bridge anything sent by the ISP Modem to the WAN interface:

- Navigate to **Firewall > Rules**, **Ethernet** tab
- Click  **Add**
- Configure the rule as follows:

Action
Pass

Quick
Checked

Interface
MODEM

Direction
in

Protocol
Any

Source
any

Destination
any

- Click **Display Advanced** and set:

Bridge To
WAN or whichever interface is attached to the ONT

- Click **Save**
- Click **Apply Changes**

35.1.4 Finish Up

The modem bypass configuration is now complete. Reboot the firewall to ensure the settings are applied completely. During the boot sequence, the modem should detect the link change and begin transmitting 802.1X authentication requests across the Layer 2 filter to the WAN interface, and the WAN interface should be able to acquire a DHCP and DHCPv6 lease.

35.2 Authenticating Users with Google Cloud Identity

Google Cloud Identity LDAP service can be used to authenticate users on pfSense® software installations.

The method varies depending on the version of pfSense software installed on the firewall. This is due to the fact that Google Cloud Identity requires a client certificate to make a secure LDAP connection.

- Firewalls running pfSense Plus software can use a client certificate directly on LDAP authentication sources.
- Firewalls running pfSense CE software or older unsupported versions require the stunnel package to make a secure LDAP connection.

Configuring a firewall running pfSense software to use G Suite LDAP authentication requires a number of steps, all of which are covered in this document.

35.2.1 Configure the LDAP Application on the G Suite admin portal

Follow the instructions from Google for configuring and enabling the G Suite LDAP application.

> **Warning:** **Follow these directions exactly**. No special provisions are required for pfSense, but please note that the LDAP application credentials (username and password) **are required**.

35.2.2 Download the certificate, key, username and password

Download the certificate, key, username and password from G Suite to a local directory on a workstation.

35.2.3 Import the certificate and key

From the web interface of a firewall running pfSense:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add/Sign** to display the certificate import interface
- Configure the entry as follows:

  Method
  Import an existing certificate

  Descriptive name
  `G Suite LDAP`

  Certificate data
  Copy and paste the contents of the downloaded certificate

Private Key data
Copy and paste the contents of the downloaded key

- Click **Save**

The certificate is now available for use by the firewall.

The next step depends on the version of pfSense software installed on the firewall.

pfSense CE software requires the stunnel package to make secure LDAP connections. For these environments, proceed to *Install the stunnel package (pfSense CE software)*.

pfSense Plus software can use a client certificate directly when communicating with LDAP authentication sources. Skip ahead to *Configure LDAP authentication on pfSense software*.

35.2.4 Install the stunnel package (pfSense CE software)

From the web interface on pfSense:

- Navigate to **System > Package manager**, **Installed Packages** tab
- Check the list for **stunnel** and if it is listed as installed
- If the package is installed and up-to-date, with a version of **5.37** or later, no action is required
- If the package is installed but out of date
  - Update the package by clicking  for the **stunnel** entry
  - Click  **Confirm** to confirm the package update
- If the **stunnel** package is not installed
  - Navigate to the **Available packages** tab
  - Locate the **stunnel** package in the list, or use the search bar
  - Click  **Install** for the **stunnel** package entry
  - Click  **Confirm** to confirm the package installation

35.2.5 Configure the stunnel package (CE or 2.4.4-RELEASE)

From the web interface on pfSense:

- Navigate to **Services > STunnel**
- Click  **Add** to create a new profile
- Configure the profile as follows:

Description
Text describing this connection, such as `G Suite`

Client Mode
Check

Listen on IP
`127.0.0.1`

Listen on port
`1636`

Certificate
The entry imported previously, in this case *G Suite LDAP*

Redirects to IP
`ldap.google.com`

Redirects to port
`636`

- Click **Save**

35.2.6 Configure LDAP authentication on pfSense software

From the web interface on pfSense:

- Select **System > User manager**, **Authentication servers** tab
- Click  **Add** to create a new entry
- Enter a **Descriptive name** for this LDAP server, such as `G Suite`
- Configure the basic settings for the server as follows:

Type
LDAP

Protocol version
3

Server timeout
`25`

Search scope
Entire tree

The remaining settings depend on which version of pfSense software is installed:

For pfSense Plus:

Hostname or IP address
`ldap.google.com`

Port value
`636`

Transport
SSL - Encrypted

Peer Certificate Authority
Global Root CA List

Client Certificate
The entry imported previously, in this case *G Suite LDAP*

For pfSense CE or when using stunnel:

Hostname or IP address
`127.0.0.1`

Port value
`1636`

Transport
TCP-Standard

The next few settings are **UNIQUE TO THE DOMAIN**. For this example, assume that is `example.com`.

> **Warning:** Substitute the actual domain when entering these values!

Base DN
The domain name in DN format, for example `dc=example,dc=com`

Authentication containers
Base DN prepended by the `Users` organizational unit, for example: `ou=Users,dc=example,dc=com`

Bind anonymous
Unchecked to show the **Bind Credentials** fields

Bind credentials
The G Suite LDAP username and password that were created with the certificate and key

The remaining attributes are not specific to the domain, or are defaults

User naming attribute
`uid`

Group naming attribute
`cn`

Group member attribute
`memberOf`

35.2.7 Create a Group

Using a remote authentication server to manage administrative logins to services on pfSense software requires a matching group to be present on both the authentication source server and on the firewall. The existing `admins` group could be used, but since the name is so general it may conflict with other desired permissions in G Suite.

This example uses a new group called `fwadmins`.

First, create the `fwadmins` group in G Suite and assign users to the group. The exact details will vary based on the domain and its organization.

Next, create a group on the firewall running pfSense software. This does not require local users, only a group entry. The group entry must have appropriate permissions.

To create the group on pfSense:

- Navigate to **System > User Manager**, **Groups** tab
- Click  **Add** to make a new group
- Configure the group as follows:

Group name
Name of the group, in this example: `fwadmins`

Scope
Remote

Description
Remote Firewall Administrators

- Click **Save**

Now the group needs privileges:

- Click  on the row for the newly created group
- Click  **Add** in the **Assigned Privileges** section
- Select the desired permissions for the group, for example: `WebCfg - All pages`

> **Warning:** **Do not select every item** in this list! Doing so will also select the `User - Config: Deny Config Write` privilege which will prevent users in this group from making changes to the firewall configuration.

- Click **Save** to store the privileges

35.2.8 Test G Suite Authentication

With the complete configuration described above, it is now possible to authenticate against Google G Suite LDAP. First, test the authentication to ensure it is working properly.

- Navigate to **Diagnostics > Authentication**
- Set the **Authentication server** to the name used for the LDAP Server entry, such as *G Suite*
- Enter a known username and password on the domain that G Suite controls

---

Note: By default, only the username part of the login is checked against the configured LDAP base DN. If a username is submitted with a domain part, for example `user@example.com`, the `@example.com` part is ignored.

- Click  **Test**

The user should show as authenticating successfully, and if the user entered is a member of the `fwadmins` group in G Suite, that should also be reflected in the test output.

If the test succeeds, the service is ready for use. pfSense software can use it as an authentication source for the GUI, for VPNs, or anywhere the user manager authentication servers work.

If the test fails, check the main system log for error messages from LDAP. Start from the beginning of this document and compare all settings between this document, G Suite, and pfSense software. Most common problems are with parameters being input incorrectly, such as selecting the wrong certificate, using an incorrect LDAP attribute name, or not using correct bind credentials.

35.2.9 Use G Suite for Administrative Logins

If all is well and the user authenticated as expected:

- Navigate to **System > User manager**, **Settings**
- Set the **Authentication server** to *G Suite*
- Click **Save**

After saving, firewall users will be authenticated against Google Cloud Identity.

Note: pfSense software automatically falls back to local authentication if it cannot authenticate using the chosen LDAP server.

35.3 Configuring BIND as an RFC 2136 Dynamic DNS Server

If the DNS service for a domain is directly controlled on a BIND server, RFC 2136 Dynamic DNS can be configured for use by pfSense® software. This section shows how to configure BIND to provide this service.

The exact location of the configuration directory for BIND varies by operating system. It could be in `/usr/local/etc/namedb/`, `/etc/namedb/`, or elsewhere.

See also:

See *Configuring RFC 2136 Dynamic DNS updates* for more information on RFC 2136 Dynamic DNS.

35.3.1 Configure the BIND Server

On the server in `named.conf`, add the following block:

```
include "/etc/namedb/dns.keys.conf";
zone "dyn.example.com" {
     type master;
     file "dynamic/dyn.example.com";
     update-policy { grant *.dyn.example.com. self . A AAAA; };
};
```

Then create the initial zone file. BIND requires read-write access to this file and the directory in which it resides so that the zone and its journal may be updated.

Warning: BIND will rewrite this zone file, which is why a subdomain is used in the example.

From there, create the zone file for the new dynamic zone, `dynamic/dyn.example.com`

```
$ORIGIN      .
$TTL 30      ; 30 seconds
dyn.example.com             IN SOA  ns.example.com. hostmaster.example.com. (
                            2016062801 ; serial
                            3600       ; refresh (1 hour)
                            600        ; retry (10 minutes)
                            2600       ; expire (43 minutes 20 seconds)
```

(continues on next page)

(continued from previous page)

```
                            30         ; minimum (30 seconds)
                            )
                    NS      ns.example.com.
                    NS      ns2.example.com.
```

Reload the named service using `rndc reload` or a similar command, and then if any slave name servers are in place, add a zone to those servers as well:

```
zone "dyn.example.com" {
     type slave;
     file "dynamic/dyn.example.com";
     masters{ 192.0.2.5; };
};
```

For BIND 9.16+ create an entry using `tsig-keygen`:

```
# tsig-keygen -a hmac-sha256 myhost.dyn.example.com
key "myhost.dyn.example.com" {
     algorithm hmac-md5;
     secret "/0/4bxF9A08n/zke/vANyQ==";
};
```

Add that key to `dns.keys.conf` manually or by redirecting command output:

```
# tsig-keygen -a hmac-sha256 myhost.dyn.example.com >> /etc/namedb/dns.keys.conf
```

For BIND version < 9.16, follow the next steps.

On the master name server, make the `keys` directory:

```
# mkdir -p /etc/namedb/keys
```

And now generate a host key. The second line is the output of the command, *not* part of the command itself.

```
# /usr/sbin/dnssec-keygen -K /etc/namedb/keys -a HMAC-MD5 -b 128 -n HOST myhost.dyn.
↪example.com.
Kmyhost.dyn.example.com.+157+32768
```

The output `Kmyhost.dyn.example.com.+157+32768` is the first part of the filename for the key, it will append `.private` to one file and `.key` to another. Both contain the same data in different formats.

Now read the key from the new key file:

```
# /usr/bin/grep ^Key: /etc/namedb/keys/Kmyhost.dyn.example.com.+157+32768.private | /usr/
↪bin/awk '{ print $2; }'
/0/4bxF9A08n/zke/vANyQ==
```

And then add that key to `dns.keys.conf`:

```
key "myhost.dyn.example.com" {
     algorithm hmac-md5;
     secret "/0/4bxF9A08n/zke/vANyQ==";
};
```

This can be automated with a simple script, `make-ddns-host.sh`:

```
#!/bin/sh
KEY_NAME=${1}
KEY_DIR=/etc/namedb/keys
KEYS_CONFIG=/etc/namedb/dns.keys.conf
/bin/mkdir -p ${KEY_DIR}
cd ${KEY_DIR}
KEY_FILE_NAME=`/usr/sbin/dnssec-keygen -K ${KEY_DIR} -a HMAC-MD5 -b 128 -n HOST ${KEY_
↪NAME}.`
KEY_TEXT=`/usr/bin/grep "^Key:" ${KEY_DIR}/${KEY_FILE_NAME}.private | /usr/bin/awk '{␣
↪print $2; }'`
echo "key ${KEY_NAME}. {" >> ${KEYS_CONFIG}
echo "       algorithm hmac-md5;" >> ${KEYS_CONFIG}
echo "       secret \"${KEY_TEXT}\";" >> ${KEYS_CONFIG}
echo "};" >> ${KEYS_CONFIG}
echo "Key for ${KEY_NAME} is: ${KEY_TEXT}"
```

After making the file, make it executable:

```
# chmod u+x make-ddns-host.sh
```

To use the script:

```
# ./make-ddns-host.sh mynewhost.dyn.example.com
# rndc reload
```

35.3.2 Configuring a Client in pfSense Software

To add a DynDNS entry in the pfSense software GUI:

- Navigate to **Services > Dynamic DNS**, **RFC 2136** tab
- Click  **Add** to create a new entry with the following settings:

  Enable
  Checked

  Interface
  WAN

  Hostname
  The fully qualified hostname, e.g. `xxxxx.dyn.example.com`

  TTL
  `30`

  Key Name
  The fully qualified hostname again, exactly: `xxxxx.dyn.example.com`

  Key algorithm
  HMAC-SHA256

  Key
  Secret key for this hostname

  Server
  The IP address or hostname of the BIND server

Protocol
Unchecked

Description
`My DynDNS Entry`

- Click **Save**

Assuming the firewall has connectivity to the name server, and there are no other access policies that would prevent the update, RFC 2136 DynDNS service is now working. If the update does not work, check the BIND log and the system log on the firewall.

35.4 Blocking Websites

There are several options for blocking websites with pfSense® software, some of which are described on this article. This is not an exact science, but these solutions typically function well enough for a majority of use cases.

See also:

The pfBlockerNG package (*pfBlocker-NG Package*) offers mechanisms which can be useful in this area, such as DNSBL, geographic IP address blocking, and automation of AS lookups.

35.4.1 Using DNS

If the built-in DNS Resolver or Forwarder are active an override can be entered there to resolve the unwanted website to an invalid IP address such as `127.0.0.1`.

> **Warning:** Do not use DNS override functionality as the only means of blocking access to sites.
>
> Blocking via DNS requires that local clients utilize the firewall as their only DNS source. See *Redirecting Client DNS Requests* and *Blocking External Client DNS Queries* for suggestions on ensuring clients get their DNS responses from the firewall. It will stop non-technical users, but it is easy to circumvent for those with more technical aptitude.

With the DNS Resolver, additional methods are possible via custom options.

This first example will prevent any host under the given zone from being resolved by clients:

```
server:
local-zone: "movie.edu" static
```

When the firewall enforces DNS resolution in this way, the firewall must also force clients to resolve DNS using the firewall. Otherwise, clients could bypass the restrictions by using alternate DNS servers. See *Redirecting Client DNS Requests* for details.

This can be limited in scope using custom views. This example is similar to the above, but only blocks access for `10.6.0.100`:

```
server:
access-control-view: 10.6.0.100/32 blocksites

view:
name: "blocksites"
local-zone: "movie.edu" static
```

35.4.2 Using Firewall Rules

If a website rarely changes IP addresses, then it can be blocked by an alias. Create an alias containing its IP addresses and then use this alias in firewall rules.

> **Warning:** This is not a feasible solution for sites that return low TTLs and spread the load across many servers and/or datacenters, such as Google and similar large sites. Most small to mid-sized websites can be effectively blocked using this method as they rarely change IP addresses.

A hostname can also be inside a network alias. The firewall will resolve the hostname periodically and update the alias as needed. This is more effective than manually looking up the IP addresses, but will still fall short if the site returns DNS records in a way that changes rapidly or randomizes results from a pool of servers on each query, which is common for large sites.

Another option is finding all the IP subnet allocations for a site. Create an alias with those networks and block traffic to those destinations. This is especially useful with sites such as Facebook that spread large amounts of IP space, but are constrained within a few net blocks. Using regional registry sites such as ARIN can help track down those networks. For example, all the networks used by Facebook in the region covered by ARIN can be found at http://whois.arin.net/rest/org/THEFA-3.html under "Related Networks". Companies may have other addresses in different regions, so check other regional sites as well, such as RIPE, APNIC, etc.

As an alternative to looking up the IP blocks manually, locate the BGP Autonomous System (AS) number for the target company by doing a `whois` lookup on one of their IP addresses. For example, the AS number for Facebook is `AS32934` and the following command will locate all of their allocations:

```
# whois -h whois.radb.net -- '-i origin AS32934' | awk '/^route:/ {print $2;}' | sort |␣
↪uniq
```

Copy the results of that command into a new alias, and it will cover all of their currently allocated networks. Check the results periodically for updates.

35.4.3 Using a Proxy

In modern environments a client proxy is not effective. HTTPS can sometimes be filtered via peek/splice to inspect SNI and similar aspects of connections, but even that fails with modern security practices like encrypted SNI. Using proxies for these tasks is no longer a recommended practice.

35.4.4 Prevent Bypassing Restrictions

With any of the above methods, there are many ways to get around the defined blocks. The easiest and likely most prevalent is using any number of proxy websites. Finding and blocking all of these individually and keeping the list up to date is impossible. The best way to ensure these sites are not accessible is using an external proxy or content filtering capable of blocking by category.

To further maintain control, use a restrictive egress ruleset and only allow traffic out to specific services and/or hosts. For example, only allow DNS access to the firewall or the DNS servers specifically used for LAN clients (*Redirecting Client DNS Requests*). Also, if a proxy is in use on the network, make sure to disallow direct access to HTTP and HTTPS through the firewall and only allow traffic to and/or from the proxy server.

35.5 Changing Credentials and Keys

Organizations may have guidelines about how often to change credentials such as passwords and encryption keys. Different types of credentials and keys are stored in various places in the configuration, and changing them can range from trivial or cumbersome depending on the type and how the firewall uses them.

These types of guidelines can vary widely based on various industry recommendations, certification standards, and other criteria. Due to these variations, this document won't make any specific commentary on timing of such changes.

See also:

- *Password Storage Security Policies*
- *Caveats and Gotchas*

- *User Manager Accounts*
  - *Unprivileged Users*
- *Authentication Servers*
  - *LDAP*
  - *RADIUS*
- *Interfaces*
  - *PPP type WANs*
  - *Wireless WANs*
  - *Wireless APs*
- *Certificate Data and Private Keys*
- *VPNs*
  - *IPsec Pre-Shared Keys*
  - *IPsec Certificates*
  - *WireGuard*
  - *OpenVPN*
  - *PPPoE Server Users*
  - *L2TP Server Users*
- *Notification Services*
  - *SMTP*
  - *Telegram*
  - *Pushover*
  - *Slack*
- *Upstream Proxy*
- *Captive Portal*
- *DHCP Server*
  - *OMAPI*

    - *Dynamic DNS Key*
- *DNS Resolver*
- *Dynamic DNS*
    - *Dynamic DNS Service Clients*
    - *RFC 2136 Clients*
- *NTP*
- *Packages*
    - *ACME*
    - *FRR*
    - *FreeRADIUS*
    - *NET-SNMP*
    - *HAProxy*
    - *Stunnel*
    - *Zabbix*
    - *Others*

35.5.1 User Manager Accounts

Administrators can change the password for their own account and for accounts of other users in the User Manager:

- Navigate to **System > User Manager**
- Find the user account in the list
- Click  at the end of the row to edit the user account
- Enter a new **Password** and enter it again in the **Confirm Password** field.
- Click **Save**

See also:

User Management and Authentication

Unprivileged Users

Non-administrator users with accounts in the user manager who have the "WebCfg - System: User Password Manager" privilege can log in to the GUI with their existing username and password and change the password for their own account to a new value in the same place (**System > User Manager**).

The GUI displays a simplified form for these users with only the password change fields.

Warning: Do not expose the GUI to the Internet. Only allow users to reach the firewall GUI from a local interface or using a secure VPN.

35.5.2 Authentication Servers

User authentication might be handed off to external authentication servers. In this case the user credentials must be changed on the authentication server. However, the credentials the firewall itself uses when communicating with the authentication servers might also need to be changed.

LDAP

LDAP servers may have two items to update, depending on the configuration:

TLS Certificate

If communication with the LDAP server users TLS (STARTTLS or dedicated TLS port) then there may be occasions when the server CA expires or needs to be changed. If there is a new server CA, replace the CA data in the *Certificate Manager*:

- Navigate to **System > Certificates**, **CAs** tab
- Find the LDAP server CA in the list
- Click  to edit the CA
- Paste in the new CA certificate and/or private key data in PEM format
- Click **Save**

Alternately, import a new CA entry, then edit the LDAP authentication server entry and switch to the new CA.

Bind Credentials

LDAP authentication servers may use authenticated or anonymous binds when validating LDAP users. If the entry uses bind credentials, these may change over time.

To update the bind credentials:

- Navigate to **System > User Manager**, **Authentication Servers** tab
- Find the LDAP server entry in the list
- Click  to edit the LDAP server entry
- Enter the new **Bind credentials** (**User DN**, **Password**)
- Click **Save**

RADIUS

RADIUS authentication servers have a **Shared Secret** (sometimes called a NAS password) which allows the firewall to perform authentication requests. If the shared secret for the firewall is changed on the RADIUS server, update the RADIUS authentication server entry on the firewall to match:

- Navigate to **System > User Manager**, **Authentication Servers** tab
- Find the RADIUS server entry in the list
- Click  to edit the RADIUS server entry
- Enter the new **Shared Secret**
- Click **Save**

35.5.3 Interfaces

A few interface types can have credentials as well. For WANs, these should be changed with the ISP first and then updated on the firewall to match.

PPP type WANs

These types of WANs (e.g. PPPoE, L2TP, PPP/Cellular) can be changed in either of two places: The interface configuration or the PPPs configuration.

Choose **one** of the following methods and update the credentials there.

Interface Method

- Navigate to **Interfaces > <Name>** for the interface in question
- Enter the new **Username**, **Password**, and **Confirm Password**
- Click **Save**
- Click **Apply Changes**

PPPs Configuration Method

- Navigate to **Interfaces > Assignments**, **PPPs** tab
- Find the entry in the list
- Click  to edit the entry
- Enter the new **Username**, **Password**, and **Confirm Password**
- Click **Save**

Wireless WANs

If the upstream wireless provider changes the pre-shared key or 802.x/EAP passphrase, a wireless WAN must change to match the new value(s):

- Navigate to **Interfaces > <Name>** for the interface in question
- Enter the new **WPA Pre-Shared Key** and/or **Inner Authentication Passphrase**
- Click **Save**
- Click **Apply Changes**

Wireless APs

WPA2

Access Points using WPA/WPA2 should periodically change the pre-shared key:

- Navigate to **Interfaces > <Name>** for the interface in question
- Enter the new **WPA Pre-Shared Key**
- Click **Save**
- Click **Apply Changes**

Clients will need to enter the new key to reconnect.

802.1x

Access Points using 802.1x/EAP may need to update the RADIUS shared secret if it changes on the RADIUS server.

- Navigate to **Interfaces > <Name>** for the interface in question
- Enter the new **Shared Secret** for the **Primary 802.1x server** and/or **Secondary 802.1x server**
- Click **Save**
- Click **Apply Changes**

35.5.4 Certificate Data and Private Keys

CA and Certificate entries have built-in expiration dates and can be easily renewed in the GUI. On occasion, it may be necessary to change the private key for a CA or certificate as well.

> **Warning:** When renewing a CA, changing the private key or serial number will invalidate any certificate signed by the CA as these are part of the data used to validate the trust chain.

To renew a CA or certificate and generate a new private key:

- Navigate to **System > Certificates**, **CAs** or **Certificates** tab
- Find the entry to renew
- Click  to start the renewal process

- Uncheck **Reuse Key**

  With this option unchecked, the renewal process will generate a new private key for the certificate. This is generally safe for server and user certificates but can be problematic for CAs as mentioned previously.

- Inspect the **Certificate Properties**

  If any are considered weak, consider checking **Strict Security** to use the recommended properties for stronger security.

- Click **Renew/Reissue**

If the entry is a CA, send the new CA to clients that need it. Assuming the key and serial were reused they can continue using the old CA until it expires, but they will need to replace their local copy of the CA before that date arrives.

Note: When renewing the self-signed GUI certificate, it is safe to replace the key and using a new serial number is required.

35.5.5 VPNs

For optimal security, changing VPN keys periodically is a good practice. The practicality of doing so largely depends on the size of the VPN (e.g. a simple two-site point-to-point link vs a remote access setup with dozens of users)

In any case, the most important thing is to coordinate the change with the remote peer(s) so all parties are using the correct keys.

IPsec Pre-Shared Keys

Changing pre-shared key values is fairly simple but must be done in a coordinated fashion. As soon as one side changes the key, the other side will fail to negotiate the tunnel the next time it attempts to authenticate.

Tunnels

- Navigate to **VPN > IPsec**
- Locate the VPN tunnel in the list
- Click  to edit the tunnel Phase 1 entry
- Enter a new **Pre-Shared Key**
- Click **Save**
- Click **Apply Changes**

Remote Access (PSK and EAP)

- Navigate to **VPN > IPsec**, **Pre-Shared Keys** tab
- Locate the key entry in the list
- Click  to edit the entry
- Enter a new **Pre-Shared Key**
- Click **Save**
- Click **Apply Changes**

IPsec Certificates

An IPsec tunnel using certificate-based authentication will have two certificates that may need changed: **My Certificate** which is used by this firewall to identify itself, and **Peer Certificate Authority** which this firewall uses to authenticate the peer.

If the CA or certificate was created on this firewall, the entry can be renewed as described in *Certificate Data and Private Keys*.

To update a CA or Certificate entry in-place with data from a remote source:

- Navigate to **System > Certificates**, **CAs** or **Certificates** tab
- Find the entry in the list
- Click  to edit the entry
- Paste in the new certificate and/or private key data in PEM format
- Click **Save**

Alternately, renew, create, or import a new CA/Certificate, then select the new entry:

- Navigate to **VPN > IPsec**
- Locate the VPN tunnel in the list
- Click  to edit the tunnel Phase 1 entry
- Select the new **My Certificate** and/or **Peer Certificate Authority** entries
- Click **Save**
- Click **Apply Changes**

WireGuard

To update a new key for a tunnel:

- Navigate to **VPN > WireGuard**, **Tunnels** tab
- Find the entry in the list
- Click  to edit the entry
- Enter a new private key

  Alternately, click  **Generate** to automatically generate a new key pair.
- Click **Save Tunnel**

Warning: The tunnel will be down until the remote peer(s) update their configurations with the new public key.

If a peer changes their key, edit the peer and update:

- Navigate to **VPN > WireGuard**, **Peers** tab
- Find the entry in the list
- Click  to edit the entry
- Enter a new **Public Key**
- Click **Save Tunnel**

Another strategy for remote access setups is to make a new tunnel with settings similar to the old one. It must have a unique port number and interface addresses, but it can otherwise use the same settings. Place the new keys on the new server and exchange new keys with clients. Once all clients are on the new server, retire the old entry and remove all of its contents.

OpenVPN

Warning: Shared Key mode is deprecated, move to certificate-based tunnels.

OpenVPN tunnels using certificates can, for the most part, be handled by methods already covered in this document. For example, if the CA or certificate was created on this firewall, the entry can be renewed as described in *Certificate Data and Private Keys* and the warnings there still apply.

Consider generating a new TLS auth/encryption key periodically as well:

- From an SSH or console shell prompt, or **Diagnostics > Commands**, run the following shell command:

```
$ openvpn --genkey secret
#
# 2048 bit OpenVPN static key
#
-----BEGIN OpenVPN Static key V1-----
```

(continues on next page)

(continued from previous page)

```
aed37de925b750e934efbcca1f342267
a428f6ee2bd677a25c0f2815f04c1d53
ba5c7a268c6b351312ee8753fa757204
50c274de24a70b199e5f4f2c094f48cf
3fbcdca14bbb1344ca042288766201ca
f057300e97c70b78aaec9385877d87d5
ad4e8a3bda1d528a03130117d63d1ed6
cd5fdfda0ee41d1d039dcbd458a58666
793a72d3393d57b906b2ee2d03748516
6a401d162c1d0da2b83d689eb5cb9a12
285b17b2de3bb816eb927e890696350e
ae6328485b4d02e4adbe4f867a4871c4
61af2d62e4693e4a334f5e540d9b5e9c
82e6f9c9b833ac8b2f83f025e48822cd
c1f8a7cc57cfb60a5adda5a3287d128c
f0a29f14cb4e1e7fda8174c4a7226252
-----END OpenVPN Static key V1-----
```

- Copy the resulting text including the # header lines and --- armor lines.
- Navigate to **VPN > OpenVPN**, **Servers** tab
- Find the entry in the list
- Click  to edit the entry
- Erase the existing TLS key
- Paste in the new TLS key
- Click **Save**

Warning: Coordinate with all clients as changing it on the server will require clients to use the new key immediately!

Another popular strategy for remote access setups is to make a new VPN server with settings similar to the old one. It must have a unique port number and tunnel network, but it can otherwise use the same settings. Place the new TLS key, CA, and certificates on this new server and deliver the new files to clients. Once all clients are on the new server, retire the old entry and remove all of its contents.

PPPoE Server Users

If the firewall is acting as a PPPoE server, it has a separate configuration for users and/or RADIUS authentication.

Local Users

- Navigate to **Services > PPPoE Server**
- Locate the server entry in the list
- Click  to edit the server
- Locate user in the **User table**
- Set a new **Password**
- Click **Save**

RADIUS

- Navigate to **Services > PPPoE Server**
- Locate the server entry in the list
- Click  to edit the server
- Set and confirm a new **Primary RADIUS Server Shared Secret** and/or **Secondary RADIUS Server Shared Secret**
- Click **Save**

L2TP Server Users

If the firewall is acting as an L2TP server, it has a separate configuration for users and/or RADIUS authentication.

Local Users

- Navigate to **VPN > L2TP Server**, **Users** tab
- Locate the entry in the list
- Click  to edit the entry
- Set and confirm a new **Password**
- Click **Save**

RADIUS

- Navigate to **VPN > L2TP Server**
- Set and confirm a new **Secret** in the **RADIUS** section
- Click **Save**

35.5.6 Notification Services

Remote notification types frequently require credentials to contact the remote server and deliver messages. If the credentials change on the server, the firewall needs the new credentials to continue delivering remote notifications.

SMTP

- Navigate to **System > Advanced**, **Notifications** tab
- Enter a new **Notification E-mail auth password** in the **SMTP** Section
- Click **Save**
- Click **Test SMTP Settings**

Telegram

- Navigate to **System > Advanced**, **Notifications** tab
- Enter a new **API key** in the **Telegram** Section
- Click **Save**
- Click **Test Telegram Settings**

Pushover

- Navigate to **System > Advanced**, **Notifications** tab
- Enter a new **API key** and/or **User Key** in the **Pushover** Section
- Click **Save**
- Click **Test Pushover Settings**

Slack

- Navigate to **System > Advanced**, **Notifications** tab
- Enter a new **API key** in the **Slack** Section
- Click **Save**
- Click **Test Slack Settings**

35.5.7 Upstream Proxy

If the firewall must authenticate to an upstream proxy for its own outgoing HTTP/HTTPS requests and those credentials change, the firewall will need the new password for outbound connectivity:

- Navigate to **System > Advanced**, **Miscellaneous** tab
- Enter a new **Proxy Password**
- Click **Save**

35.5.8 Captive Portal

Uses the *User Manager* for local users, *Authentication Servers* for remote users, and *certificates from the certificate manager*. See those respective sections for information on updating credentials.

35.5.9 DHCP Server

OMAPI

To change the OMAPI key:

- Navigate to **Services > DHCP Server**
- Paste in a new **OMAPI Key** or check **Generate New Key** to create a new key automatically
- Click **Save**

Dynamic DNS Key

If the DHCP server securely sends dynamic DNS updates to an upstream DNS server, those credentials can be changed if the server updates the key:

- Navigate to **Services > DHCP Server**
- Paste in a new **DNS Domain key secret** from the DNS server
- Click **Save**

35.5.10 DNS Resolver

Uses *certificates from the certificate manager* when acting as a DNS over TLS server. See that section for information on updating credentials.

35.5.11 Dynamic DNS

Dynamic DNS Service Clients

The type and format of credentials for Dynamic DNS clients vary by provider.

- Navigate to **Services > Dynamic DNS**, **Dynamic DNS Clients** tab
- Find the entry in the list
- Click  to edit the entry
- Enter the new **Password** or equivalent credential
- Click **Save**

  Alternately, click **Save and Force Update** to also force an update to test the new settings.

RFC 2136 Clients

- Navigate to **Services > Dynamic DNS**, **RFC 2136 Clients** tab
- Find the entry in the list
- Click  to edit the entry
- Pick a new **Key algorithm** if it changed
- Enter the new **Key**
- Click **Save**

  Alternately, click **Save and Force Update** to also force an update to test the new settings.

35.5.12 NTP

If NTPv3 authentication is enabled and the server key changes, the firewall must be changed to match.

- Navigate to **Services > NTP**
- Enter a new **Authentication Key**
- Pick a new **Digest Algorithm** if it changed
- Click **Save**

35.5.13 Packages

Add-on packages can have their own sets of credentials. This section is not comprehensive. Check each add-on package individually.

ACME

Certificates

These certificates generally only have a 90-day lifetime, therefore it is unusual to renew them sooner as their relatively short life is quite secure.

To forcefully renew an ACME certificate:

- Navigate to **Services > ACME Certificates**, **Certificates** tab
- Locate the entry in the list
- Click **Issue/Renew**

Private Keys

Changing a private key for a certificate isn't a common need, the package does not currently have a good method to forcefully generate a new private key of the same type and size.

- Navigate to **Services > ACME Certificates**, **Certificates** tab
- Locate the entry in the list
- Click  to edit the entry
- Change the **Private Key** selection to a different length or type (RSA vs ECDSA)

  Alternately, generate a new private key externally and then choose *Custom* and paste the private key data in PEM format.
- Click **Save**
- Locate the entry in the list
- Click **Issue/Renew**

When the package renews the certificate it will change the private key to the new size and/or type.

Account Keys

The account key identifies a particular ACME user, but it can be changed if needed.

- Navigate to **Services > ACME Certificates**, **Account Keys** tab
- Locate the entry in the list
- Click  to edit the entry
- Click  **Create new account key**
- Click  **Register ACME account key**
- Click **Save**
- Navigate to **Services > ACME Certificates**, **Certificates** tab
- Locate each entry in the list using this account key
- Click **Issue/Renew** on each entry using the key

FRR

The FRR package can optionally store credentials used to communicate with peers.

Global

The master password is only used internally by FRR. Changing it is typically unnecessary but simple to do as it does not need to be updated anywhere else.

- Navigate to **Services > FRR Global/Zebra**
- Enter a new **Master Password**

BGP

- Navigate to **Services > FRR BGP**, **Neighbors** tab
- Locate the entry in the list
- Click  to edit the entry
- Enter a new **Password**
- Click **Save**

OSPF

- Navigate to **Services > FRR OSPF**, **Interfaces** tab
- Locate the entry in the list
- Click  to edit the entry
- Enter a new **Password**
- Click **Save**

FreeRADIUS

The FreeRADIUS package contains credentials for users as well as for its clients and other servers.

Note: If using a database, change data there instead.

Users

- Navigate to **Services > FreeRADIUS**, **Users** tab
- Locate the entry in the list
- Click  to edit the entry
- Enter a new **Password**
- Enter a new One-Time Password **Init-Secret** and/or **PIN** if needed
- Click **Save**

NAS / Clients

- Navigate to **Services > FreeRADIUS**, **NAS / Clients** tab
- Locate the entry in the list
- Click  to edit the entry
- Enter a new **Client Shared Secret**
- Click **Save**

EAP Certificates

EAP uses certificates from the *Certificate Manager*. To change those, see that section instead.

SQL

If this FreeRADIUS configuration pulls its data from an SQL server, it will have stored credentials for communicating with the SQL server. To update those credentials:

- Navigate to **Services > FreeRADIUS**, **SQL** tab
- Enter a new **Database Password** in the section for server 1 and optionally for server 2.
- Click **Save**

LDAP

If this FreeRADIUS configuration passes authentication requests to an LDAP server, it may have stored credentials for communicating with the LDAP server. To update those credentials:

- Navigate to **Services > FreeRADIUS**, **LDAP** tab
- Enter a new **Password** in the section for server 1 and optionally for server 2.
- Click **Save**

If communication with the LDAP server uses certificates, they are stored in the *Certificate Manager*. To change those, see that section instead.

NET-SNMP

NET-SNMP SSL/TLS options use entries from the *Certificate Manager*. To change those, see that section instead.

Users

- Navigate to **Services > NET-SNMP**, **Users** tab
- Locate the entry in the list
- Click  to edit the entry
- Enter a new **Password** and/or **Passphrase**
- Click **Save**

Communities

- Navigate to **Services > NET-SNMP**, **Communities** tab
- Locate the entry in the list
- Click  to edit the entry
- Enter a new **Community Name**
- Click **Save**

HAProxy

HAProxy SSL/TLS options use entries from the *Certificate Manager*. To change those, see that section instead.

Stunnel

HAProxy SSL/TLS options use entries from the *Certificate Manager*. To change those, see that section instead.

Zabbix

- Navigate to **Services > Zabbix Agent** or **Services > Zabbix Proxy**
- Paste a new **TLS PSK** value from the server

SSL/TLS options use entries from the *Certificate Manager*. To change those, see that section instead.

Others

The contents of the package system change on an ongoing basis. There may be additional third party packages with credentials not mentioned in this document. Check the settings of any other installed package for local or remote credentials.

35.6 Diagnostic Data

pfSense® software includes a function to gather diagnostic data. The data collected by this function can be useful when diagnosing issues, especially when seeking technical assistance from Netgate TAC or community members on the Netgate Forum. This information may also be referred to as "Status Output".

This function is handled by an unlinked script at `/usr/local/www/status.php`. The script produces output and an optional archive containing a wide array of information from an installation.

Warning: Though the script makes an effort to remove private information such as passwords, keys, and other secrets, the output may still contain sensitive data. This is especially true when it comes to packages which may use non-standard names for items in the configuration.

Always inspect the output data to ensure it does not contain any private information.

There are multiple ways to invoke the script, both in the GUI and through the shell.

- *View and Download Diagnostic Data in the GUI*
- *Generate Diagnostic Data Archive in the Shell*
- *Script Options*
- *Copying the Diagnostic Data Archive*
  - *Download from the GUI*
  - *Download Using SCP*
  - *Copy to a USB Drive*

35.6.1 View and Download Diagnostic Data in the GUI

This script is not linked from any menu as it's not a function users would need to access on a regular basis.

To load the script, first *connect to the GUI* and then add `/status.php` to the end of the URL.

For example: `https://x.x.x.x/status.php`

Note: The page can take a while to load depending on the hardware.

The script will create an archive with the diagnostic data along with a  **Download** button to easily download the archive file.

By default, the script also outputs all the diagnostic data to the browser which can make it easy for users to inspect. This output can be suppressed, see *Script Options* for details.

35.6.2 Generate Diagnostic Data Archive in the Shell

The script can also be invoked from the shell if the GUI is unavailable. Additionally, in some cases the script may fail to run in the GUI but succeed from the shell. The script can be run from the console (video or serial) or via SSH.

To start, *connect to the firewall console* or *enable SSH* and *connect using an SSH client*.

Note: This script requires `admin` or `root` privileges, so use either of those usernames or use the *Sudo Package* to run the script from another user.

When connected to the console or SSH use option `8` from the menu to load a shell prompt and then execute the script:

```
php -f /usr/local/www/status.php
```

The script will gather the data and create an archive file at `/tmp/status_output.tgz`. This file can then be copied off the firewall using one of the methods covered later in this document.

35.6.3 Script Options

When invoking the script there a couple options that control its behavior:

archiveonly
: When set, suppresses the output of the script in the GUI so that it only generates an archive for download purposes.

  This option is only valid in the GUI as the console method only offers archive output.

nocleanup
: When set, the script leaves behind all the individual diagnostic data files in `/tmp/status_output/` which is handy for inspecting the data directly in the shell without using the archive file.

To use the options in the GUI, they must be passed as query parameters, for example:

- `https://x.x.x.x/status.php?archiveonly`
- `https://x.x.x.x/status.php?nocleanup`
- `https://x.x.x.x/status.php?archiveonly&nocleanup`

When running the script in the shell, pass the option the script filename:

```
php -f /usr/local/www/status.php nocleanup
```

35.6.4 Copying the Diagnostic Data Archive

After using the script users can copy the archive file (`/tmp/status_output.tgz`) off the firewall in a variety of ways.

Download from the GUI

The easiest way to download the archive from the GUI is to use `/status.php` directly in the GUI. If that is not viable, use the *Download function* on **Diagnostics > Command Prompt** to download the `/tmp/status_output.tgz` archive file after generating it in the shell.

Download Using SCP

If the SSH daemon is *enabled on the firewall*, then an SCP client can copy files from the firewall remotely. Alternately, SCP can copy files from the firewall to a remote SSH server from the shell.

See *Accessing Files on the Firewall with SCP* for information on setting up and using SCP.

Copy to a USB Drive

Alternately, bypass the network and copy the archive file to a USB drive by following the procedure in *Copy Files to a USB Drive*.

35.7 Blocking External Client DNS Queries

This procedure configures the firewall to block DNS requests from local clients to servers outside the local network. With no other accessible DNS servers, clients are forced to send DNS requests to the DNS Resolver or DNS Forwarder on pfSense® software for resolution.

Note: Blocking is effective but does not gracefully handle the situation. Clients must manually adjust their configuration to use the firewall for DNS. Redirecting DNS requests to the firewall is a more seamless solution. See *Redirecting Client DNS Requests* for details.

- Navigate to **Firewall > Rules**, **LAN** tab
- Create the block rule as the first rule in the list:
  - Click  **Add** to create a new rule at the top of the list
  - Fill in the following fields on the rule:

    Action
    Reject

    Interface
    LAN

    Protocol
    TCP/UDP

    Destination
    Any

    Destination Port Range
    DNS (53)

    Description
    `Block DNS to Everything Else`

- Create the pass rule to allow DNS to the firewall, above the block rule:
  - Click  **Add** to create a new rule at the top of the list
  - Fill in the following fields on the rule:

    Action
    Pass

    Interface
    LAN

    Protocol
    TCP/UDP

    Destination
    LAN Address

    Destination Port Range
    DNS (53)

    Description
    `Pass DNS to the Firewall`

- Click  **Apply Changes** to reload the ruleset

When complete, there will be two rule entries that look like the following picture:

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
✓	0 /0 B	IPv4 TCP/UDP	*	*	LAN address	53 (DNS)	*	none		Pass DNS to the Firewall	
	0 /0 B	IPv4 TCP/UDP	*	*	*	53 (DNS)	*	none		Block DNS to Everything Else	

Certain local PCs could be allowed to use other DNS servers by placing a pass rule for them above the block rule.

35.7.1 DNS over TLS

Another concern is that clients could use DNS over TLS to resolve hosts. DNS over TLS sends DNS requests over an encrypted channel on an alternate port, `853`.

This traffic can be blocked with a firewall rule for port `853` using the same procedure used for `53`. Though if the firewall will not be providing DNS over TLS service to clients, do not add the pass rule.

35.7.2 DNS over HTTPS

Similar to DNS over TLS, clients may also use DNS over HTTPS (DoH). This is harder to block as it uses port `443`. Blocking port `443` on common public DNS servers may help (e.g. `1.1.1.1`, `8.8.8.8`).

Some browsers automatically attempt to use DNS over HTTPS because they believe it to be more secure and better for privacy, though that is not always the case. Each browser may have its own methods of disabling this feature. Firefox uses a "canary" domain `use-application-dns.net` by default if the user has not manually enabled DNS over HTTPS. If Firefox cannot resolve this name, Firefox disables DNS over HTTPS.

To prevent Firefox from using DNS over HTTPS, add the following to the DNS Resolver custom options:

```
server:
local-zone: "use-application-dns.net" always_nxdomain
```

35.8 Configuring DNS over TLS

Several popular public DNS providers provide encrypted DNS service using DNS over TLS. This prevents intermediate parties from viewing the content of DNS queries and can also assure that DNS is being provided by the expected DNS servers.

35.8.1 Requirements

This feature is only implemented in the *DNS Resolver*. If the firewall is currently using the *DNS Forwarder*, convert to the DNS Resolver before starting this procedure.

Pick a DNS over TLS upstream provider, such as a private upstream DNS server or a public service like Cloudflare, Quad9, or Google public DNS. Note the addresses of the servers and their associated hostnames.

35.8.2 Configure DNS Servers

First, configure the DNS servers on the firewall.

> **Warning:** When the firewall uses DNS over TLS, **every** DNS server the firewall uses **must** provide DNS over TLS service.

- Navigate to **System > General**
- Locate the **DNS Server Settings** Section
- Add or replace entries in the **DNS Servers** section such that only the chosen DNS over TLS servers are in the list

  Address
  : IP address of an upstream DNS Server providing DNS over TLS service (e.g. `1.1.1.1`).

  Hostname
  : Hostname of the same upstream DNS Server in the **Address** field, used for TLS certificate validation (e.g. `cloudflare-dns.com`).

  > **Warning:** The hostname is technically optional but dangerous to omit. The DNS Resolver must have the hostname to validate that the correct server is providing a given response. The response is still encrypted without the hostname, but the DNS Resolver has no way to validate the response to determine if the query was intercepted and answered by a third party server (Man-in-the-Middle attack).

- Click  **Add DNS Server** and repeat the previous step as needed for each available DNS server
- Uncheck **Allow DNS server list to be overridden by DHCP/PPP on WAN**

  DNS servers obtained dynamically may not provide DNS over TLS.

- Set **DNS Resolution Behavior** to *Use local DNS (127.0.0.1), ignore remote DNS Servers*

  This makes the firewall itself use only the DNS Resolver, and it will not attempt to contact the DNS servers directly. This prevents DNS requests from the firewall being leaked unencrypted on port `53` if the resolver is temporarily unavailable (*DNS Resolution Behavior*).

- Click **Save**

Use *Example DNS Server list for DNS over TLS from Cloudflare* as a reference for the settings on the page.

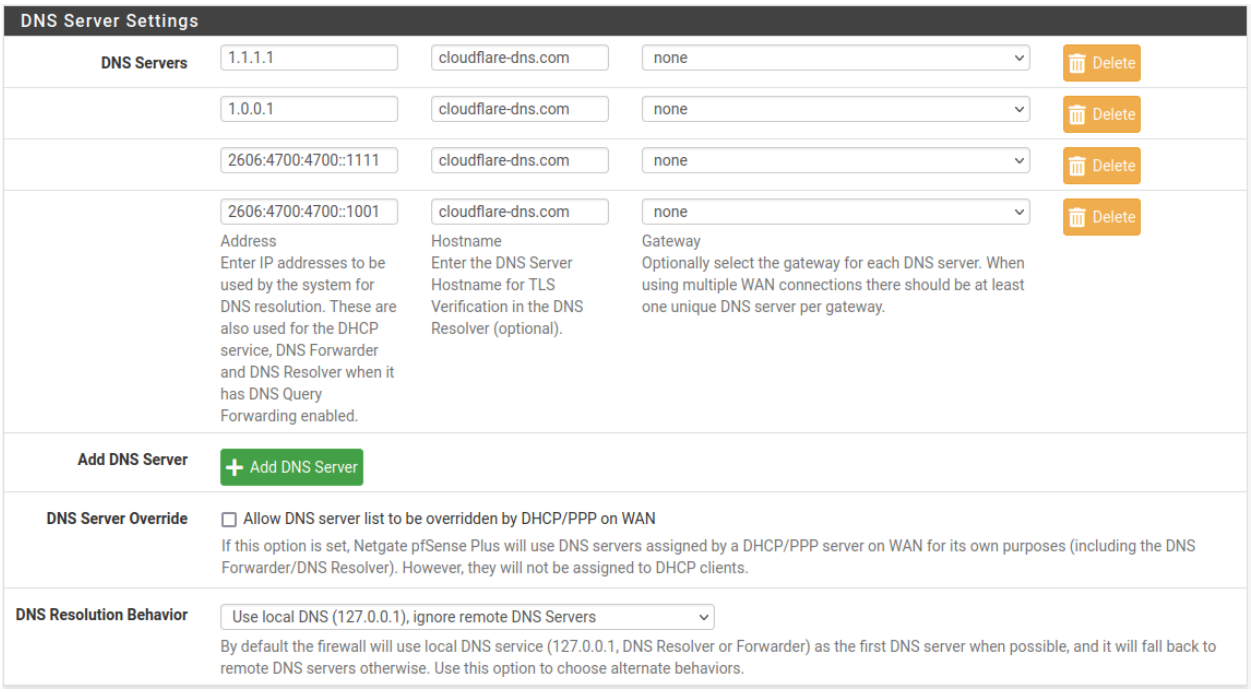


Fig. 2: Example DNS Server list for DNS over TLS from Cloudflare

35.8.3 Enable DNS over TLS for Forwarded Queries

Next, configure the DNS Resolver to use DNS over TLS for outgoing queries.

- Navigate to **Services > DNS Resolver**
- Uncheck **Enable DNSSEC Support**

> **Note:** DNSSEC is not generally compatible with forwarding mode, with or without DNS over TLS.

- Check **Enable Forwarding Mode**
- Check **Use SSL/TLS for outgoing DNS Queries to Forwarding Servers**
- Click **Save**
- Click **Apply Changes**

Use *Example DNS Resolver configuration for outgoing DNS over TLS* as a reference for the settings on the page.

The DNS Resolver will now send queries to all upstream forwarding DNS servers using SSL/TLS on the default port of `853`.

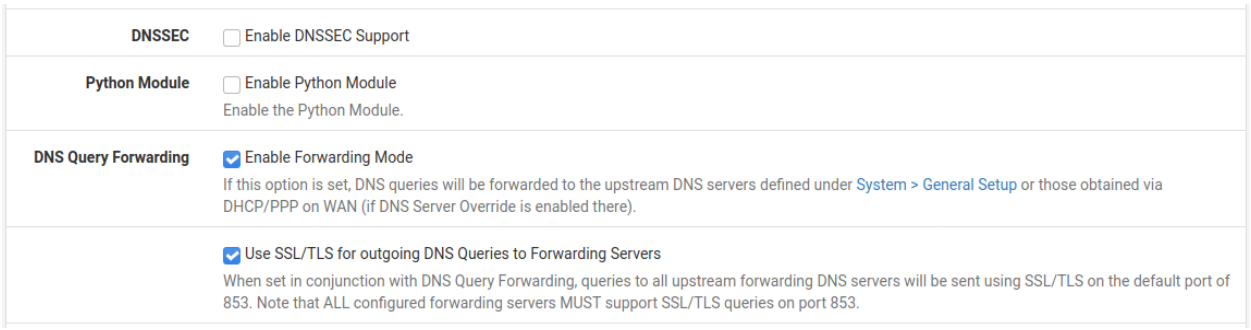


Fig. 3: Example DNS Resolver configuration for outgoing DNS over TLS

35.8.4 Testing DNS over TLS

There are several ways to validate that outbound queries are using DNS over TLS.

- Test via **Diagnostics > DNS Lookup** (*DNS Lookup*) and ensure the results from `127.0.0.1` are correct.
- Check for states using port `853` going to the DNS servers in the configuration (*Firewall States*) like those in *Example State Table contents for DNS over TLS queries*.
- Packet capture port `853` (*Packet Capturing*) and inspect the capture in Wireshark. The contents of the query are not visible, but the TLS exchange is, and any TLS errors in negotiation should be visible in the capture.

States

Interface	Protocol	Source (Original Source) -> Destination (Original Destination)	State	Packets	Bytes
WAN	tcp	198.51.100.23:51344 -> 1.0.0.1:853	FIN_WAIT_2:FIN_WAIT_2	11 / 9	952 B / 4 KiB
WAN	tcp	198.51.100.23:31300 -> 1.1.1.1:853	FIN_WAIT_2:FIN_WAIT_2	11 / 9	947 B / 4 KiB

Fig. 4: Example State Table contents for DNS over TLS queries

35.8.5 Enable DNS over TLS Server (optional)

The DNS Resolver can also act as a DNS over TLS server, though this function **does not** affect outbound/forwarded queries, so this section is optional.

Enable this feature only when *local clients* may need to communicate with the *DNS Resolver* using DNS over TLS queries.

- Navigate to **Services > DNS Resolver**
- Check **Respond to incoming SSL/TLS queries from local clients**
- Select a valid server certificate in **SSL/TLS Certificate**

> **Note:** Clients may reject this certificate if it is self-signed, consider using a certificate from *ACME*.

- Leave **SSL/TLS Listen Port** at the default (empty or `853`)
- Click **Save**
- Click **Apply Changes**

Use *Example DNS Resolver configuration for acting as a DNS over TLS Server* as a reference for the settings on the page.

Now the DNS Resolver will listen for DNS over TLS queries from local clients on TCP port 853.

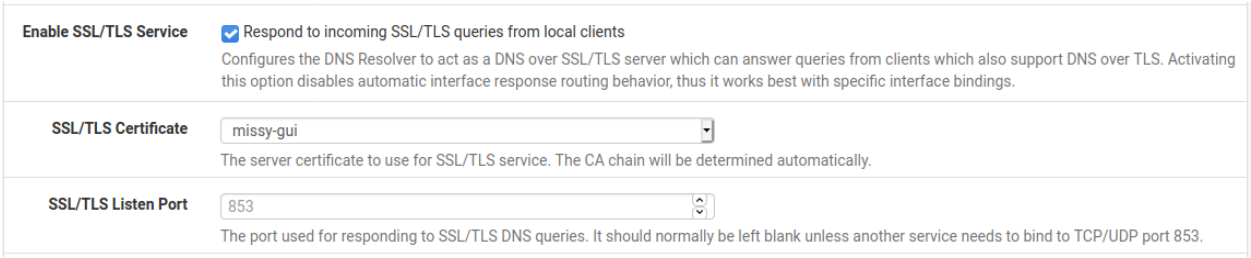


Fig. 5: Example DNS Resolver configuration for acting as a DNS over TLS Server

35.8.6 Caveats

Clients can make their own connections to DNS over TLS servers, so block them on TCP/UDP ports 53 and 853 to ensure they only query the DNS Resolver (*Blocking External Client DNS Queries*).

Redirecting DNS over TLS queries to the DNS Resolver may or may not work, depending on the clients. Set up the *DNS over TLS server* and add port forward redirects for TCP/UDP ports 53 and 853 to redirect DNS queries to the firewall (*Redirecting Client DNS Requests*).

Note: Though clients may reject the DNS over TLS server certificate since it would not match their intended server, this could still have the intended result. The client may fall back to traditional DNS queries if DNS over TLS validation fails.

35.9 Redirecting Client DNS Requests

To restrict client DNS to only the DNS Resolver or Forwarder on pfSense® software, use a port forward to capture all client DNS requests.

Note: Either The DNS Resolver or DNS Forwarder must be active, and it must bind to and answer queries on *Localhost*, or *All* interfaces.

See also:

- *Blocking External Client DNS Queries*
- *Blocking Web Sites Using DNS*

The following example uses the LAN interface, but the same technique will work with any local interface.

- Navigate to **Firewall > NAT**, **Port Forward** tab
- Click  **Add** to create a new rule
- Fill in the following fields on the port forward rule:

Interface
LAN

Protocol
TCP/UDP

Destination
Invert Match *checked*, *LAN Address*

Destination Port Range
DNS (53)

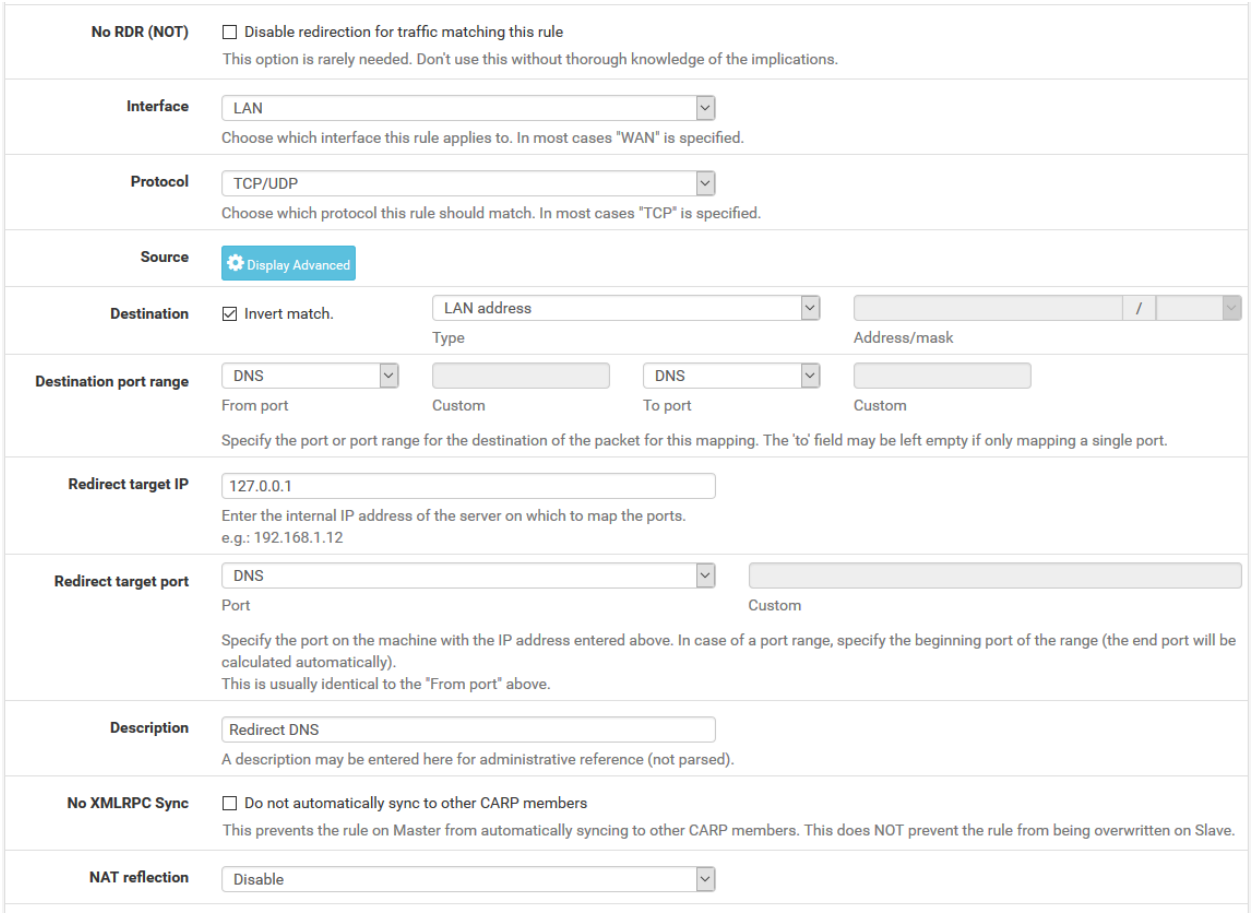
Redirect Target IP
`127.0.0.1`

Redirect Target Port
DNS (53)

Description
`Redirect DNS`

NAT Reflection
Disable

When complete, the port forward must appear as follows:



Note: If DNS requests to other DNS servers are blocked, such as by following *Blocking External Client DNS Queries*,

ensure the rule to pass DNS to `127.0.0.1` is above any rule that blocks DNS.

With this port forward in place, DNS requests from local clients to **any** external IP address will result in the query being answered by the firewall itself. Access to other DNS servers on port 53 is impossible.

Tip: This can be adapted to allow access to only a specific set of DNS servers by changing the Destination network from "LAN Address" to an alias containing the allowed DNS servers. The **Invert match** box should remain checked.

Warning: Clients using DNS over TLS or DNS over HTTPS could circumvent this protection. Redirecting or blocking port `853` may help with DNS over TLS, depending on the clients.

See *Blocking External Client DNS Queries* for additional advice.

35.10 Dynamic Routing Protocol Basics

Three routing protocols are available in pfSense® software using the FRR package:

- *BGP* (Border Gateway Protocol)
- *OSPF* (Open Shortest Path First v2, for IPv4).
- *OSPF6* (Open Shortest Path First v3, for IPv6).

An in depth discussion of routing protocols is outside the scope of this documentation.

All the available dynamic routing protocols are handled by the *FRR package*. To install FRR:

- Navigate to **System > Package Manager**
- Click **Available Packages**
- Locate **FRR** in the list, or search for it
- Click the  **Install** to the right of the **FRR** package entry.
- Click  **Confirm**
- Wait for the installation to complete
- Navigate to **Services > FRR Global/Zebra**

35.10.1 BGP

The *FRR package documentation* contains a basic *BGP example* which can be used as the basis for most common configurations.

The general form of configuration for FRR BGP is:

- Add a route map to allow routes to be filtered
- Configure the BGP options for the local AS
- Add neighbors

35.10.2 OSPF

The *FRR package documentation* contains a basic *OSPF example* which can be used as the basis for most common configurations.

The general form of configuration for OSPF in FRR is:

- Add interfaces as needed, with local interfaces being marked passive, and those facing other OSPF routers as active.
- Configure the general settings as needed with the router ID, area ID, and so on.

See also:

OpenVPN Site-to-Site with Multi-WAN and OSPF contains an example configuration of OSPF.

35.10.3 OSPF6

The *FRR package documentation* contains a basic *OSPF6 example* which can be used as the basis for most common configurations.

The general form of configuration for OSPF6 in FRR is:

- Add interfaces as needed, with local interfaces being marked passive, and those facing other OSPF6 routers as active.
- Configure the general settings as needed with the router ID, area ID, and so on.

35.11 Basic Firewall Configuration Example

This article is designed to describe how pfSense® software performs rule matching and a basic strict set of rules. The approach described in this document is not the most secure, but will help show how rules are set up.

Rules on the **Interface** tabs are matched on the **incoming** interface.

See also:

Read the *Aliases* article as it will make management of rules easier.

35.11.1 Basic lock down of the LAN and DMZ outgoing rules

Outbound LAN

Make sure the **Default LAN > any** rule is either disabled or removed.

1. Allowing DNS access:
   - If pfSense is the DNS server:
     - Allow **TCP/UDP 53** (DNS) from LAN subnet to **LAN Address**.
   - If using Upstream DNS Servers:
     - Allow **TCP/UDP 53** (DNS) from LAN subnet to **Upstream DNS Servers**.
   - Otherwise:
     - Allow **TCP/UDP 53** (DNS) from LAN subnet to **anywhere**.
2. Allowing all users to browse web pages anywhere:

- Allow **TCP 80** (HTTP) from LAN subnet to **anywhere**.

3. Allowing users to browse secure web pages anywhere:
   - Allow **TCP 443** (HTTPS) from LAN subnet to **anywhere**.
4. Allowing users to access FTP sites anywhere:
   - Allow **TCP 21** (FTP) from LAN subnet to **anywhere**.
5. Allowing users to access SMTP on a mail server somewhere:
   - Allow **TCP 25** (SMTP) from LAN subnet to **anywhere**.
6. Allowing users to access POP3 on a mail server somewhere:
   - Allow **TCP 110** (POP3) from LAN subnet to **anywhere**.
7. Allowing users to access IMAP on a mail server somewhere:
   - Allow **TCP 143** (IMAP) from LAN subnet to **anywhere**.
8. Allowing remote connections to an outside Windows server for remote administration:
   - Allow **TCP/UDP 3389** (Terminal server) from LAN subnet to **IP address of remote server**.
9. Allowing LAN to access windows shares on the DMZ, via NETBIOS/Microsoft-DS:
   - Allow **TCP/UDP 137** from LAN subnet (NETBIOS) to **DMZ subnet**.
   - Allow **TCP/UDP 138** from LAN subnet (NETBIOS) to **DMZ subnet**.
   - Allow **TCP/UDP 139** from LAN subnet (NETBIOS) to **DMZ subnet**.
   - Allow **TCP 445** from LAN subnet (NETBIOS) to **DMZ subnet**.

Outbound DMZ

By default, there are no rules on **OPT** interfaces.

1. Allowing servers to use Windows update or browse the **WAN**:
   - Allow **TCP 80** from DMZ subnet (HTTP) to **anywhere**.
   - Allow **TCP 443** from DMZ subnet (HTTP) to **anywhere**.
2. Allow users to connect to an external DNS server:
   - Allow **TCP/UDP 53** from DMZ subnet (DNS) to **IP address of the upstream DNS server(s)**
3. Allowing servers to use a remote NTP (time) server:
   - If using an upstream remote NTP server:
     - Allow **UDP 123** from DMZ subnet (NTP) to **IP address of remote NTP server**.
   - Otherwise:
     - Allow **UDP 123** from DMZ subnet (NTP) to **any**.

35.11.2 Setup isolating LAN and DMZ, each with unrestricted Internet access

The following setup can be used instead if outbound access is more lenient, but still controlled between local interfaces. This assumes all local networks are privately numbered, and that interfaces have already been configured.

Create an alias, **Firewall > Aliases** from the main menu, called `RFC1918` containing `192.168.0.0/16`, `172.16.0.0/12`, and `10.0.0.0/8`.

LAN Configuration

1. For DNS from the firewall:
   - Allow **TCP/UDP** from LAN subnet to **LAN Address port 53**.
2. For accessing the GUI:
   - Allow **TCP** from LAN subnet to **LAN address port 443**.
3. To ping the firewall from the LAN:
   - Allow **ICMP** from LAN subnet to **LAN address**.
4. If there is any traffic required from LAN to DMZ:
   - Allow any traffic required from **LAN** to **DMZ**.
5. Do not allow LAN to reach DMZ or other private networks:
   - Reject **Any** from LAN subnet to **RFC1918**.
6. For internet access:
   - Allow **Any** from LAN subnet to **any**.

DMZ Configuration

1. For DNS from the firewall:
   - Allow **TCP/UDP** from DMZ subnet to **DMZ Address port 53**.
2. For accessing the GUI (optional):
   - Allow **TCP** from DMZ subnet to **DMZ address port 443**.
3. To ping the firewall from the DMZ:
   - Allow **ICMP** from DMZ subnet to **DMZ address**.
4. If there is any traffic required from DMZ to LAN:
   - Allow any traffic required from **DMZ** to **LAN**.
5. Do not allow DMZ to reach LAN or other private networks:
   - Reject **Any** from DMZ subnet to **RFC1918**.
6. For Internet access:
   - Allow **Any** from DMZ subnet to **any**.

Additional Interfaces

Repeat the above pattern as needed.

35.12 External User Authentication Examples

There are countless ways to configure the user manager to connect to an external RADIUS or LDAP server, but there are some common methods that can be helpful to use as a guide. The following are all tested/working examples, but the server setup will likely vary from the example.

See also:

Authentication Servers

35.12.1 RADIUS Server Example

This example was made against FreeRADIUS but doing the same for Windows Server would be identical. See *Authenticating from Active Directory using RADIUS/NPS* for info on setting up a Windows Server for RADIUS.

This assumes the RADIUS server has already been configured to accept queries from this firewall as a client with a shared secret.

Descriptive Name
`ExCoRADIUS`

Type
Radius

Hostname or IP Address
`192.2.0.5`

Shared Secret
`secretsecret`

Services Offered
Authentication and Accounting

Authentication Port
`1812`

Accounting Port
`1813`

Authentication Timeout
`10`

35.12.2 OpenLDAP Example

In this example, the firewall is connecting back to an OpenLDAP server for the company.

Descriptive Name
`ExCoLDAP`

Type
LDAP

Hostname or IP Address
`ldap.example.com`

Port
`636`

Transport
SSL - Encrypted

Peer Certificate Authority
ExCo CA

Protocol Version
3

Search Scope
Entire Subtree, `dc=pfsense,dc=org`

Authentication Containers
`CN=pfsgroup;ou=people,dc=pfsense,dc=org`

Bind Credentials
Anonymous binds *Checked*

Initial Template
OpenLDAP

User Naming Attribute
`cn`

Group Naming Attribute
`cn`

Group Member Attribute
`memberUid`

RFC2307 Groups
Checked

Group Object Class
`posixGroup`

UTF8 Encode
Checked

Username Alterations
Unchecked

35.12.3 Active Directory LDAP Example

In this example, the firewall connects to an Active Directory structure in order to authenticate users for a VPN. The results are restricted to the **VPNUsers** group. Omit the **Extended Query** to accept any user.

Descriptive Name
`ExCoADVPN`

Type
LDAP

Hostname or IP Address
`192.0.2.230`

Port
`389`

Transport
TCP - Standard

Protocol Version
3

Search Scope
Entire Subtree, `DC=domain,DC=local`

Authentication Containers
`CN=Users,DC=domain,DC=local`

Extended Query
`memberOf=CN=VPNUsers,CN=Users,DC=example,DC=com`

Bind Credentials
Anonymous binds *Unchecked*

User DN
`CN=binduser,CN=Users,DC=domain,DC=local`

Password
`secretsecret`

Initial Template
Microsoft AD

User Naming Attribute
`samAccountName`

Group Naming Attribute
`cn`

Group Member Attribute
`memberOf`

This example uses plain TCP, but if the Certificate Authority for the AD structure is imported under the Certificate Manager, the connection can also use SSL as well by selecting that option and choosing the appropriate CA from the **Peer Certificate Authority** drop down, and setting the **Hostname** to the match the server certificate.

35.13 Using External Wireless Access Points

Most SOHO-style wireless routers can be used as an access point if a true Access Point (AP) is not available. If pfSense® software replaced an existing wireless router, the old router can still be used to handle the wireless portion of the network.

This type of deployment is popular for wireless because it is easier to keep the access point in a location with better signal coverage and take advantage of newer wireless hardware without relying on drivers in pfSense software. This way a network can still provide 802.11ac or later wireless connectivity secured by pfSense software at the edge, even though pfSense software itself is not capable of using wireless hardware for those newer standards.

This technique is also common with wireless equipment running *WRT, Tomato, or other custom firmware for use as dedicated access points rather than edge routers.

35.13.1 Turning a wireless router into an access point

To convert the wireless router into a wireless access point, follow these generic steps for any device. To find specifics for a particular wireless router, refer to its documentation.

Disable the DHCP server

Disable the DHCP server on the wireless router to prevent a conflict. pfSense software will handle this function for the network, and having two DHCP servers on the same broadcast domain will not function correctly.

Change the LAN IP address

A functional, unique, IP address on the access point is required for management purposes.

Change the LAN IP address on the wireless router to an unused IP address in the subnet where it will reside (commonly LAN). If the firewall running pfSense software replaced this wireless router, then the wireless router was probably using the same IP address now assigned to the firewall LAN interface, which conflicts.

Plug in the LAN interface

Most wireless routers bridge their wireless network to an internal LAN port or switch ports. This means the wireless segment will be on the same broadcast domain and IP subnet as the wired ports. For routers with an integrated switch, any of the LAN switch ports will typically work.

Note: Do not plug in the WAN or Internet port on the wireless router!

This will put the wireless network on a different broadcast domain from the rest of the network and the wireless router will perform NAT on the traffic between the wireless and LAN. This also results in double NAT of traffic between the wireless network and the Internet. This is an ugly design, and will lead to problems in some circumstances, especially if communication must occur between wireless and wired LAN clients.

Deciding where to connect the LAN interface from the wireless router depends on the chosen network design. The next sections cover options and considerations for selecting the best deployment style.

35.13.2 Bridging wireless to the LAN

One common means of deploying wireless is to plug the access point directly into the same switch as the LAN hosts, where the AP bridges the wireless clients onto the wired network. This works well, but offers limited control over the ability of the wireless clients to communicate with internal hosts.

See also:

See *Choosing Routing or Bridging* for details on bridging in this role.

35.13.3 Bridging wireless to an OPT interface

To keep wireless and wired networks on the same IP subnet and broadcast domain while also increasing control over wireless clients, add an OPT interface to the firewall for the access point and bridge the OPT interface to the LAN interface.

> **Warning:** Though bridging offers increased control over traffic, it also results in lower performance as all wireless traffic must pass through and be processed by the firewall. Typically, wireless speeds are low enough that this is not a major concern, but as wireless speeds improve the severity of the problem also increases.

This scenario is functionally equivalent to plugging the access point directly into the LAN switch, except pfSense software can filter traffic from the wireless network to provide protection to LAN hosts and vice versa.

Note: A configuration with the bridge assigned as LAN is optimal here, rather than only having the OPT bridged to the existing wired LAN.

35.13.4 Routed segment on an OPT interface

The wireless network can also be placed on a separate IP subnet if desired. This is done without bridging the OPT interface on pfSense, instead assigning it with an IP address in a separate subnet different from the LAN. This enables routing between internal and wireless networks, as permitted by the firewall ruleset. This is commonly done on larger networks, where multiple access points are plugged into a switch that is then plugged into the OPT interface on pfSense software. It is also preferable when wireless clients will be forced to connect to a VPN before allowing connections to internal network resources.

35.14 Using Software from FreeBSD

pfSense® software is based on FreeBSD, thus many familiar FreeBSD packages are available for use by veteran FreeBSD system administrators.

> **Warning:** Installing software this way **will** have unintended side effects.
>
> This action is **not recommended or supported** by Netgate.

Many parts of FreeBSD are not included in the base installation of pfSense software, so library and other issues can occur when attempting to use software installed in this manner. The pfSense software base installation does not include a compiler in the base system for many reasons, and as such software cannot be built locally. However, packages can be installed from FreeBSD the package repository.

35.14.1 Concerns/Warnings

Several important concerns must be considered by any administrator before deciding to install additional software, especially software that is not obtained from Netgate package repositories.

Security Concerns

Any extra software added to a firewall is a security problem, and must be evaluated fully before installation. If the need outweighs the risk, it may be worth taking. Official packages for pfSense software are not immune to this problem either. Any additional service is another potential attack vector.

Performance Concerns

Most hardware running pfSense software can handle the traffic load with which it is tasked. If the firewall hardware has horsepower to spare, it may not hurt performance to add additional software. That said, be mindful of the resources consumed by the added software.

Conflicting Software

If an installed package duplicates functionality found in the base system, or replaces a base system package with a newer version, it could cause unpredictable system instability. Ensure that the software does not already exist in pfSense.

Lack of Integration

Any extra software installed will not have GUI integration. For some, this is not a problem, but there have been people who expected to install a package and have a GUI magically appear for its configuration. These packages must be configured by hand. If this is a service, that means also making sure that any startup scripts accommodate the methods used by pfSense software.

Software can also install additional web pages that are not protected by the pfSense software authentication process. Test any installed software to ensure that it protects and filters access appropriately.

Lack of Backups

Packages installed in this manner must have any configuration or other needed files backed up manually.

These files will not be backed up during a normal backup and could be lost or changed during a firmware update. The add-on package described in *Backup Files and Directories with the Backup Package* is capable of backing up arbitrary files such as these.

35.14.2 Installing Packages

To install a package, the proper package site must be used. pfSense software is compiled against a specific FreeBSD branch, and has only a specific set of packages hosted on Netgate servers.

Packages located in the Netgate package repository, including some FreeBSD software packages that are not a part of the pfSense software distribution, can be installed using `pkg install` directly:

```
# pkg install screen
```

Or use a full URL to a `pkg add` to add them from the FreeBSD package servers:

```
# pkg add http://pkg.freebsd.org/FreeBSD:11:amd64/latest/All/tshark-3.2.6.txz
```

The pkg utility will download and install the package, along with its required dependencies.

Additionally, the full set of FreeBSD packages can be made available by editing /usr/local/etc/pkg/repos/pfSense.conf and changing the first line to:

```
FreeBSD: { enabled: yes }
```

Next, edit /usr/local/etc/pkg/repos/FreeBSD.conf and make the same change there:

```
FreeBSD: { enabled: yes }
```

It must be enabled in both places to function.

Warning: Adding software from FreeBSD package repositories **will** introduce problems with package dependencies, especially if a package depends on another piece of software that already exists on the firewall which may have been built with conflicting options. Take extreme caution when adding packages in this way.

Custom packages can also be built on another computer running FreeBSD and then the package file can be copied and installed on a firewall running pfSense software. Due to the complexity of this topic, it will not be covered here.

35.14.3 Maintaining Packages

The following command prints a list of all currently installed packages, including packages and components of the base system of pfSense software:

```
# pkg info
```

To delete an installed package, pass its full name or use a wildcard:

```
# pkg_delete iftop-1.0.p4
# pkg_delete pstree-\*
```

35.15 Using EAP and PEAP with FreeRADIUS

35.15.1 General EAP configuration

The default EAP settings will work in most situations (EAP-MD5, EAP-TLS, EAP-TTLS, EAP-PEAP) so there is no need to change them without any need. If EAP-TTLS or EAP-PEAP is used with VLAN assignment then set **Use Tunneled Reply** to *yes*:

To make the use of certificates more secure, check the **Common Name** of the client certificate against the username entered in **FreeRADIUS > Users**. For this set **Check Client Certificate CN** to *yes*:

Another option to increase security with certificates is to check the issuer of the client certificate against the CA certificate. This can be enabled with **Check Cert Issuer**, but then it is necessary to enter country, state, province and organization – case-sensitive – to match the CA.

FreeRADIUS by default allows many EAP types for authentication. In some environments only some strong EAP types (TLS, TTLS, PEAP, MSCHAPv2) may be allowed or weak types (MD5, GTC, LEAP) may be disallowed. Disable

the weak EAP types in FreeRADIUS using **Disable weak EAP types** so that FreeRADIUS rejects users which try to authenticate using such a weak method. If these types are disabled it does **not** affect the inner tunnel session in EAP-TTLS and EAP-PEAP. Further it is no problem to use a weak or cleartext method in the inner tunnel because if the outer tunnel uses one of the above call strong encryption types.

FreeRADIUS is multi-talented. It can handle almost all authentication types hosts send. So if weak encryption types such as MD5 and others are not disabled then the following will happen:

```
A client wants to authenticate using MD5 => FreeRADIUS will do that
A client wants to authenticate using LEAP => FreeRADIUS will do that
A client wants to authenticate with TLS or TTLS or PEAP or MSCHAPv2 or ... => FreeRADIUS␣
↪will do that
```

So **disable weak encryption** is checked then this disables MD5, GTC and LEAP. This will happen:

```
A client wants to authenticate using MD5 => FreeRADIUS will not do that because that was␣
↪disabled
A client wants to authenticate using LEAP => FreeRADIUS will not do that because that␣
↪was disabled
A client wants to authenticate with TLS or TTLS or PEAP or MSCHAPv2 or... => FreeRADIUS␣
↪will do that.
```

35.15.2 PEAP and MSCHAPv2

- FreeRADIUS package configuration in the pfSense® software GUI:
  - Configure an interface in **FreeRADIUS > Interfaces**
  - Create a **CA-Certificate** and a **Server-Certificate**.

    Choose pfSense Cert-Manager or FreeRADIUS Cert-Manager but **never** use the default certificates which come with FreeRADIUS after package installation!
  - Select the certificates in **FreeRADIUS > EAP**.

    If FreeRADIUS as Cert-Manager is selected then nothing must be changed. If pfSense Cert-Manager was chosen, then it must be enabled there, and the certs must be chosen from the pull-down menu. Click **Save**.
  - Add the WLAN Access Point in **FreeRADIUS > NAS/Clients**
  - Add a username/password in **FreeRADIUS > Users**
- WLAN Access-Point Configuration:
  - Change the wireless encryption to **WPA-Enterprise** or better **WPA2-Enterprise** with **TKIP** or better **AES/CCMP**
  - Do **not** use a passphrase but select **RADIUS** or **802.1X**. Enter the IP-Address of the FreeRADIUS server on pfSense software and the shared secret according to that what was entered in **FreeRADIUS > NAS/Clients**
- WLAN Device (Supplicant) Configuration:
  - Some devices can autoconfigure the Authentication and Encryption Method. If not choose **PEAP as encryption** and **MS-CHAPv2 as Authentication**. [1]
  - Connect to WLAN Access Point and the client will be prompted for username and password
  - Some devices auto-accept the CA-Certificate as valid. Often this CA-Certificate will first need to be accepted. This is the certificate created on pfSense software.

The most part of the "command line action" which is done in these tutorials can be done from FreeRADIUS GUI.

See also:

How-to with screenshots in German.

35.15.3 EAP-TLS

- pfSense software configuration:
  - Create a **CA**, a **Server-Certificate** and a **Client-Certificate**. Using **System > Certificates** is recommended.
- FreeRADIUS configuration:
  - Create an **interface**, add a **NAS/Client** and create a **user**. For this example, use *myuser* as username and *mypass* as password.
  - The EAP default options are working - read *FreeRADIUS package*.
  - Using **pfSense Cert-Manager** and selecting the CA and the server certificate is recommended.
  - Leave the password field empty
  - Download the `CA.crt` - **not the key** - from **System > Certificates**, **CAs** tab and Client `.p12` from **System > Certificates**, **Certificates** tab
- Client Requires password on `.p12`

  If a client fails to load the `.p12` without a password (and space does not work), then export the archive with a password as described in *Export Password-Protected Files or Use Different Encryption Options*.

35.16 Using Mobile One-Time Passwords with FreeRADIUS

Using Mobile-One-Time-Password (mOTP) with the *FreeRADIUS package*.

35.16.1 Enable Mobile-One-Time-Password (OTP)

This documentation covers many parts from installation, configuration, modification, and more from here.

A one time password is a password which can be only used one time and will be only usable within a short time period (10s). So it can be compared with the handling of tokens from RSA SecureID.

This kind of password generation makes sense in some scenarios but not in all. It probably makes no sense to use these passwords in the office – there shouldn't be any attacker. But the mOTP could make sense for Road Warriors. Most of them use a state-of-the-art mobile phone and a notebook to connect to the company VPN. For more take a look at the chapter below: *Miscellaneous configuration and hints*

pfSense® software configuration:

- Enable Mobile-One-Time-Password in **FreeRADIUS > Settings**

This will install the script in `/usr/local/etc/raddb/scripts/otpverify.sh`. To execute the script it needs the additional package `bash`. This will be installed automatically.

FreeRADIUS configuration:

- Create a user in **FreeRADIUS > Users**
- Enter a **username** but **do not** enter any password!
- Check **Enable Mobile-One-Time-Password For This User**

- Enter the **Init-Secret**.

  This value is created by the client (mobile device, mobile phone)

- Enter the **PIN**.

  Every time the user wants to generate a new password with their mobile token then they must enter the PIN, then the token generates a one-time password.

- The **Offset** should be zero by default.

  If the mobile token is in a different time zone than the FreeRADIUS server then correct the offset. If the mobile device changes its time zone automatically than there is no need to do this.

Miscellaneous configuration and hints:

- The configuration for different mobile devices can be found here.
- And here is the software for mobile devices.
- Please read the limitations to make sure how secure this is.
- mOTP will probably not work with EAP, CHAP, MS-CHAP.
- A generated one-time-password can be used for ~20 seconds. For example the Windows token generator generates new tokens every ~10 seconds. Perhaps other token generators use other time intervals for generating tokens. FreeRADIUS and the OTP script accept tokens which were generated within the last 20 seconds.

35.17 Using NAT and FTP without a Proxy

Firewalls running pfSense® software do not include an FTP Proxy, but this does not affect clients and servers as much as one might think.

Important: Use of FTP is strongly discouraged. It is an outdated protocol that transmits credentials and other data openly without encryption which is very insecure.

35.17.1 Client Behind pfSense Software

FTPS, or encrypted FTP, is not affected. The proxy could not have affected its traffic before.

A client on a LAN or other internal interface behind pfSense software will likely not notice any difference. Most clients, aside from the Microsoft command line FTP program, default to passive (PASV) FTP, where clients make outbound connections to servers.

Passive mode on the client will require access to random/high ports outbound, which could run afoul of a strict outbound ruleset. Environments with a security policy that requires strict outbound firewall rules likely would not be using FTP anyhow, as it transmits credentials without encryption.

Active mode FTP through NAT will not function as that relies on a proxy or similar mechanism. Use Passive mode instead. Another option is the recently added FTP Client Proxy package which leverages in FreeBSD to allow clients on local interfaces to reach remote FTP servers with active FTP.

Active mode FTP for a client that does not involve NAT (Client has a public IP address) should work so long as WAN rules pass the appropriate traffic back to the client. The client may have a configurable active port range to make that simpler.

35.17.2 Server Behind pfSense Software

FTPS, or encrypted FTP, is not affected. The proxy could not have affected its traffic before.

A server behind pfSense software would work fine with active mode, there would be no difference here. In active mode the server would make outbound connections back to the client, so as long as the firewall rules on the interface containing the server allow outbound connections, it will work.

A server behind pfSense software running in Passive mode will function but requires a few items to be configured:

- Port forwards or 1:1 NAT to forward not only port 21, but also the passive port range in to the server
- The passive port range must be configured on the server, corresponding to the range of ports forwarded in the previous step.
- The server may also need to be configured to account for NAT. Some clients will ignore private addresses in passive responses, so this may not be necessary.

Sample Configuration for vsftpd

In `vsftpd.conf`:

```
# Do not allow the client to use PORT
port_enable=NO
# Use the hostname in the PASV response (DNS must be setup and match!)
pasv_addr_resolve=YES
# Enable Passive Mode
pasv_enable=YES
# Set the passive port range (1000 ports)
pasv_min_port=20000
pasv_max_port=20999
```

35.18 Configuring pfSense Software for Online Gaming

This page provides information on using pfSense® software with online games.

See also:

The Netgate Forum often has a wide array of threads for specific games and consoles. Search there for consoles or games if they are not listed here.

- *NAT Types*
  - *Open / Type 1 / Type A*
  - *Moderate / Type 2 / Type B*
  - *Strict / Type 3 / Type C or D*
- *General Recommendations*
  - *Static Port*
  - *Endpoint-independent Port Restricted Cone NAT*
  - *Multiple players or devices behind one NAT device*

  - *Overcome NAT issues with UPnP IGD & PCP*
- *Specific Game Consoles*
  - *Nintendo Switch and Switch 2*
  - *Steam / Steam Deck*
  - *Xbox*
  - *Playstation*

35.18.1 NAT Types

Game consoles attempt to simplify NAT troubleshooting for players by referring to different NAT scenarios with alternative terminology. Later portions of this recipe may use or refer to these terms.

These terms generally fall into the following categories:

Open / Type 1 / Type A

All external connections are passed through to the gaming device. For example, by use of UPnP or a *1:1 NAT rule on the WAN IP address* pointing to a game console and a firewall rule allowing all traffic inbound to the console.

This is the **least secure** option and should rarely be necessary. However, this is the most compatible type of NAT for gaming, allowing connections from any other NAT type.

Danger: Allowing all traffic to a gaming console is a significant security concern. Should this be necessary, consider isolating the console on its own network segment and preventing it from communicating with any other internal host.

Moderate / Type 2 / Type B

When a local client initiates an outbound connection, NAT maintains the source port for subsequent connections. This can mean NAT uses whichever port the client chose (e.g. *Static Port NAT*), or that the firewall uses the same random external port for all client connections from a given source port (e.g. *Port Restricted Cone NAT*).

In this mode, when a client initiates a UDP connection outbound to a host, return traffic from that host coming back to the same port will pass through because of the existing connection state. By taking advantage of this behavior, peers can communicate directly without sending all traffic through the server. The server informs peers about the IP address and port number of other peers, those peers then initiate packets to each other using the supplied IP addresses and ports. Because the source port is known and does not change between connections, once those connection states are established the peers can communicate directly as their packets will match the states in the table. This behavior relies on all peers having moderate NAT or Open NAT.

Strict / Type 3 / Type C or D

When a local client initiates a connection, NAT uses a random source port and/or IP address. This is the default behavior of outbound NAT on pfSense software. Clients behind strict NAT can typically only connect to peers with Open NAT, as the connections from the client will not use known or predictable source ports.

35.18.2 General Recommendations

Many consoles, online services, and games require the use of *Static Port NAT*, *Port Restricted Cone NAT*, and/or *UPnP IGD & PCP*.

Static Port

Some games require at least *Moderate NAT* to facilitate direct peer communication, and one common way to achieve that is by using *Static Port NAT*. If there is only one game console on the local network, static port NAT may be sufficient. With static port, multiple consoles can conflict if they both attempt to use the same source port when communicating with the same server.

If a game has problems establishing a connection and there is only one local game console, try configuring static port for traffic coming from that console as described in *Static Port*.

Endpoint-independent Port Restricted Cone NAT

Port Restricted Cone NAT is another method to achieve *Moderate NAT* to facilitate direct peer communication. This method is more likely to be compatible with multiple local game consoles as NAT randomizes the external port but still maintains that same external port between subsequent remote connections.

If a game has problems establishing a connection and there are multiple local consoles, try configuring Port Restricted Cone NAT for traffic coming from the consoles as described in *Endpoint-independent Port Restricted Cone NAT*.

Note: This type of NAT is currently considered experimental.

Multiple players or devices behind one NAT device

Some games have issues where multiple players or devices are behind a single NAT device. These issues are specific to NAT in general, not pfSense software.

In particular, games which require moderate NAT may have issues using static port NAT. *Port Restricted Cone NAT* is much more likely to be compatible with multiple local consoles, except in cases where a game requires clients to use a specific source port number.

Multiple games and consoles which support UPnP generally work as expected, except in the case of Playstation. See the note below in *Playstation* for details.

Overcome NAT issues with UPnP IGD & PCP

Some games and consoles support Universal Plug-and-Play (UPnP IGD) and/or Port Control Protocol (PCP) to automatically configure NAT port forwards and firewall rules which allow inbound game traffic. Enabling UPnP IGD & PCP on pfSense software will typically allow these games to work with little or no intervention. However, allowing local devices to open ports in this way is also a significant security risk.

See also:

See *UPnP IGD & PCP* for more information on configuring and using UPnP IGD & PCP, and for information on potential security concerns.

Note: There was a bug with UPnP IGD and multiple client devices on the same network in versions before pfSense Plus software version 22.05 and pfSense CE software version 2.7.0. Update to a current supported release to obtain the fix.

35.18.3 Specific Game Consoles

This section contains recommendations for specific consoles types and games. If a game or console requires special handling but is not listed here, please *submit a documentation update*. Include a link to manufacturer documentation when possible.

Note: What works to make a single console/device work from behind a firewall may not work for multiple consoles/devices behind the same firewall.

What works to allow a client to play an online game may not work for hosting an online game.

Hosting an online game and connecting a client to the same online game from behind the same firewall may not work, but as with everything else, it varies by game/console.

Nintendo Switch and Switch 2

Nintendo Switch and Switch 2 consoles require either a *Static Port NAT* or *Port Restricted Cone NAT* setup for IPv4 NAT.

Static Port and Port Restricted Cone NAT both allow the console to achieve *NAT type B* and work with most games. Port Restricted Cone NAT is more likely to work with multiple consoles behind the same router, but results may vary by game.

Nintendo Switch 2 also supports IPv6, but how well that works depends on the game and whether peers also have IPv6 connectivity.

Tip: While the official game servers for older Nintendo consoles such as Nintendo Wii, Wii U, and 3DS/2DS have been shut down, there are third-party services which work similarly. These services generally require the same type of setup as Nintendo Switch and Switch 2 consoles.

Steam / Steam Deck

Varies by game, but typically *UPnP IGD & PCP* or manual port forwards are sufficient. Some games may require *Static Port* or *Port Restricted Cone NAT*.

Xbox

Modern Xbox consoles, including multiple consoles, work well with *UPnP IGD & PCP* in many cases.

Some games or situations may require *Static Port* or *Port Restricted Cone NAT*.

See also:

Xbox Support

Playstation

Single consoles work well with *UPnP IGD & PCP* though multiple consoles can be problematic. When multiple consoles are on the same network, Playstation devices do not automatically attempt to use a different port if they cannot use their preferred port.

Tip: In a mixed environment with Playstation and other console types, start the game on Playstation first so it can get the port it wants, then start other clients which will properly notice the port is in use and shift to alternate ports.

Some games or situations may require *Static Port* or *Port Restricted Cone NAT*

35.19 High Availability Configuration Example

This recipe describes a typical pfSense® software high availability (HA) cluster configuration with two nodes (primary and secondary) containing three interfaces: WAN, LAN, and Sync. This is logically equivalent to a deployment with two interfaces (WAN and LAN), with the Sync interface carrying synchronization data between the primary and secondary nodes.

Note: This example covers both IPv4 and IPv6 configuration.

If static IPv6 assignments are not available, set IPv6 to *None* on all interfaces and skip any configuration tasks specific to IPv6.

See also:

Review *High Availability* before following this recipe.

To troubleshoot issues with HA, see:

- *Troubleshooting High Availability*
- *Troubleshooting High Availability Clusters in Virtual Environments*
- *Troubleshooting VPN Connectivity to a High Availability Secondary Node*

35.19.1 Determine IP Address Assignments

The first task is to plan IP address assignments. A good strategy is to number the addresses as follows:

- Lowest usable IP address in the subnet as the CARP VIP address.
- The next IP address as the primary node interface IP address.
- The next IP address as the secondary node interface IP address.

This specific suggestion is optional. The best practice is to use a consistent and logical scheme to make design and administration simpler, but it is ultimately a matter of preference. For example, some network designers prefer to place routers at the end of a subnet instead of the beginning.

WAN IP Addresses

Select the WAN IP addresses from those assigned by the ISP. For this example, the assignments are listed in Table *WAN Interface IP Address Assignments*.

For IPv4, the WAN subnet for the HA pair is `198.51.100.0/24`. The WAN IPv4 addresses for the cluster are `198.51.100.200` through `198.51.100.202`.

For IPv6, the WAN prefix is `2001:db8::/64` and the ISP is routing a prefix of `2001:db8:1:df30::/60` for local networks. The WAN IPv6 addresses for the cluster are `2001:db8::200/64` through `2001:db8::202/64`.

Note: Though there is a whole `/64` to utilize, this example uses the same ending values in both IPv4 and IPv6 for consistency and to make similarities more apparent.

Table 1: WAN Interface IP Address Assignments

IP Address	Usage
`198.51.100.0/24`	WAN IPv4 subnet
`198.51.100.1`	WAN ISP IPv4 Gateway
`198.51.100.200/24`	CARP shared IPv4 address
`198.51.100.201/24`	Primary node WAN IPv4 address
`198.51.100.202/24`	Secondary node WAN IPv4 address
`2001:db8::/64`	WAN IPv6 Prefix
`2001:db8::1`	WAN ISP IPv6 Gateway
`2001:db8:1:df30::/60`	Routed IPv6 Prefix
`2001:db8::200/64`	CARP shared IPv6 address
`2001:db8::201/64`	Primary node WAN IPv6 address
`2001:db8::202/64`	Secondary node WAN IPv6 address

Note: In this case the ISP would route the IPv6 prefix (`2001:db8:1:df30::/60`) to the IPv4 WAN CARP VIP, `2001:db8::200`. Requirements vary by ISP. If the ISP prefers to route to a link-local address, add a CARP VIP on the WAN interface using a link-local address for this purpose.

LAN IP Addresses

The LAN IPv4 subnet is `192.168.1.0/24` and the LAN IPv6 prefix is `2001:db8:1:df30::/64`. For this example, the LAN IP addresses will be assigned as shown in Table *LAN Interface IP Address Assignments*.

Table 2: LAN Interface IP Address Assignments

IP Address	Usage
`192.168.1.1/24`	CARP shared IPv4 address
`192.168.1.2/24`	Primary node LAN IPv4 address
`192.168.1.3/24`	Secondary node LAN IPv4 address
`2001:db8:1:df30::1/64`	CARP shared IPv6 address
`2001:db8:1:df30::2/64`	Primary node LAN IPv6 address
`2001:db8:1:df30::3/64`	Secondary node LAN IPv6 address
`fe80::1:1/64`	CARP shared IPv6 Link-Local Address

Note: The CARP IPv6 link-local address in this example uses `fe80::1:1/64` as the `fe80::1/64` address is reserved for use by pfSense software in certain scenarios and can conflict. Using a different address avoids any potential problems.

Sync Interface IP Addresses

The IP addresses on the Sync interface are used only for communication between the firewalls. As there are no other devices on the Sync network, there is no need for a shared CARP VIP on the Sync interface, so it is not necessary to add a VIP on that interface.

This example uses the IPv4 subnet `172.16.1.0/24` and IPv6 prefix `2001:db8:1:df31::/64` for the Sync interface. The example only uses two IP addresses, but the IPv4 subnet uses a `/24` to be consistent with the other internal interface (LAN). For the last octet of the IP addresses, use the same last octet as the LAN IP addresses for consistency.

Table 3: Sync Interface IP Address Assignments

IP Address	Usage
`172.16.1.2/24`	Primary node Sync IPv4 address
`172.16.1.3/24`	Secondary node Sync IPv4 address
`2001:db8:1:df31::2/64`	Primary node Sync IPv6 address
`2001:db8:1:df31::3/64`	Secondary node Sync IPv6 address

35.19.2 Example Cluster Diagram

Figure *Example High Availability Cluster Network Diagram* shows the layout of this example HA cluster. The primary and secondary nodes each have identical connections to the WAN and LAN, and a crossover cable between them to connect the Sync interfaces.

See also:

In this basic example, the WAN switch and LAN switch are still potential single points of failure. Switching redundancy is covered in *Layer 2 Redundancy*.

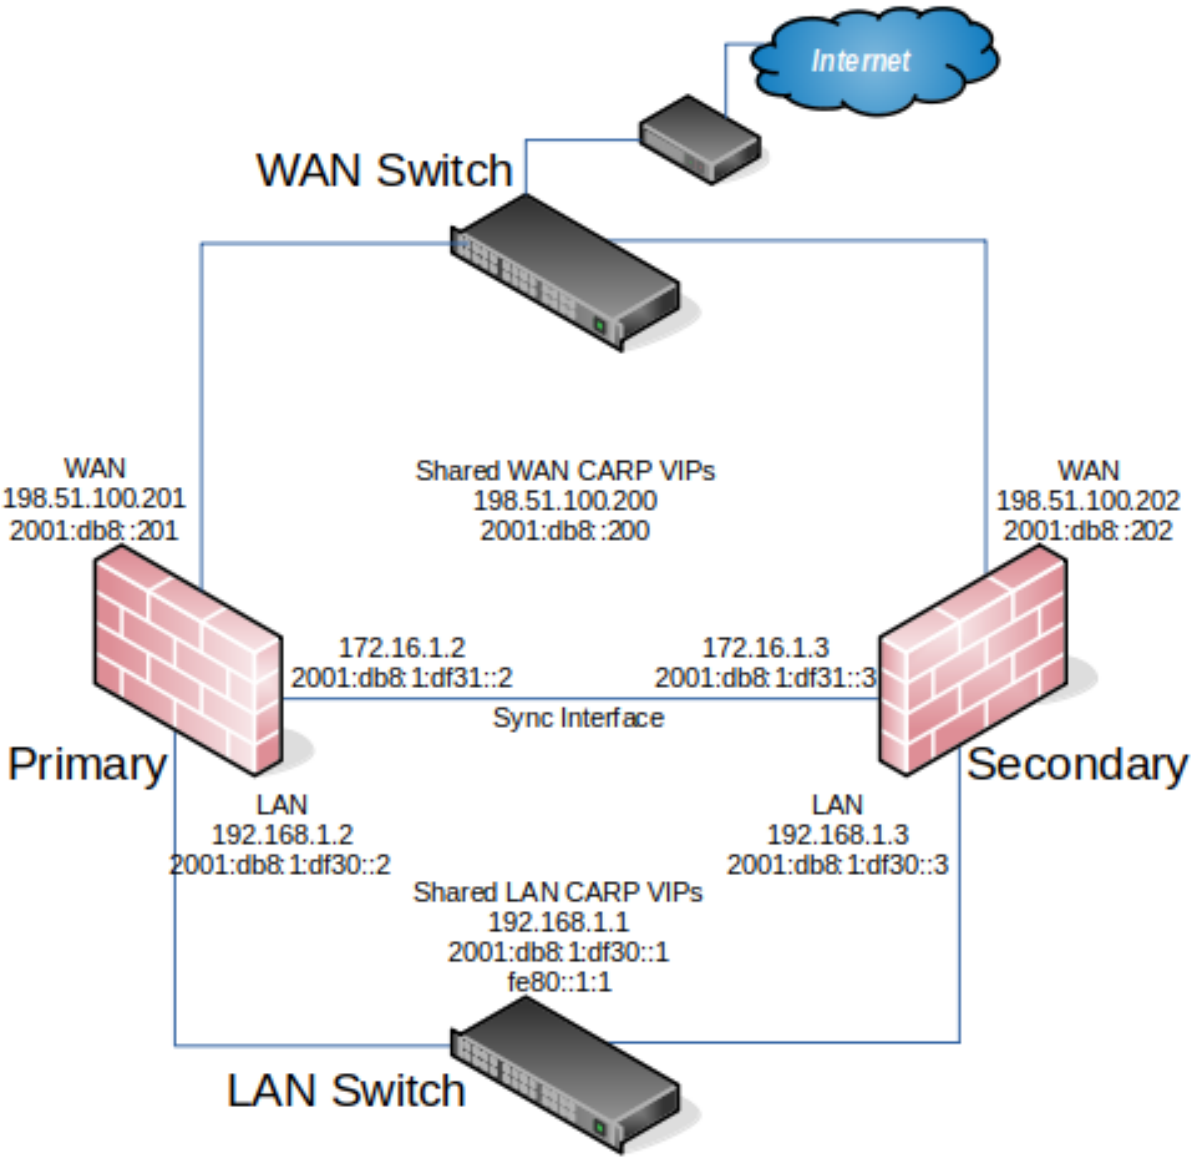


Fig. 6: Example High Availability Cluster Network Diagram

35.19.3 Cluster Configuration Basics

Both nodes require a few basic configuration steps before they are ready to be configured as a high availability cluster.

Danger: Do not connect both nodes to the same LAN (switch/layer 2) before both nodes have a non-conflicting LAN configuration.

Installation, interface assignment, and basic configuration

Install pfSense software on both nodes and assign the interfaces **identically** on both nodes.

Warning: Interfaces **must** be assigned in the same order on all nodes **exactly**. If the interface order is not identical, configuration synchronization and other tasks will not behave correctly. If any adjustments have been made to the interface assignments in the future, they **must** be replicated **identically** on both nodes.

After installation, connect to the GUI and use the Setup Wizard to configure each node with a unique hostname and non-conflicting static IP addresses.

See also:

Setup Wizard

For example, the primary node could be `firewall-a.example.com` and the secondary could be `firewall-b.example.com`, or a more personalized pair of names.

Note: Avoid naming the nodes `master` or `backup` since those are *states*, not *roles*. Consider naming them `primary` and `secondary` instead.

The default configuration on WAN is for DHCP, this must be changed to a static IP address configuration, such as the example WAN addresses in *WAN Interface IP Address Assignments*. Be sure to configure appropriate upstream gateways for IPv4 and IPv6.

The default LAN IPv4 address is `192.168.1.1`, but each node must be moved to its own unique and non-conflicting address. The default LAN IPv6 configuration is set to track WAN, which is not valid for HA, so it must be changed to a static configuration The example addresses for IPv4 and IPv6 are shown in *LAN Interface IP Address Assignments*.

Once each node has a unique LAN IPv4 and IPv6 address, then both nodes may be plugged into the same LAN switch.

Setup Sync Interface

Before proceeding, the Sync interfaces on the cluster nodes must be configured. *Sync Interface IP Address Assignments* lists the addresses to use for the Sync interfaces on each node. Once that has been completed on the primary node, perform it again on the secondary node with the appropriate address values.

To complete the Sync interface configuration, add firewall rules on both nodes to allow synchronization.

At a minimum, the firewall rules must pass configuration synchronization traffic (by default, HTTPS on port TCP `443`), pfsync traffic, and Kea DHCP HA traffic on TCP ports `8765` and `8766`. In most cases, a simple "allow all" style rule is sufficient, but using specific rules is a better practice.

When complete, the rules will look like the example in figure *Example Sync Interface Firewall Rules*, which also includes a rule to allow ICMP echo (ping) for diagnostic purposes.

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
☐	✓ 0/0 B	IPv4 TCP	SYNC subnets	*	SYNC address	443 (HTTPS)	*	none		Allow configuration synchronization	
☐	✓ 0/67 KiB	IPv4 PFSYNC	SYNC subnets	*	*	*	*	none		Allow state synchronization	
☐	✓ 0/0 B	IPv4 ICMP any	SYNC subnets	*	SYNC subnets	*	*	none		Allow ICMP for diagnostics	
☐	✓ 2/32.47 MiB	IPv4+6 TCP	SYNC subnets	*	SYNC address	8765 - 8766	*	none		Allow DHCP Sync	

Fig. 7: Example Sync Interface Firewall Rules

The secondary does not need all of these rules initially, only a rule to allow traffic to the GUI for XMLRPC to function. The full set of rules will synchronize via XMLRPC configuration synchronization later.

35.19.4 Configure State Synchronization (pfsync)

State synchronization using pfsync must be configured on both the primary and secondary nodes to function.

See also:

State Synchronization Settings (pfsync)

First on the primary node and then on the secondary, perform the following:

- Navigate to **System > High Availability**
- Check **Synchronize States**
- Set **Synchronize Interface** to *SYNC*
- Set a custom **Filter Host ID**

  For example, `1` on the primary node, `2` on the secondary node.
- Set **pfsync Synchronize Peer IP** to the Sync interface IPv4 address of the **other** node

  Set this to `172.16.1.3` on the primary node, or `172.16.1.2` on the secondary node
- Click **Save**

35.19.5 Configure Configuration Synchronization (XMLRPC)

Warning: Configuration synchronization **must only be configured on the primary node**. Never activate options in this section on the secondary node of a two-member cluster.

See also:

Configuration Synchronization Settings (XMLPRC Sync)

On the primary node **only**, perform the following:

- Navigate to **System > High Availability**
- Set **Synchronize Config to IP** to the Sync interface IPv4 address on the secondary node, `172.16.1.3`
- Set **Remote System Username** to `admin`

Note: This must either be the `admin` account or a user on both nodes with the "System - HA node sync" privilege.

- Set **Remote System Password** to the admin user account password, and repeat the value in the confirmation box.
- Check the boxes for each area to synchronize to the secondary node.

  For this guide, as with most configurations, all boxes are checked.

Tip: Use the **Toggle All** button to select all the options at once, rather than selecting them individually.

- Click **Save**

As a quick confirmation that the synchronization worked, on the secondary node navigate to **Firewall > Rules**, **SYNC** tab. The rules entered on the primary should now be present on the tab, and the temporary rule has been removed.

The two nodes are now linked for configuration synchronization! Configured areas will synchronize to the secondary node whenever a change is made on the primary node.

Warning: Do not make changes to the secondary in areas set to be synchronized! These changes will be overwritten the next time the primary node performs a synchronization.

35.19.6 Configuring the CARP Virtual IPs

With configuration synchronization in place, the CARP Virtual IP addresses need only be added to the primary node, and they will automatically copy to the secondary node.

Each VIP needs a unique VHID on a given broadcast domain/layer 2. See *CARP VIP VHID Assignments* for the VHIDs to use in this example.

Table 4: CARP VIP VHID Assignments

CARP VIP Address	Interface	VHID
`198.51.100.200/24`	WAN	200
`2001:db8::200/64`	WAN	201
`192.168.1.1/24`	LAN	1
`2001:db8:1:df30::1/64`	LAN	2
`fe80::1:1/64`	LAN	3

- Navigate to **Firewall > Virtual IPs** on the primary node to manage CARP VIPs
- Click  **Add** at the top of the list to create a new VIP.

Note: Each interface handling user traffic must have a CARP VIP, in this case WAN and LAN.

- Fill in the VIP options as described below. See *CARP VIP VHID Assignments* for specific settings unique to each VIP in this example. Refer to *VIP Configuration Options* for more information on the options in general.

  Type
  CARP

Interface
The interface upon which the VIP resides, e.g. *WAN*

Address(es)
The IP address and subnet mask/prefix value for the VIP. Consult *CARP VIP VHID Assignments* for the address values in this example.

Note: The subnet mask/prefix must match the subnet mask on the interface IP addresses. For example, the first VIP `198.51.100.200` and `24` on WAN (See *WAN Interface IP Address Assignments*).

Virtual IP Password
Password for the CARP VIP. Use a random value.

This need only match between the two nodes, which will be handled by synchronization. The password and confirm password box must both be filled in and they must match.

VHID Group
Virtual ID for the CARP VIP. Consult *CARP VIP VHID Assignments* for the VHID to use for each address in this example.

Tip: A common tactic is to make the VHID match the last octet of the IP address, e.g. `200` for an address ending in `.200`.

Advertising Frequency
How often the node sends CARP heartbeats on its interface.

Base
Whole seconds between Heartbeats, typically *1*.

Skew
Fractions of a second (1/256th increments).

Note: A primary node is typically set to *0* or *1*, secondary nodes will be *100* or higher. This adjustment is handled automatically by XMLRPC synchronization.

Description
Text to identify the VIP, such as `WAN CARP VIP`.

The above description used the WAN VIP as an example. Repeat the procedure for every required CARP VIP (See *CARP VIP VHID Assignments*).

If there are any additional IP addresses in the WAN subnet that will be used for purposes such as 1:1 NAT, port forwards, VPNs, etc., add them now.

Tip: For HA setups designed with multiple CARP VIPs of the same address family on the same interface, instead consider adding the additional CARP VIPs as IP alias VIPs which use a single CARP VIP as their parent interface. This reduces the number of CARP VIP heartbeats on a network segment and also allows the VIPs to fail as a group. For details, see *Using IP Aliases to Reduce Heartbeat Traffic*.

Click **Apply Changes** after completing the VIP configuration.

After adding VIPs, check **Firewall > Virtual IPs** on the secondary node to ensure that the VIPs synchronized as expected.

The Virtual IP addresses on both nodes will look like *CARP Virtual IP Address List* if the process was successful.

Virtual IP Address

Virtual IP address	Interface	Type	Description	Actions
198.51.100.200/24 (vhid: 200)	WAN	CARP	WAN CARP IPv4 VIP	
192.168.1.1/24 (vhid: 1)	LAN	CARP	LAN CARP IPv4 VIP	
2001:db8::200/64 (vhid: 201)	WAN	CARP	WAN CARP IPv6 VIP	
2001:db8:1:df30::1/64 (vhid: 2)	LAN	CARP	LAN CARP IPv6 VIP	
fe80::1:1/64 (vhid: 3)	LAN	CARP	LAN CARP IPv6 Link-Local VIP	

Fig. 8: CARP Virtual IP Address List

35.19.7 Configure Outbound NAT for CARP

The next step is to configure outbound NAT so that the firewall translates IPv4 traffic from clients on the LAN to the shared IPv4 CARP VIP address on WAN as the address as it exits.

- Navigate to **Firewall > NAT**, **Outbound** tab
- Click to select **Hybrid Outbound NAT rule generation**
- Click **Save**

Hybrid mode allows the default NAT rules to stay in place, but also allows the default rules to be overridden by manually created rules. This way it becomes unnecessary to manage entries for the firewall itself (which **should not** use the CARP VIP for NAT) and it can also make the process less disruptive when adding new interfaces later.

Create outbound NAT rules for internal subnet sources to work with the CARP IP address.

- Click  below the **Mappings** section to add a new rule
- Configure the rule as follows:

  Interface
  WAN

  Address Family
  IPv4

  Source
  LAN Subnets

  Translation, Address
  Choose the WAN CARP IPv4 VIP from the **Address** drop-down, *198.51.100.200*

  Description
  `LAN to WAN`

  Leave all other options at the default values.

- Click **Save**
- Click  below the **Mappings** section to add another new rule to the top of the list for IPsec pass-through.

Note: This is optional and may be omitted if there will not be any LAN clients connecting to external IPsec servers which lack NAT-T capabilities.

- Configure the rule as follows:

  Interface
  WAN

  Address Family
  IPv4

  Source
  LAN Subnets

  Destination

  Type
  Any

  Port
  `500`

  Translation

  Address
  Choose the WAN CARP IPv4 VIP from the **Address** drop-down, *198.51.100.200*

  Port or Range
  Check **Static Port**

  Description
  `ISAKMP - LAN to WAN`

  Leave all other options at the default values.

- Click **Save**
- Click **Apply Changes**

Warning: If an additional local interface is added later, such as a second LAN, DMZ, etc., and that interface uses private IP addresses, then additional outbound NAT rules must be added to translate those sources. Similarly, if additional WANs are added later, they will also need a set of outbound NAT rules.

Danger: Do not use an overly permissive source on outbound NAT rules! If an outbound NAT rule matches traffic from the firewall using a VIP, it can break several protocols and cause instability. Use only local/internal networks as the source on NAT rules, or an alias containing those local private networks.

Any rules using **Localhost** as a source must use an interface address for translation and **not** a CARP VIP.

When complete, the rule changes will look like those found in *Outbound NAT Rules for LAN with CARP VIP*

Mappings

	Interface	Source	Source Port	Destination	Destination Port	NAT Address	NAT Port	Static Port	Description	Actions
☐ ✓	WAN	LAN subnets	*	*	500 (ISAKMP)	198.51.100.200 (WAN CARP IPv4 VIP)	*	✓	ISAKMP - LAN to WAN	
☐ ✓	WAN	LAN subnets	*	*	*	198.51.100.200 (WAN CARP IPv4 VIP)	*		LAN to WAN	

Fig. 9: Outbound NAT Rules for LAN with CARP VIP

35.19.8 Modifying the DHCP Server

The DHCP server daemons on the cluster nodes need adjustments so that they can work together. The changes will synchronize from the primary to the secondary, so as with the VIPs and Outbound NAT, these changes need only be made on the primary node.

Note: This guide assumes the Kea DHCP Backend is active as it is the only backend which is capable of HA for both IPv4 and IPv6 (*Server Backend*).

Set the Backend

On **both** nodes, ensure the Kea DHCP backend is active:

- Navigate to **System > Advanced**, **Networking** tab
- Set the **Server Backend** to **Kea DHCP** in the **DHCP Options** section if it is not already selected
- Click **Save** if the value was changed

Configure DHCP for IPv4

On the **primary** node only, enable HA for IPv4 DHCP in the *Kea Settings*:

- Navigate to **Services > DHCP Server**, **Settings** tab
- Check **Enable high availability**
- Set **Node Role** to **Primary**
- Set **Local Name** to a custom value, e.g. `ha-primary`
- Set **Local Address** to the IP address of the sync interface on the primary node, e.g. `172.16.1.2`
- Set **Remote Name** to the name of the secondary node, e.g. `ha-secondary`.
- Set **Remote Address** to the IP address of the Sync interface on the secondary node, e.g. `172.16.1.3`
- Configure TLS to encrypt the DHCP HA traffic (optional)

  The Sync interface is directly connected in this example, so there is no risk of the traffic being intercepted, but using encryption is still a best practice. Consult *TLS Transport* for more information on configuring TLS for DHCP HA. Note that TLS settings do not synchronize and must be configured separately on each node.
- Click **Save**

Note: XMLRPC configuration synchronization will adjust these options appropriately when copying the settings to the secondary node.

Still on the **primary** node only, configure the LAN interface DHCP settings to use the LAN CARP IPv4 VIP:

- Navigate to **Services > DHCP Server**, **LAN** tab
- Set the **DNS Server** to the LAN CARP VIP, here `192.168.1.1`
- Set the **Gateway** to the LAN CARP VIP, here `192.168.1.1`
- Click **Save**

Setting the **DNS Server** and **Gateway** to a CARP VIP ensures that the local clients are talking to the failover address and not directly to either node. This way if the primary fails, the local clients will continue talking to the secondary node.

Configure DHCP for IPv6

Configure DHCPv6 HA

On the **primary** node only, enable HA for IPv6 DHCP in the *Kea Settings*:

- Navigate to **Services > DHCPv6 Server**, **Settings** tab
- Check **Enable high availability**
- Set **Node Role** to **Primary**
- Set **Local Name** to a custom value, e.g. `ha-primary`
- Set **Local Address** to the IP address of the sync interface on the primary node, e.g. `2001:db8:1:df31::2`
- Set **Remote Name** to the name of the secondary node, e.g. `ha-secondary`.
- Set **Remote Address** to the IP address of the Sync interface on the secondary node, e.g. `2001:db8:1:df31::3`
- Configure TLS to encrypt the DHCP HA traffic (optional)

  The Sync interface is directly connected in this example, so there is no risk of the traffic being intercepted, but using encryption is still a best practice. Consult *TLS Transport* for more information on configuring TLS for DHCP HA. Note that TLS settings do not synchronize and must be configured separately on each node.
- Click **Save**

Note: XMLRPC configuration synchronization will adjust these options appropriately when copying the settings to the secondary node.

Configure DHCPv6 Interface Settings

Still on the **primary** node only, configure the LAN interface DHCPv6 settings to use the LAN IPv6 CARP VIP for DNS.

- Navigate to **Services > DHCPv6 Server**, **LAN** tab
- Set the **DNS Server** to the LAN CARP VIP, here `2001:db8:1:df30::1`
- Click **Save**

Configure IPv6 Router Advertisements

Finally, configure IPv6 Router Advertisements (RA) to use the LAN CARP IPv6 Link-Local VIP for the IPv6 gateway:

- Navigate to **Services > Router Advertisement**, **LAN** tab
- Set **Router Mode** to *Managed* so that clients will use DHCPv6 to obtain addresses and other settings.
- Set **RA Interface** to the LAN CARP IPv6 Link-Local VIP, *fe80::1:1*
- Check **Enable DNS**

- Check **Mirror DHCPv6** to have the RA daemon advertise the same DNS settings as the DHCPv6 server
- Click **Save**

35.19.9 Finish Up & Test

At this point the cluster is configured and ready to use. Review the testing procedure in *Verifying Failover Functionality* to confirm it is operating properly.

Once everything tests OK, make configuration backups of both nodes.

35.20 Converting High Availability DHCP from ISC to Kea

This document covers how to convert an existing pfSense® software High Availability (HA) setup from the ISC DHCP backend to the Kea DHCP backend and also explains differences in how each of these backends implements HA failover.

Warning: The ISC DHCP backend is deprecated and will eventually be removed, so converting to the Kea DHCP backend is the best practice.

See also:

- *Server Backend*
- *High Availability Configuration Example*

The Kea implementation of HA is fundamentally different from ISC and as a consequence, the HA/failover settings are not compatible between ISC DHCP and Kea DHCP. The conversion process is manual and largely the same as setting up high availability for Kea DHCP from scratch. However, since the Kea HA configuration is much simpler, it is not a lengthy or difficult process.

Converting High Availability DHCP from ISC to Kea

- *Converting High Availability DHCP from ISC to Kea*
  - *DHCP Backend Differences*
    - *Backend Difference Summary*
  - *Status Differences*
  - *Conversion Process*
    - *Converting DHCPv4 from ISC to Kea*
    - *Adding DHCPv6 HA after switching to Kea*

35.20.1 DHCP Backend Differences

There are important differences in high availability operation between the ISC DHCP backend and the Kea DHCP backend, including:

- Kea DHCP operates HA in "hot standby" mode where the primary node serves leases exclusively unless it fails, in which case the standby node serves leases until the primary recovers. Both nodes share lease data as needed. This is similar to the HA behavior of pfSense software in general.

  ISC DHCP behaves closer to an "active/active" load balanced mode which has each node coordinate and serve part of each pool simultaneously and exchange lease data in both directions.

- Kea has a global setting declaring the role of a node as primary or standby.

  ISC attempts to automatically determine the role by inspecting the properties of CARP VIPs which is not always accurate.

- Kea HA serves leases from a secondary node if it boots while the primary node is offline.

  ISC fails to serve leases from a secondary unless it sees the primary node online first.

- Kea DHCP provides HA for IPv6 in DHCPv6.

  ISC is not capable of providing IPv6 HA.

  - As a part of HA for DHCPv6, XMLRPC configuration synchronization can copy Kea DHCPv6 settings between HA nodes.
  - XMLRPC configuration synchronization for Kea DHCPv6 also synchronizes Router Advertisement settings.

  Warning: If an HA cluster had IPv6 configured separately on each node before converting to Kea, the configuration on the secondary node will be overwritten once XMLRPC configuration synchronization is enabled for DHCPv6 on the primary node. If a pool was split between two nodes in this manner, it can be combined into one large pool after converting.

- Kea has one set of global settings for HA per node.

  ISC requires per-interface failover configuration.

  - Despite being configured in different ways, XMLRPC synchronization of HA DHCP settings for Kea and ISC both automatically adjust the values needed for settings on the primary and secondary.

- Kea uses one address per node for HA traffic (lease synchronization, heartbeats, etc.).

  ISC uses manually configured failover addresses on every DHCP interface separately.

  - The best practice with Kea is to use the Sync interface for DHCP HA traffic since it uses a single address. This is safer and more secure as it does not need to flow over every interface separately where it may be exposed to end-user networks.
  - Similarly, firewall rules for Kea only need to be on one interface (e.g. Sync), where ISC DHCP may need firewall rules on every DHCP server interface.

- Kea can use IPv4 or IPv6 for HA traffic in both DHCPv4 and DHCPv6.

  ISC can only use IPv4.

- Kea lease synchronization copies hostnames between nodes.

  ISC lease synchronization does not copy hostnames.

- Kea tracks HA status globally for each node.

  ISC tracks HA status separately for each address pool.

- Kea can optionally encrypt HA traffic using TLS and can also optionally require a client certificate to validate it is speaking to the expected peer.

  ISC does not offer encryption or authentication of failover traffic.

  - The TLS settings for HA do not synchronize via XMLRPC configuration synchronization as each node may need to use different certificates.

Backend Difference Summary

The following table summarizes the differences between the backends:

Table 5: High Availability DHCP Backend Comparison

Feature	Kea	ISC
Supported Upstream	Yes	No (Deprecated)
HA Style	Hot Standby	Active/Active
Active/Passive Role	Manual	Automatic
Secondary Only Boot	Yes	No
DHCPv6 HA	Yes	No
DHCPv6+RA XMLRPC Sync	Yes	No
HA Config	Global	Per-Interface
IP Addresses for HA	One per Node	One per Interface per Node
HA Interfaces	Sync	All Enabled for DHCP
HA Transports	IPv4 or IPv6	IPv4 Only
HA Copies Hostnames	Yes	No
HA Status	Per Node	Per Pool
TLS Encryption	Yes (Optional)	No
TLS Authentication	Yes (Optional)	No

35.20.2 Status Differences

The DHCP Lease status for failover information changed significantly between ISC DHCP and Kea DHCP.

The ISC DHCP Lease status page shows a list of **Failover Groups** at the top of the page, as described in *Pool Status (HA/Failover) – ISC DHCP Only*. Each pool has a state for both nodes in the group and the last time the state changed. The information is fairly plain, with only the text of the state fields changing to indicate status.

In contrast, the Kea DHCP HA status is much more user-friendly. It has moved to the bottom of the DHCP Lease status page, as shown in *High Availability Status – Kea DHCP Only*. The Kea DHCP HA status section includes an easy to interpret icon indicating the current state along with information about each node and its status, and a timer indicating the last received heartbeat.

35.20.3 Conversion Process

Converting DHCPv4 from ISC to Kea

- Switch the backend to Kea (*Server Backend*)
- Add firewall rules on the primary node Sync interface which pass IPv4 Kea HA traffic (*Setup Sync Interface*)
- Enable and configure Kea DHCP HA on the primary node (*Configure DHCP for IPv4*)
- Check the DHCP status and ensure the two nodes are communicating properly (*High Availability Status – Kea DHCP Only*)
- Remove any old firewall rules which explicitly passed ISC DHCP failover traffic (TCP ports `519` and `520`). These could be on interface tabs, interface group tabs, floating rules, etc.

Adding DHCPv6 HA after switching to Kea

- Add firewall rules on the primary node Sync interface which pass IPv6 Kea HA traffic (*Setup Sync Interface*)
- Add an IPv6 Link-Local CARP VIP for LAN if one does not already exist, such as `fe80::1:1/64` (*LAN IP Addresses*)
- Enable and configure Kea DHCPv6 HA on the primary node (*Configure DHCPv6 HA*)
- Configure DHCPv6 interfaces to use a CARP VIP for DNS (*Configure DHCPv6 Interface Settings*)
- Configure Router Advertisements for HA using the link-local CARP VIP as the **RA Interface** (*Configure IPv6 Router Advertisements*)
- Enable XMLRPC Synchronization for DHCPv6 (*Configure Configuration Synchronization (XMLRPC)*)
- Adjust the DHCPv6 configuration to account for changes in the pool addresses if the HA cluster had DHCPv6 configured manually without synchronization.
- Check the DHCP status and ensure the two nodes are communicating properly (*High Availability Status – Kea DHCP Only*)

35.21 High Availability Configuration Example with Multi-WAN

High availability (HA) can be deployed for firewall redundancy in a multi-WAN configuration. This section details the VIP and NAT configuration needed for a dual WAN HA deployment.

Note: This recipe is a supplement to *High Availability Configuration Example*. Read through that recipe before proceeding. This document only covers topics specific to high availability and multi-WAN and is not a complete HA configuration guide.

See also:

- *Multiple WAN Connections*

35.21.1 Determine IP Address Assignments

This example uses four IPv4 addresses on each WAN. Each firewall needs an IP address, plus one CARP VIP for Outbound NAT, plus an additional CARP VIP for a 1:1 NAT entry that will be used for an internal mail server in the DMZ segment.

WAN and WAN2 IP Addresses

Table *WAN IP Addresses* shows the IP addresses for both WANs. In most environments these will be public IP addresses.

Table 6: WAN IP Addresses

IP Address	Usage
`198.51.100.200`	Shared CARP VIP for Outbound NAT
`198.51.100.201`	Primary firewall WAN
`198.51.100.202`	Secondary firewall WAN
`198.51.100.203`	Shared CARP VIP for 1:1 NAT

Table 7: WAN2 IP Addresses

IP Address	Usage
`203.0.113.10`	Shared CARP VIP for Outbound NAT
`203.0.113.11`	Primary firewall WAN2
`203.0.113.12`	Secondary firewall WAN2
`203.0.113.13`	Shared CARP VIP for 1:1 NAT

LAN Addresses

The LAN subnet is `192.168.1.0/24`. For this example, the LAN IP addresses are assigned as described in Table *LAN IP Address Assignments*:

Table 8: LAN IP Address Assignments

IP Address	Usage
`192.168.1.1`	CARP shared LAN VIP
`192.168.1.2`	Primary firewall LAN
`192.168.1.3`	Secondary firewall LAN

DMZ Addresses

The DMZ subnet is `192.168.2.0/24`. For this example, the DMZ IP addresses are assigned as described in Table *DMZ IP Address Assignments*:

Table 9: DMZ IP Address Assignments

IP Address	Usage
`192.168.2.1`	CARP shared DMZ VIP
`192.168.2.2`	Primary firewall DMZ
`192.168.2.3`	Secondary firewall DMZ

Sync Interface Addressing

There will be no shared CARP VIP on the Sync interface because there is no need for one. These IP addresses are used only for communication between the HA nodes. For this example, `172.16.1.0/24` is the Sync subnet. The two nodes only require one IP address each (two total), but the example uses a `/24` size subnet to be consistent with the other internal interfaces. For the last octet of the IP addresses, use the same last octet as the LAN IP address on that node for consistency.

Table 10: Sync IP Address Assignments

IP Address	Usage
`172.16.1.2`	Primary firewall Sync
`172.16.1.3`	Secondary firewall Sync

35.21.2 NAT Configuration

The NAT configuration when using HA with Multi-WAN is the same as HA with a single WAN, except the rules are repeated, so there is a set for each WAN. Ensure that only CARP VIPs are used for inbound traffic or routing.

See also:

See *Network Address Translation* for more information on NAT configuration.

35.21.3 Firewall Configuration

With Multi-WAN a firewall rule must be in place to pass traffic to local networks using the default gateway. Otherwise, when traffic attempts to reach the CARP address or from LAN to DMZ it will instead go out a WAN connection.

Add a rule at the top of the firewall rules for all internal interfaces which will direct traffic for all local networks to the default gateway. The important part is the gateway **must** be *default* for this rule and not one of the failover or load balance gateway groups. The destination for this rule would be the local LAN network, or an alias containing any locally reachable networks.

See also:

- *Policy Routing Configuration*

35.21.4 Multi-WAN HA with DMZ Diagram

Due to the additional WAN and DMZ elements, a diagram of this layout is much more complex as can be seen in Figure *Diagram of Multi-WAN HA with DMZ*.

35.22 High Availability Configuration Example without NAT

High Availability (HA) can provide redundancy for routed public subnets. This section describes this type of configuration, which is common in large networks, ISP, wireless ISP networks, and data center environments.

Note: This recipe is a supplement to *High Availability Configuration Example*. Read through that recipe before proceeding. This document only covers topics specific to high availability without NAT and is not a complete configuration guide.

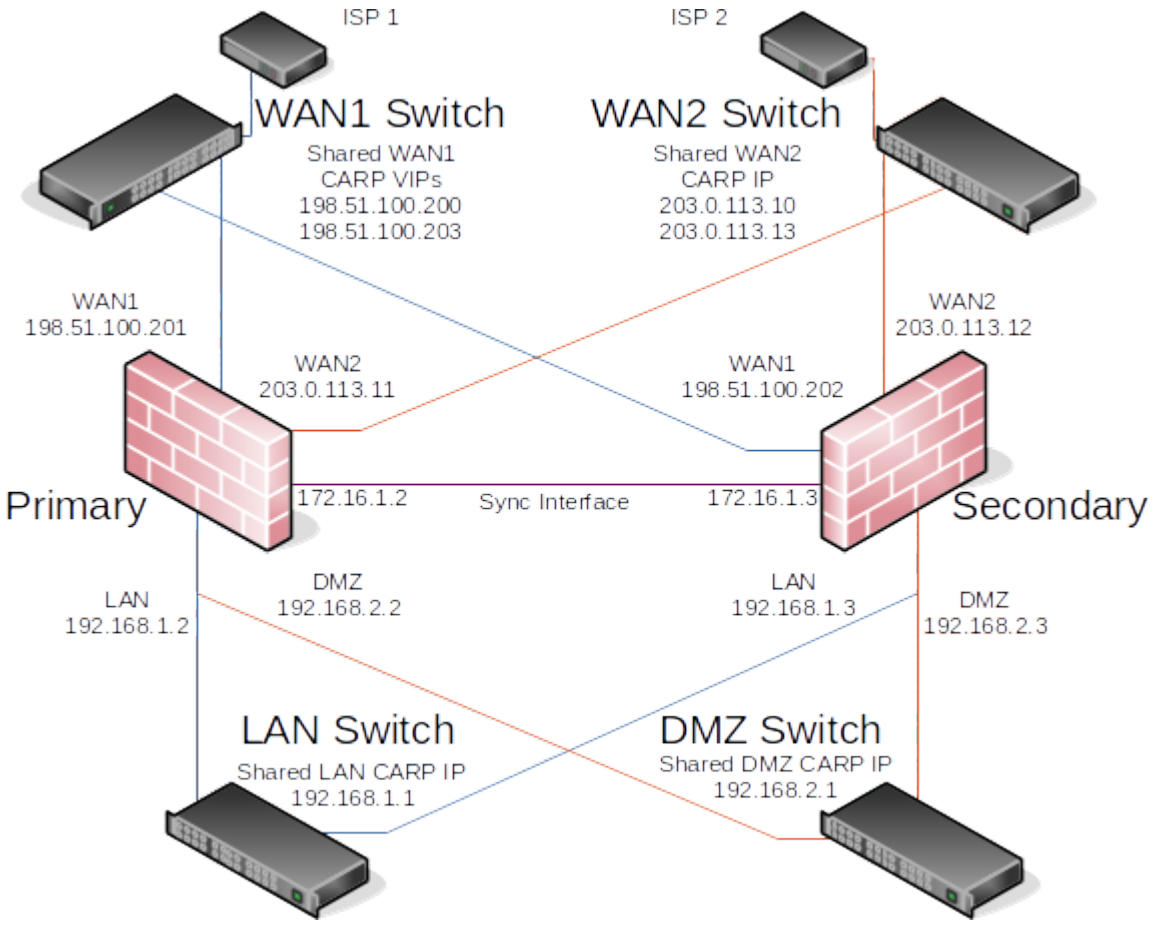


Fig. 10: Diagram of Multi-WAN HA with DMZ

As mentioned in *High Availability Configuration Example*, CARP VIPs are the only type of VIPs which provide redundancy for addresses directly handled by the firewall, and CARP VIPs can only be used in conjunction with NAT or services on the firewall itself. However, sources can send traffic to the CARP VIPs, such as delivering traffic from routed subnets, and it will reach the active node for processing even if it isn't involved in NAT or local services.

Note: This recipe only covers IPv4 without NAT as the IPv6 configuration in *High Availability Configuration Example* does not involve NAT, so there are no differences to address.

35.22.1 Public IPv4 Address Assignments

HA requires at least a `/29` block of public IPv4 addresses for the WAN side of the firewall, which provides six usable IPv4 addresses. Each firewall requires one IPv4 address, and at least one CARP VIP on the WAN side. A two node deployment only requires three addresses, but this is the smallest subnet that will accommodate three usable IPv4 addresses since the upstream router consumes one, and two are lost to the network and broadcast addresses.

Additional public IPv4 address subnets for internal use must be routed to one of the WAN CARP VIPs by the ISP, data center, or upstream router. When the upstream source routes additional subnets to a CARP VIP, the routing will not be dependent upon a single node. The example configuration depicted in this section uses a `/24` public IPv4 subnet and splits it into two `/25` subnets.

35.22.2 Network Overview

The example network depicted here is a data center environment consisting of two HA nodes with four interfaces each: WAN, WEB (LAN), DBDMZ, and Sync. This network contains a number of web and database servers. It is not based on any real network, but there are countless similar production deployments.

WAN

The WAN side connects to the upstream network, which may be an ISP, a data center, upstream router, or other means of external connectivity.

WEB Network

The WEB segment in this network uses the "LAN" interface but renamed. It contains web servers, so it has been named WEB, but it could be called DMZ, SERVERS, or anything desired.

DBDMZ Network

This segment is an OPT interface and contains database servers. It is common to segregate web and database servers into separate networks in hosting environments. The database servers typically do not require direct access from the Internet, and hence are better protected against remote compromise than web servers.

Sync Network

The HA nodes use the Sync network in this diagram to replicate configuration changes via XMLRPC and for state synchronization to replicate state table changes between the two firewalls. As described in other HA documentation sections, the best practice is to use a dedicated interface for this purpose.

Network Layout

Figure *Diagram of HA with Routed IPv4 Subnets* illustrates this network layout, including all routable IP addresses, the WEB network, and the Database DMZ.

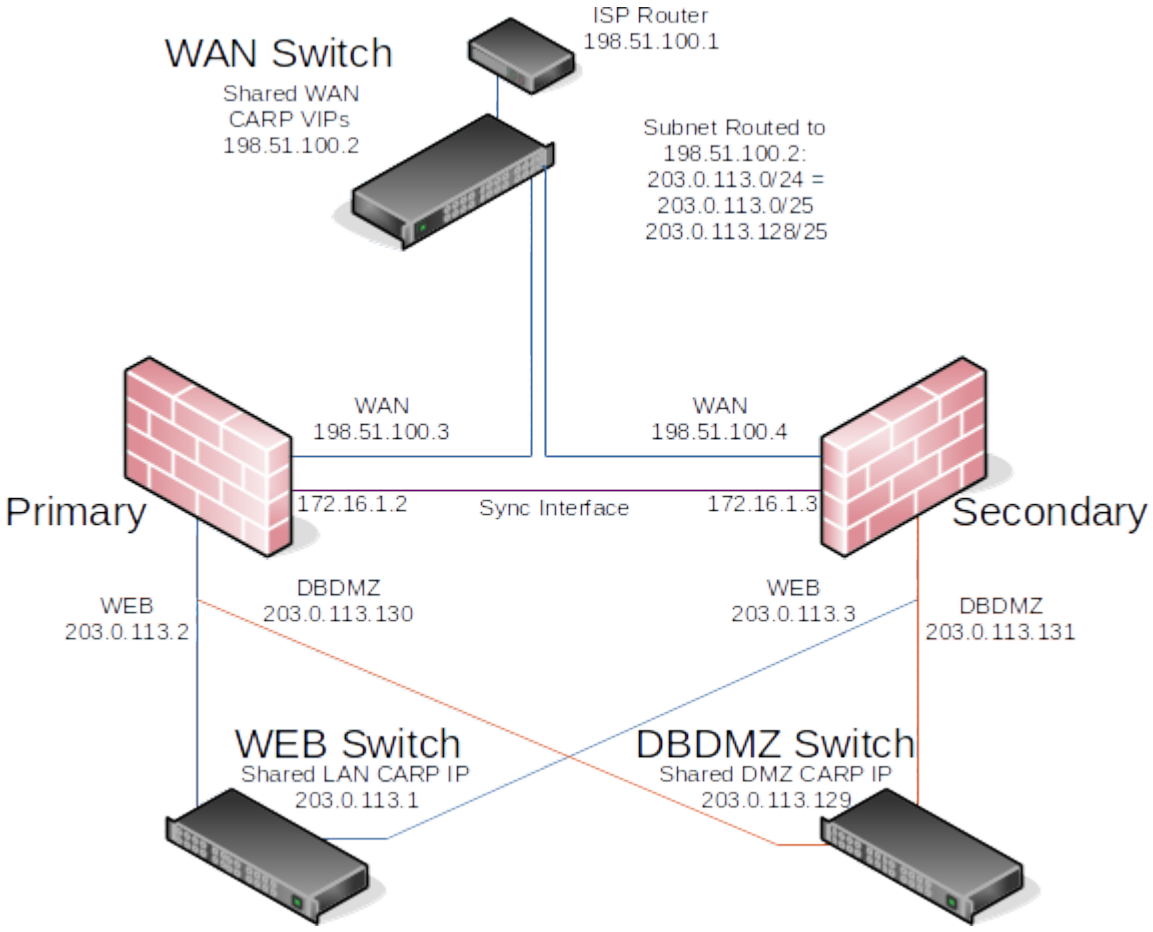


Fig. 11: Diagram of HA with Routed IPv4 Subnets

Note: Segments containing database servers typically do not need to be publicly accessible, and hence would more commonly use private IP subnets, but the example illustrated here can be used regardless of the function of the two internal subnets.

Outbound NAT

The **Outbound NAT Mode** in **Firewall > NAT** on the **Outbound NAT** tab must be set to **Manual** to properly accommodate this scenario. When the mode is set to **Manual**, saving the settings creates a set of outbound NAT rules equivalent to the rules which were utilized in **Automatic** or **Hybrid** modes.

In the list of rules created by switching to **Manual Outbound NAT**, delete any rules which translate sources using the routed subnets (e.g. WEB Networks, DBDMZ Networks). This effectively disables NAT for the routed subnets.

Warning: **Do not delete** outbound NAT rules that translate from a source of Localhost (`127.0.0.1`) using the WAN interface address(es) as these are necessary to ensure traffic from the firewall itself can egress properly.

See also:

- *Outbound NAT*

Firewall Rules

Firewall rules in this scenario are much simpler as there is no need to consider NAT translation. Allow traffic directly to the necessary destinations as needed.

Warning: As there is no NAT involved, carefully crafting WAN rules becomes more critical as being overly permissive with WAN rules not only can allow remote attackers access to the firewall itself, but all local networks.

Ensure the rules properly segment the DBDMZ from the WEB network; Block or reject any traffic between the segments except what is absolutely required for necessary functionality. For example, only allow connections from web servers to the database port(s) and nothing else.

Note: Best practices and requirements vary wildly in this regard depending on the operating systems and other software involved. Consult documentation for the servers to determine requirements.

See also:

- *Firewall*

35.23 IPsec Remote Access VPN Example Using IKEv1 with Pre-Shared Keys

This article describes how to set up mobile IPsec in pfSense® software with a Pre-Shared Key.

Note: The current best practice is to use IKEv2 with EAP authentication for IPsec Remote Access on modern clients. See *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2* for details.

Warning: There are very few remaining clients compatible with this type of configuration because it is considered weak compared to other options such as IKEv2 with EAP.

35.23.1 IPsec Server Setup

This is the setup for the pfSense® software side of the connection.

Mobile Clients

- Navigate to **VPN > IPsec**, **Mobile Clients** tab
- Set the options as follows:

  Enable IPsec Mobile Client Support
  Checked

  User Authentication
  Local Database

  Provide a virtual IP address to clients
  Checked

  Enter an unused subnet in the box (e.g. `10.11.200.0`), pick a subnet mask (e.g. `24`)
- Set other options if desired
- Click **Save**
- Click **Apply Changes**
- Click  **Create Phase 1** at the top of the screen if it appears

Phase 1 settings

- Navigate to **VPN > IPsec**
- Locate the Mobile Phase 1 in the list
- Click  to edit the Mobile Phase 1
- Enter the following settings:

  Description
  `Mobile IPsec PSK`

  Key Exchange Version
  Auto to allow both IKEv1 and IKEv2 connections. If all clients are compatible with IKEv2, use that instead.

  Note: Some clients, such as the native Android client, require options which only work with IKEv2.

  Authentication method
  Mutual PSK

  Negotiation mode
  Aggressive or *Main* depending on client requirements.

  My identifier
  My IP address

Encryption Algorithm
Create several entries which match values for common clients. Add them in order of preference with the most secure options listed first. For example:

- **Algorithm** *AES128-GCM*, **Hash** *SHA256*, **DH Group** *16* (if using IKEv2 only, required for Android)
- **Algorithm** *AES 256*, **Hash** *SHA512*, **DH Group** *14*
- **Algorithm** *AES 256*, **Hash** *SHA256*, **DH Group** *14*
- **Algorithm** *AES 256*, **Hash** *SHA1*, **DH Group** *14*
- **Algorithm** *AES 128*, **Hash** *SHA256*, **DH Group** *2*
- **Algorithm** *AES 128*, **Hash** *SHA1*, **DH Group** *2*

Life Time
`86400`

NAT Traversal
Force

- Click **Save**

Phase 2 settings

- Click **Show Phase 2 Entries** inside the Mobile phase 1 to expand its phase 2 list
- Click **Add P2** to create a new phase 2 entry
- Enter the following settings:

Description
`Mobile IPsec`

Mode
Tunnel IPv4

Local Network
The network on the firewall site which the clients must reach, e.g. *LAN Subnet*, or *Network* `0.0.0.0/0` to send all traffic over the VPN.

Protocol
ESP

Encryption Algorithms
AES 128

Hash Algorithms
SHA256

PFS key group
off

Lifetime
`28800`

- Add additional phase 2 entries for local networks if necessary
- Click **Save**

- Click **Apply Changes**

User Settings

Create pre-shared keys to identify users for the VPN

- Navigate to **VPN > IPsec**, **Pre-shared keys** tab.
- Click  **Add** to create a new entry
- Enter the settings as follows:

  Identifier
  Any identifier may be used so long as it is unique to the person using the account.

  Tip: E-mail addresses are commonly used as they are more unique than first or last names.

  Secret Type
  PSK

  Pre-Shared Key
  Generate a long/random Pre-Shared Key. The longer and more complex the key, the more secure it is.

  Note: Some clients, such as Linux network manager, require a minimum key length of 20 characters.

- Click **Save**
- Click **Apply Changes**

Firewall Rules

Add firewall rules to pass traffic from clients

- Navigate to **Firewall > Rules**, **IPsec** tab
- Add rules that match traffic to allow from mobile clients or add a rule to pass any protocol/any source/any destination to allow everything.

See also:

IPsec and firewall rules

The firewall configuration is complete.

35.23.2 Client Configuration

Android

Android 11.x and later contain a client compatible with a pre-shared key configuration provided that it uses IKEv2 only. See the inline notes above for additional requirements.

Note: The settings below are from pure Android 11.x. These exact settings may not present on all Android devices, depending on the Android version and changes made by the OEM.

See *Remote Access Mobile VPN Client Compatibility* for additional details.

- Swipe down twice from the top of the screen
- Tap the **Settings** cog
- Tap **Networks & Internet**, **Advanced**, **VPN**
- Tap +
- Enter the connection settings as follows:

  Name
  `pfSense Mobile VPN` or another suitable description

  Type
  IKEv2/IPsec PSK

  Server Address
  The address of the server.

  IPsec Identifier
  The identifier on the pre-shared key for this user (e.g. a username or e-mail address)

  Pre-Shared Key
  The PSK value associated with the identifier for this user.

- Tap **Save**

35.24 IPsec Remote Access VPN Example Using IKEv1 with Xauth

This document covers IPsec using Xauth and a mutual Pre-Shared Key.

Note: The current best practice is to use IKEv2 for IPsec Remote Access on modern clients. See *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2* for details.

This setup has been tested and working on various *Android* and *iOS* devices. Other clients may work as well.

35.24.1 IPsec Server Setup

This is the setup for the pfSense® software side of the connection

Mobile Clients

- Navigate to **VPN > IPsec**, **Mobile Clients** tab
- Set the options as follows:

  Enable IPsec Mobile Client Support
  Checked

  User Authentication
  Local Database

  Provide a virtual IP address to clients
  Checked

  Enter an unused subnet in the box (e.g. `10.11.200.0`), pick a subnet mask (e.g. `24`)
- Set other options if desired
- Click **Save**
- Click **Apply Changes**
- Click  **Create Phase 1** at the top of the screen if it appears

Phase 1 settings

- Navigate to **VPN > IPsec**
- Locate the Mobile Phase 1 in the list
- Click  to edit the Mobile Phase 1
- Enter the following settings:

  Description
  `Mobile IPsec PSK + Xauth`

  Key Exchange Version
  IKEv1

  Authentication method
  Mutual PSK + Xauth

  Negotiation mode
  Aggressive or *Main* depending on client requirements.

  My identifier
  My IP address

  Peer identifier
  User fully qualified domain name / E-mail, `vpnusers@example.com`

  Pre-Shared Key
  A long/random pre-shared key suitable for giving to users.

Encryption Algorithm
Create several entries which match values for common clients. Add them in order of preference with the most secure options listed first. For example:

- **Algorithm** *AES 256*, **Hash** *SHA512*, **DH Group** *14*
- **Algorithm** *AES 256*, **Hash** *SHA256*, **DH Group** *14*
- **Algorithm** *AES 256*, **Hash** *SHA1*, **DH Group** *14*
- **Algorithm** *AES 128*, **Hash** *SHA1*, **DH Group** *2*

Life Time
`86400`

NAT Traversal
Force

- Click **Save**

Phase 2 settings

- Click  **Show Phase 2 Entries** inside the Mobile phase 1 to expand its phase 2 list
- Click  **Add P2** to create a new phase 2 entry
- Enter the following settings:

Description
`Mobile IPsec`

Mode
Tunnel IPv4

Local Network
The network on the firewall site which the clients must reach, e.g. *LAN Subnet*, or *Network* `0.0.0.0/0` to send all traffic over the VPN.

Protocol
ESP

Encryption Algorithms
AES 128

Hash Algorithms
SHA1

PFS key group
off

Lifetime
`28800`

- Add additional phase 2 entries for local networks if necessary
- Click **Save**
- Click **Apply Changes**

User Settings

- Navigate to **System > User Manager**
- *Add a user*
- Edit the user and grant them the *User - VPN - IPsec xauth Dialin* privilege or add them to a group with this privilege.

Note: Xauth uses both this per-user password and the value of the pre-shared key for different types of authentication. The pre-shared key is used to authenticate the tunnel itself and the per-user password ensures that a particular user is authorized to access the tunnel.

Firewall Rules

Add firewall rules to pass traffic from clients

- Navigate to **Firewall > Rules**, **IPsec** tab
- Add rules that match traffic to allow from mobile clients or add a rule to pass any protocol/any source/any destination to allow everything.

See also:

IPsec and firewall rules

35.24.2 Device Setup (Android)

Note: The settings below are from pure Android 11.x. These exact settings may not present on all Android devices, depending on the Android version and changes made by the OEM.

See *Remote Access Mobile VPN Client Compatibility* for additional details.

- Swipe down twice from the top of the screen
- Tap the **Settings** cog
- Tap **Networks & Internet**, **Advanced**, **VPN**
- Tap +
- Enter the connection settings as follows:

  Name
  `pfSense Mobile VPN` or another suitable description

  Type
  IPsec Xauth PSK

  Server Address
  The address of the server.

  IPsec Identifier
  If the mobile IPsec phase 1 is set for *Aggressive* fill in the identifier set in phase 1 (e.g. `vpnusers@example.com`).

  If the mobile IPsec phase 1 is set for *Main*, leave this at the default empty value of *(not used)*.

Pre-Shared Key
The value of the pre-shared key from the mobile phase 1 entry.

Username
The username for this xauth user

Password
The password for this xauth user

- Tap **Save**

35.24.3 Device Setup (iOS)

- Tap **Settings** > **VPN** or **Settings** > **General** > **VPN**
- Tap **Add VPN Configuration**
- Set **Type** to **IPsec**
- Enter the settings as follows:

Description
`pfSense Mobile VPN` or another suitable description

Server
The address of the server.

Account
The username for this xauth user

Password
The password for this xauth user (or leave blank to be prompted every time)

Group Name
The identifier set in phase 1 (e.g. `vpnusers@example.com`).

Secret
The value of the pre-shared key from the mobile phase 1 entry.

35.24.4 Troubleshooting

By default, iOS will tunnel all traffic over the VPN including traffic going to the Internet. If Internet sites are inaccessible once connected, a DNS server may need to be pushed to the client for it to use. This could be the LAN IP address of the firewall if the DNS resolver is enabled or a public DNS server such as `8.8.8.8` and/or `8.8.4.4`.

The reason for the above is that the cellular provider is likely giving mobile devices DNS servers that are only accessible from their network. Once connected to the VPN the DNS servers are now being accessed via the VPN instead of the provider network, thus the queries are likely to be dropped. Supplying a local or public DNS server will work around this problem.

35.25 Configuring IPsec IKEv2 Remote Access VPN Clients

Most operating systems include native clients compatible with IPsec IKEv2 VPN connections, and others typically have an app or add-on package which adds the capability.

This section covers IPsec IKEv2 client configuration for several popular operating systems.

Tip: The *ipsec-profile-wizard* package on pfSense Plus software generates a set of files which can automatically import VPN settings into Apple macOS and iOS (**VPN > IPsec Export: Apple Profile**) as well as Windows clients (**VPN > IPsec Export: Windows**).

This feature allows much greater flexibility in settings as it will configure clients to match what is set on the server specifically rather than making the server accommodate the default settings on various operating systems.

This package is exclusive to pfSense Plus software and is not available on the community edition.

If the package is not already installed, add it using the *Package Manager*.

See also:

Hangouts Archive to watch the September and October 2015 Hangouts on Remote Access VPNs which covers client installations for most operating systems.

35.25.1 Configuring IPsec IKEv2 Remote Access VPN Clients on Windows

---

Tip: The *ipsec-profile-wizard* package on pfSense Plus software generates a set of files which can automatically import VPN settings into Apple macOS and iOS (**VPN > IPsec Export: Apple Profile**) as well as Windows clients (**VPN > IPsec Export: Windows**).

This feature allows much greater flexibility in settings as it will configure clients to match what is set on the server specifically rather than making the server accommodate the default settings on various operating systems.

This package is exclusive to pfSense Plus software and is not available on the community edition.

If the package is not already installed, add it using the *Package Manager*.

Windows 8 and newer can easily connect to IKEv2 VPNs. The procedure in this section was performed on Windows 10 20H2 but earlier versions are similar.

Import the CA to the Client (All EAP types)

This step is necessary for all EAP types (EAP-MSCHAPv2, EAP-RADIUS, EAP-TLS).

- Export the CA Certificate from the pfSense® software GUI and download or copy it to the client PC:
  - Navigate to **System > Certificates**, **Certificate Authorities** tab on the firewall
  - Click  by the CA to download only the certificate
- Locate the downloaded file on the client PC (e.g. `VPNCA.crt`) as seen in Figure *Downloaded CA Certificate*
- Double-click the CA file
- Click **Install Certificate…** as shown in *Certificate Properties*

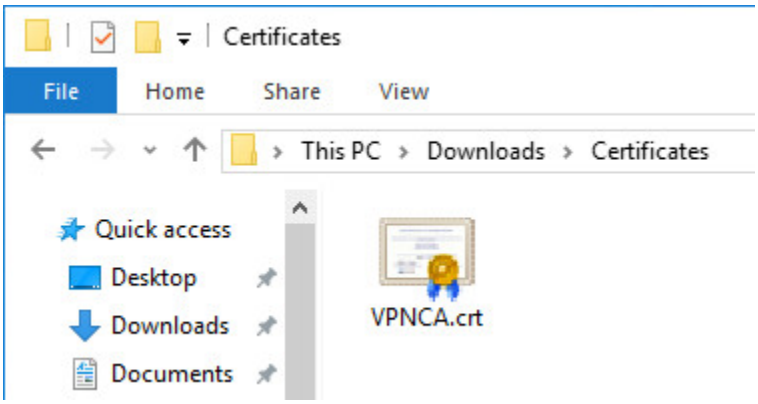


Fig. 12: Downloaded CA Certificate

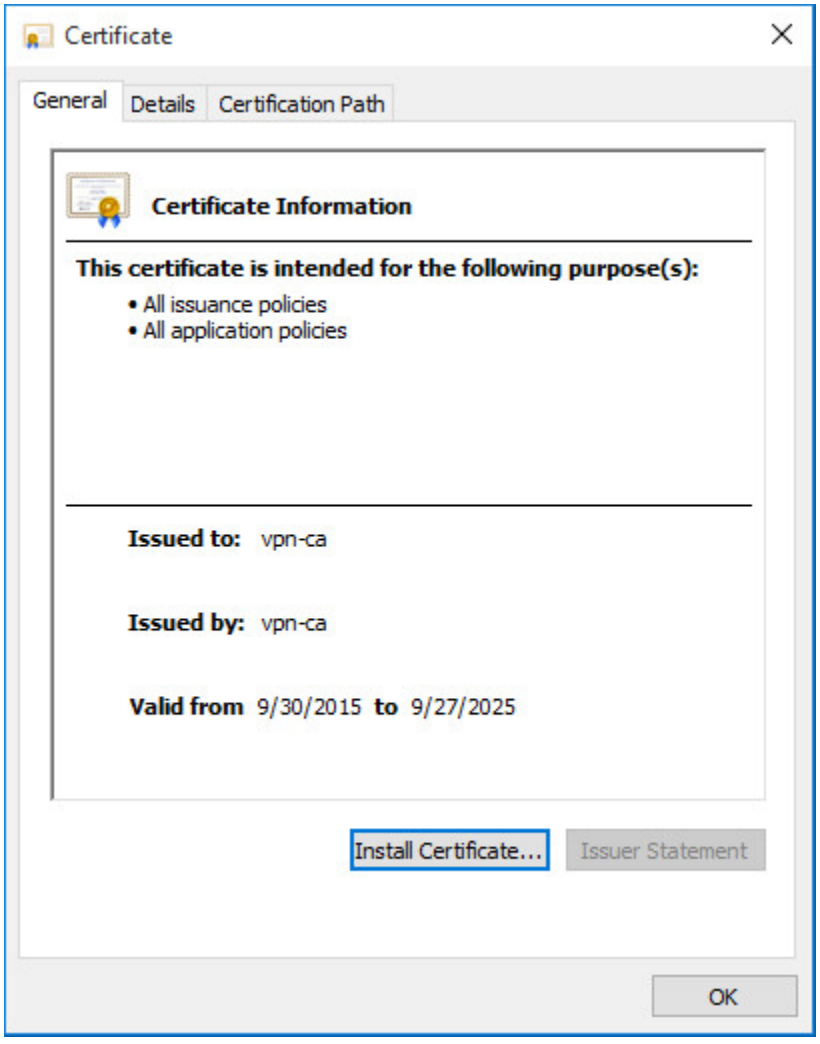


Fig. 13: Certificate Properties

- Select **Local Machine** as shown in *Certificate Import Wizard - Store Location*

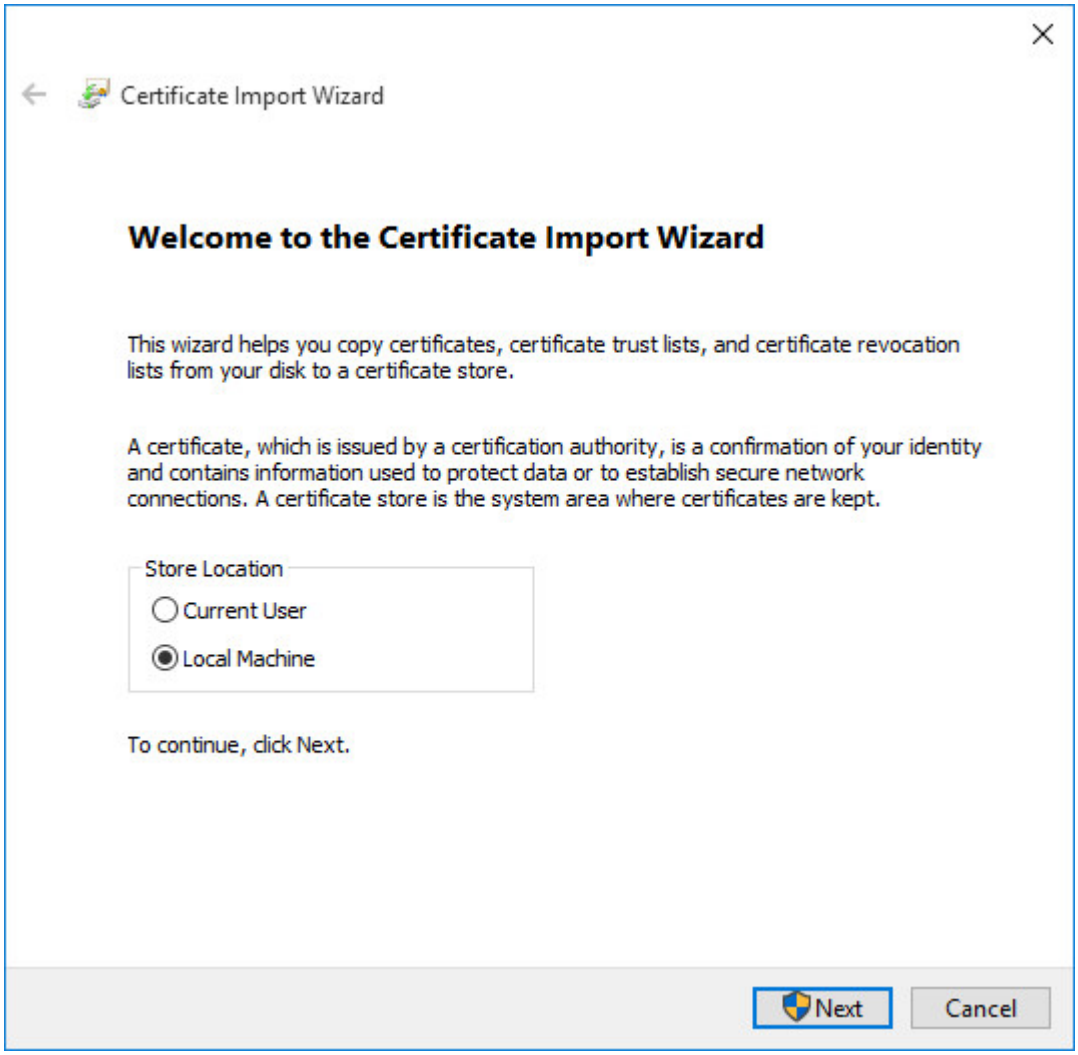


Fig. 14: Certificate Import Wizard - Store Location

- Click **Next**
- Click **Yes** at the UAC prompt if it appears
- Select **Place all Certificates in the following store** as shown in Figure *Certificate Import Wizard - Browse for the Store*
- Click **Browse**
- Click **Trusted Root Certification Authorities** as shown in Figure *Select Certificate Store*
- Click **Next**
- Review the details, they should match those in Figure *Completing the Certificate Import Wizard*
- Click **Finish**
- Click **OK** (2x)

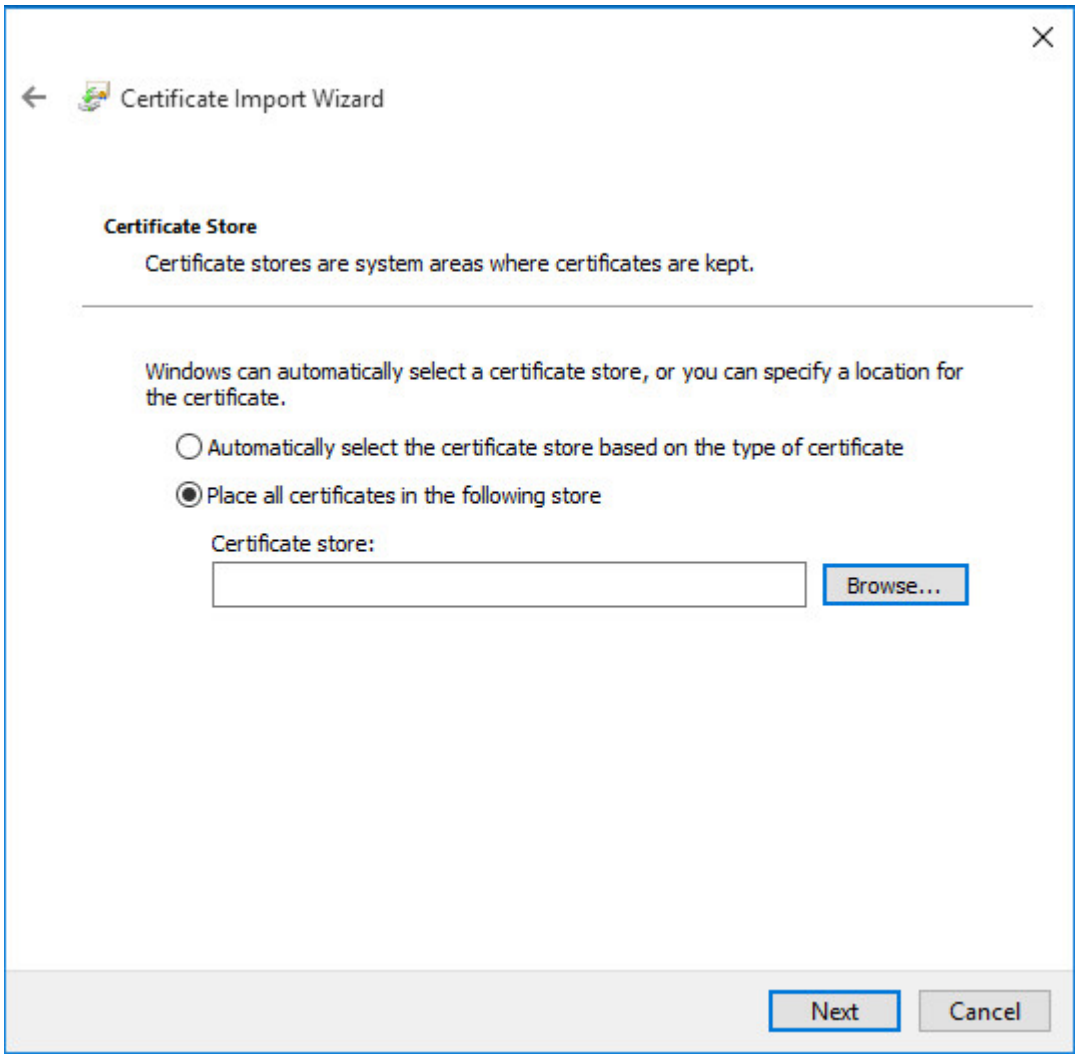


Fig. 15: Certificate Import Wizard - Browse for the Store

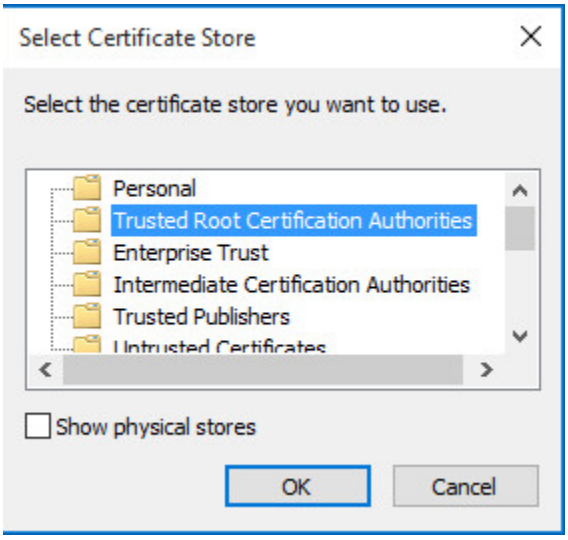


Fig. 16: Select Certificate Store

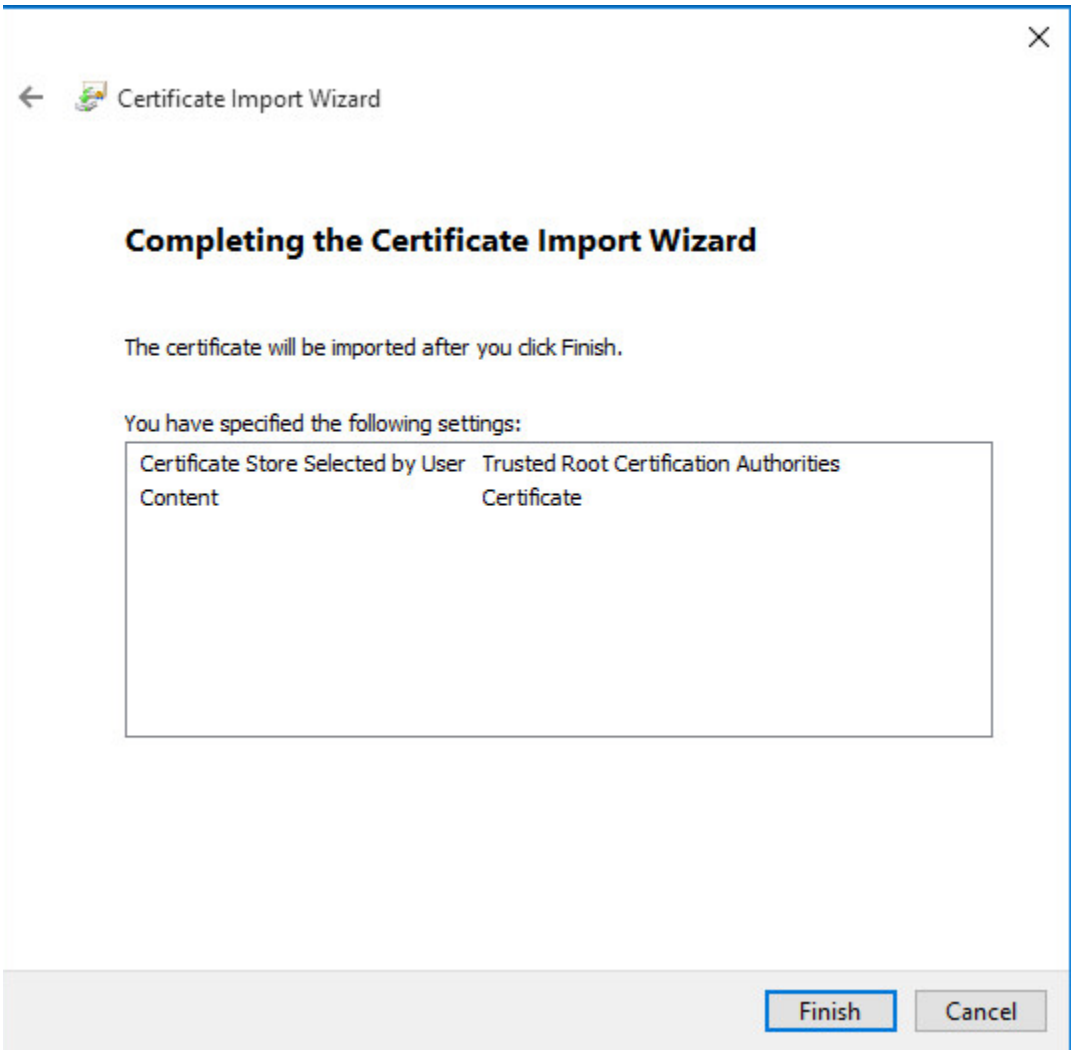


Fig. 17: Completing the Certificate Import Wizard

Import the CA and Client Certificate to the Client (EAP-TLS Only)

This process is only required for EAP-TLS which uses per-user client certificates. For EAP-MSCHAPv2 or EAP-RADIUS, skip to the next section.

- Export client certificate from the firewall and download it to the client PC
  - Navigate to **System > Certificates**, **Certificates** tab
  -  to edit the user certificate
  - Enter an **Export Password** known to the end user which will encrypt the sensitive contents of the archive file
  - Click  **Export PKCS#12** to download a `.p12` file containing the client certificate and key
- Locate the downloaded file on the client PC (e.g. `client1.p12`)
- Double-click client certificate `.p12` file
- Select *Current User*
- Click **Next**
- Click **Yes** at the UAC prompt if it appears
- Confirm the proper file is selected
- Click **Next**
- Enter the same **Password** used when exporting the `.p12` file
- Click **Next** (2x)
- Click **Finish**
- Click **Yes** to confirm adding the certificate data
- Click **OK**

Set up the VPN Connection

Once the certificate has been properly imported it is time to create the client VPN connection. The exact steps will vary depending on the version of Windows being used by the client, but will be close to the following procedure:

- Open **Network & Internet Settings** on the client PC
- Click **VPN** on the left side
- Click **+ Add a VPN connection**
- Set the fields as follows:

  Example values are shown in Figure *Windows IKEv2 VPN Connection Setup Screen*:

  VPN Provider
  Windows (built-in)

  Connection Name
  `ExampleCo Mobile VPN`

Server Name or Address
`vpn.example.com`

> **Warning:** This value **must** match the contents of the server certificate!

VPN type
IKEv2

Type of sign-in info
Username and password for EAP-MSCHAPv2 or EAP-RADIUS

Certificate for EAP-TLS

Username, Password
Fill in values for this client when using EAP-MSCHAPv2 or EAP-RADIUS. Leave blank to be prompted by Windows.

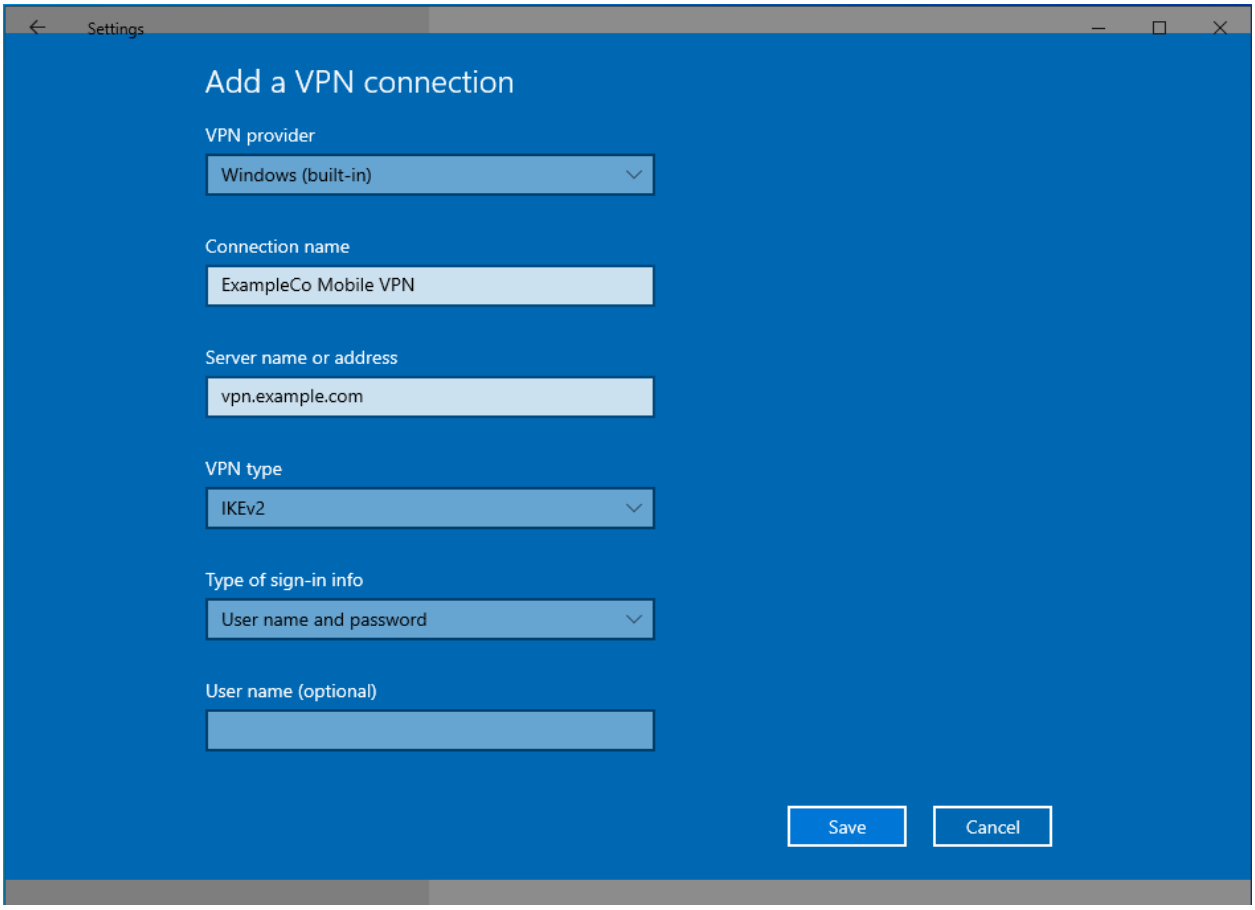


Fig. 18: Windows IKEv2 VPN Connection Setup Screen

- Click **Save**

The connection is now ready to use.

Note: When making the first connection Windows may prompt to approve the server certificate. Check the certificate

and then choose to proceed when prompted.

Disable EKU Check

Windows expects IKEv2 server certificates to contain the IKE intermediate extended key usage attribute (`1.3.6.1.5.5.8.2.2`), among others. Creating a CA and a server certificate in the Certificate Manager will add the correct set of attributes for this usage (*Certificate Settings*).

If the server certificate is created with the wrong settings, or the certificate is generated elsewhere (e.g. via ACME), the certificate may lack these attributes and clients will fail to connect.

To accommodate such certificates, the Extended Key Usage check can be disabled on Windows.

Warning: Disabling this check also disables validation of the certificate common name and SAN fields, so it is potentially dangerous. Any certificate from the same CA could be used for the server when this is disabled, so proceed with caution.

To disable the extended key usage checks:

- Open up **Registry Editor** on the Windows client
- Navigate to the following location in the client registry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\RasMan\Parameters\
```

- Add a new **DWORD** entry with the following attributes:

  Name
  `DisableIKENameEkuCheck`

  Value
  1

- Reboot the client PC to ensure the new setting is activated

Advanced Windows IPsec settings

With Windows 10 PowerShell Cmdlets it is possible to change various advanced settings. The available commands are explained on the Microsoft PowerShell VpnClient module reference.

Routes

Enable split tunneling so that the client does not send all of its traffic across the VPN:

```
PS> Set-VPNconnection -name "ExampleCo Mobile VPN" -SplitTunneling $true
```

Add a VPN connection route to send a specific subnet through the VPN, use:

```
PS> Add-VpnConnectionRoute -ConnectionName "ExampleCo Mobile VPN" -DestinationPrefix 10.
↪4.0.0/24
```

Replace `ExampleCo Mobile VPN` with the actual connection name, and replace `10.4.0.0/24` with the desired destination network. Repeat the add command for each network to route over the VPN.

See also:

For more information, see PowerShell VpnClient module reference

35.25.2 Configuring IPsec IKEv2 Remote Access VPN Clients on Android

---

Note: Android considers using a VPN an action that must be secure. When activating any VPN option the OS will force the user to add a lock method to the device if one is not already present. It does not matter which type of lock is chosen (PIN lock, Pattern lock, Password, etc.) but it will not allow a VPN to be configured until a secure lock has been added.

On Android devices with Face lock, that is not available as a secure lock type on its own.

There are two methods to configuring IKEv2 on Android: Natively on Android 11.x and later, or using the strongSwan app from the Play Store.

Native IKEv2 on Android

Android 11.x and later now include several IKEv2 client options compatible with mobile IPsec on pfSense® software. This example covers EAP-MSCHAPv2 which also works with EAP-RADIUS.

Note: The settings below are from pure Android 11.x. These exact settings may not present on all Android devices, depending on the Android version and changes made by the OEM.

Import the Server CA

To validate the server, the client needs to know the server certificate CA. The Android IKEv2 client will only validate against CA entities imported by the user.

Note: Though this validation is optional it is the best practice as otherwise the client cannot verify it is connecting to the correct server.

Warning: Installing a self-signed root CA into Android in this manner carries some danger as the CA could also be used to impersonate other servers. The danger is lower since this is controlled by firewall administrators, but the warnings presented when this is done may still confuse and worry end users.

If this is unacceptable, use the strongSwan application instead. It can validate against an existing root CA as well as validating a CA without installing it into the operating system trust store.

- Copy the CA certificate to the device
- Swipe down twice from the top of the screen
- Tap the **Settings** cog
- Tap **Security**, **Encryption & Credentials**

- Tap **Install a certificate**
- Tap **CA certificate**
- Read the warning text
- Tap **Install anyway**
- Locate and tap the CA certificate which was copied to the device

Set up the VPN Connection

- Swipe down twice from the top of the screen
- Tap the **Settings** cog
- Tap **Networks & Internet**, **Advanced**, **VPN**
- Tap +
- Enter the connection settings as follows:

  Name
  `ExampleCo Mobile VPN` or another suitable description

  Type
  IKEv2/IPsec MSCHAPv2

  Server Address
  The address of the server.

  Note: This must match a value in the server certificate. For example, a hostname or IP address in a certificate SAN entry.

  IPsec Identifier
  The identifier on the EAP pre-shared key for this user (e.g. a username or e-mail address)

  IPsec CA Certificate
  Select the imported CA (optional, but the best practice)

  Username
  The identifier for this user again.

  Password
  The EAP key value associated with the identifier for this user.

- Tap **Save**

Connecting and Disconnecting

To Connect:

- Swipe down twice from the top of the screen
- Tap the **Settings** cog
- Tap **Networks & Internet**, **Advanced**, **VPN**
- Tap the name of the VPN
- Tap **Connect**

Android displays a key icon in the notification bar near the network status icons and clocks while a VPN is connected.

To Disconnect:

- Swipe down twice from the top of the screen
- Tap the **Settings** cog
- Tap **Networks & Internet**, **Advanced**, **VPN**
- Tap the name of the VPN
- Tap **Disconnect**

strongSwan App on Android

Before starting, install the strongSwan app from the Play Store:

Set up the VPN Connection

- Copy the CA Certificate to the device
- Open the strongSwan app
- Import the CA:
  - Tap the settings icon (Three vertical dots in the upper right)
  - Tap **CA Certificates**
  - Tap the settings icon (Three vertical dots in the upper right)
  - Tap **Import Certificate**
  - Locate the CA Certificate copied earlier and tap it.
- Tap **Add VPN Profile**
- Enter the address of the firewall as the **Gateway** (e.g. `vpn.example.com`)
- Select *IKEv2 EAP (Username/Password)* for the **Type**
- Enter the **Username**
- Enter the **Password** to have it be remembered or leave it blank to prompt for the password on each connection.
- Check Select automatically under CA Certificate
- Enter a **Profile Name** (optional, if left blank, the gateway address will be used)
- Compare the settings to Figure *Android strongSwan Client Settings*

Connecting and Disconnecting

To Connect:

- Open the strongSwan app
- Tap the desired VPN
- Check **I trust this application** at the security prompt as shown in *Android strongSwan Client Settings*
- Tap OK

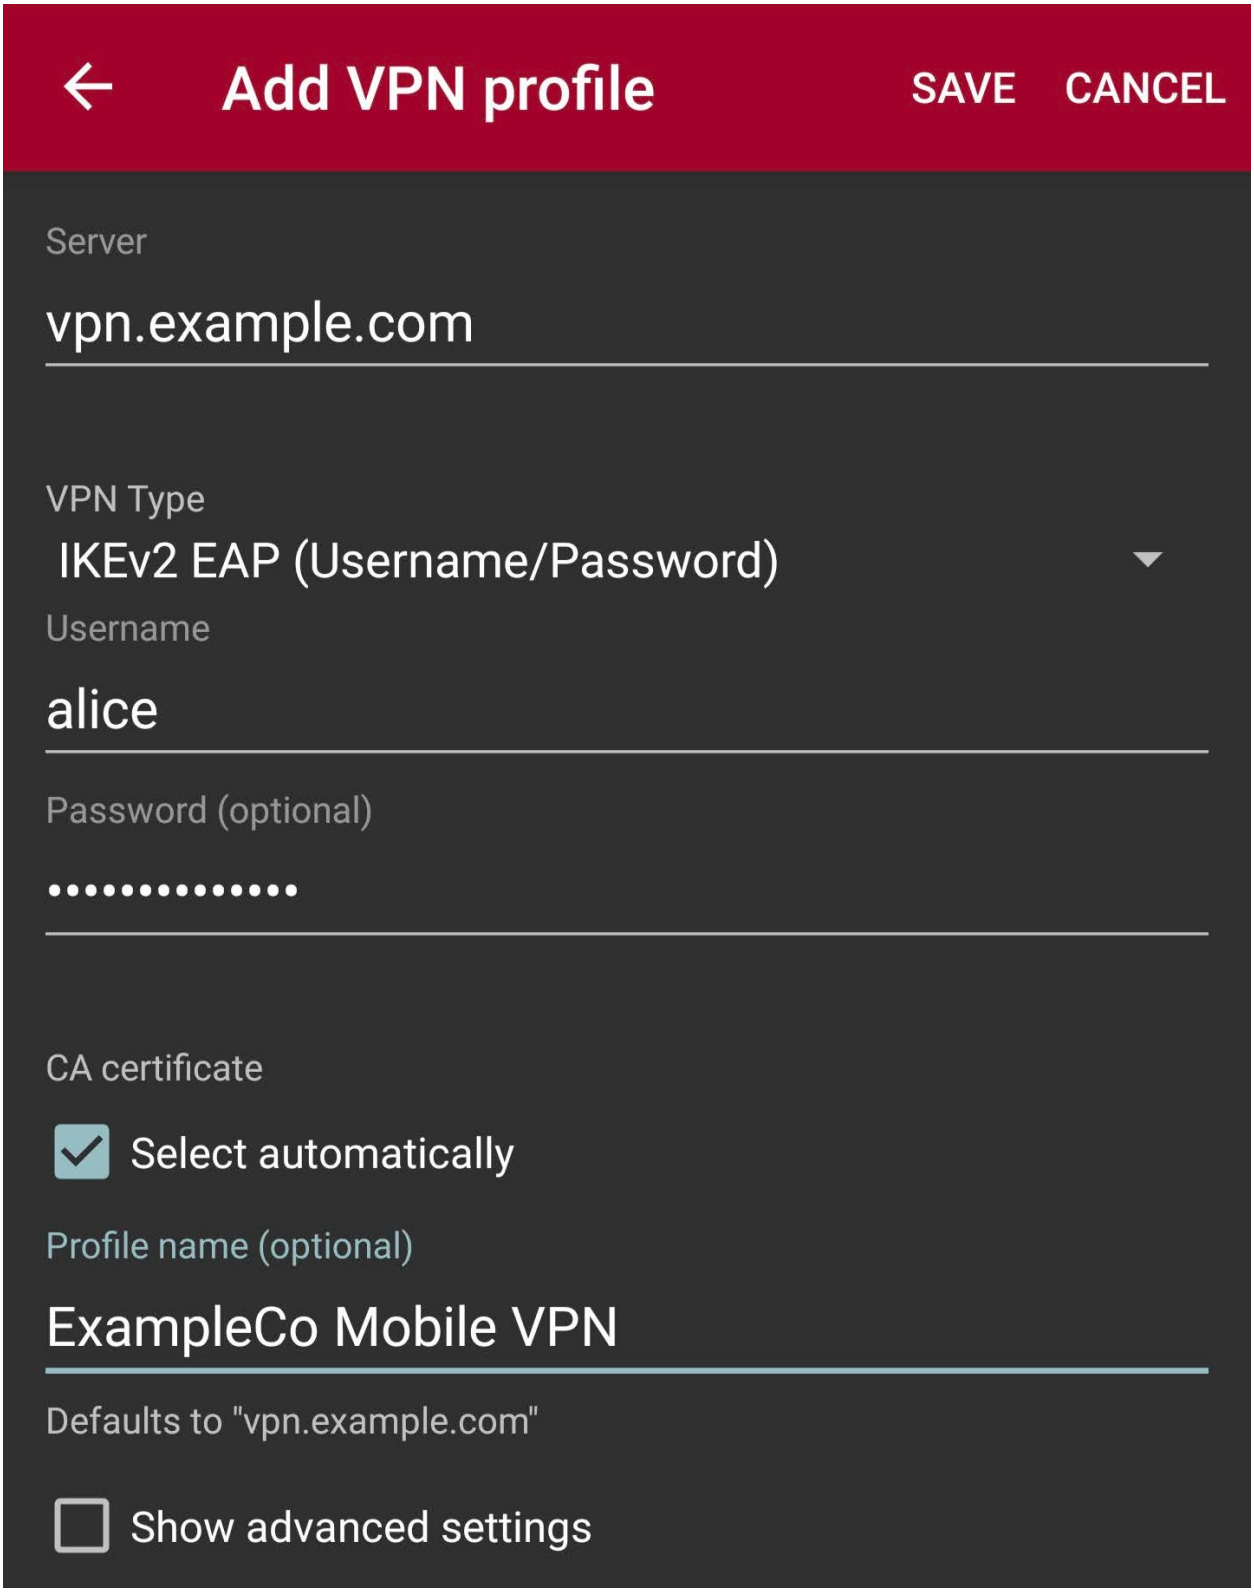


Fig. 19: Android strongSwan Client Settings

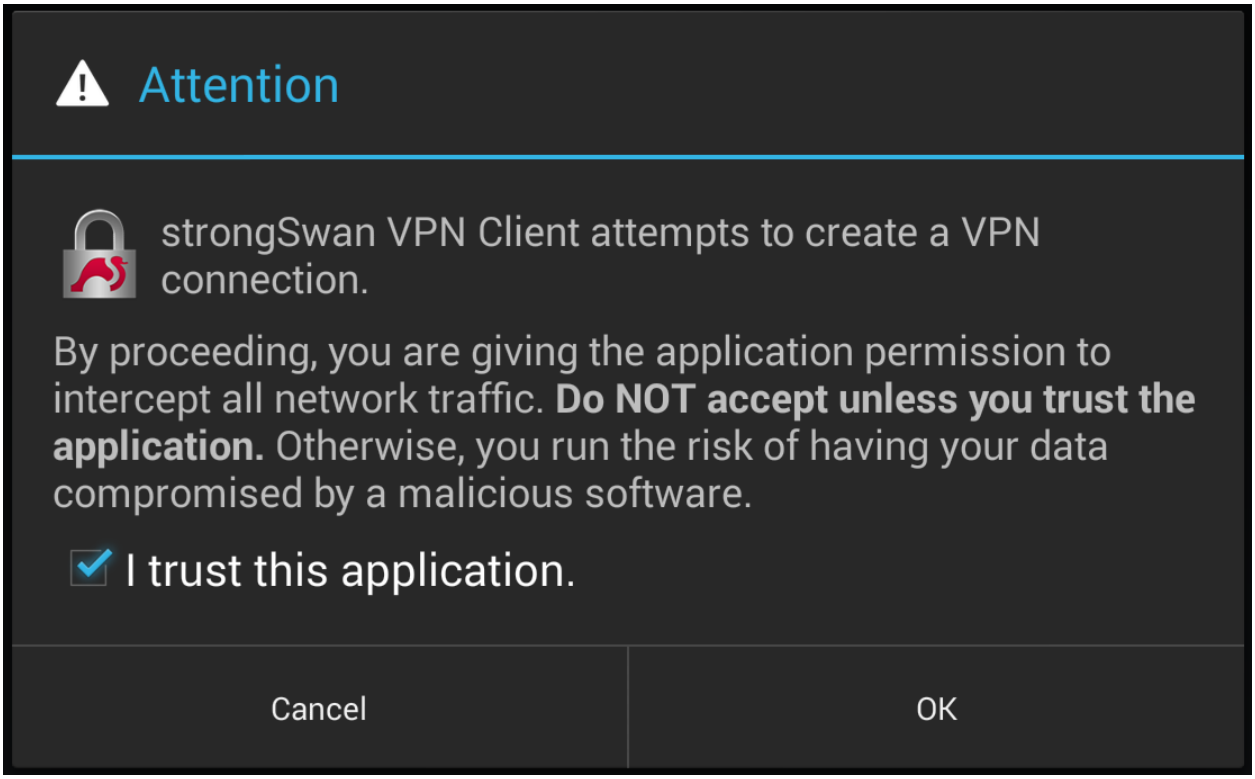


Fig. 20: Android strongSwan Client Settings

To Disconnect:

- Swipe down from the top notification bar
- Tap the strongSwan entry in the notification list
- Tap **Disconnect**

Alternately:

- Open the strongSwan app
- Tap **Disconnect** on the desired VPN

35.25.3 Configuring IPsec IKEv2 Remote Access VPN Clients on macOS

It is possible to configure an IKEv2 type VPN manually in the macOS GUI without needing a VPN Profile configuration file. Basic configuration for IKEv2 is integrated into the network management settings the same as other connections, but it is quite limited.

Tip: The *ipsec-profile-wizard* package on pfSense Plus software generates a set of files which can automatically import VPN settings into Apple macOS and iOS (**VPN > IPsec Export: Apple Profile**) as well as Windows clients (**VPN > IPsec Export: Windows**).

This feature allows much greater flexibility in settings as it will configure clients to match what is set on the server specifically rather than making the server accommodate the default settings on various operating systems.

This package is exclusive to pfSense Plus software and is not available on the community edition.

If the package is not already installed, add it using the *Package Manager*.

Warning: The best practice is to use a VPN profile, such as from the Apple Configurator or *IPsec Export Package*. A profile requires less configuration on the client and can use options which are more secure and offer faster performance than the client attempts by default. Additionally, without using a profile it may not be possible to create a mobile IPsec configuration which can natively allow connections from different client types.

Import the CA Certificate into macOS

The VPN Server CA Certificate must be imported before a client can connect.

- Copy the CA Certificate to the macOS system
- Double-click the CA Certificate File in Finder (Figure *macOS Certificate File in Finder*), which opens Keychain Access

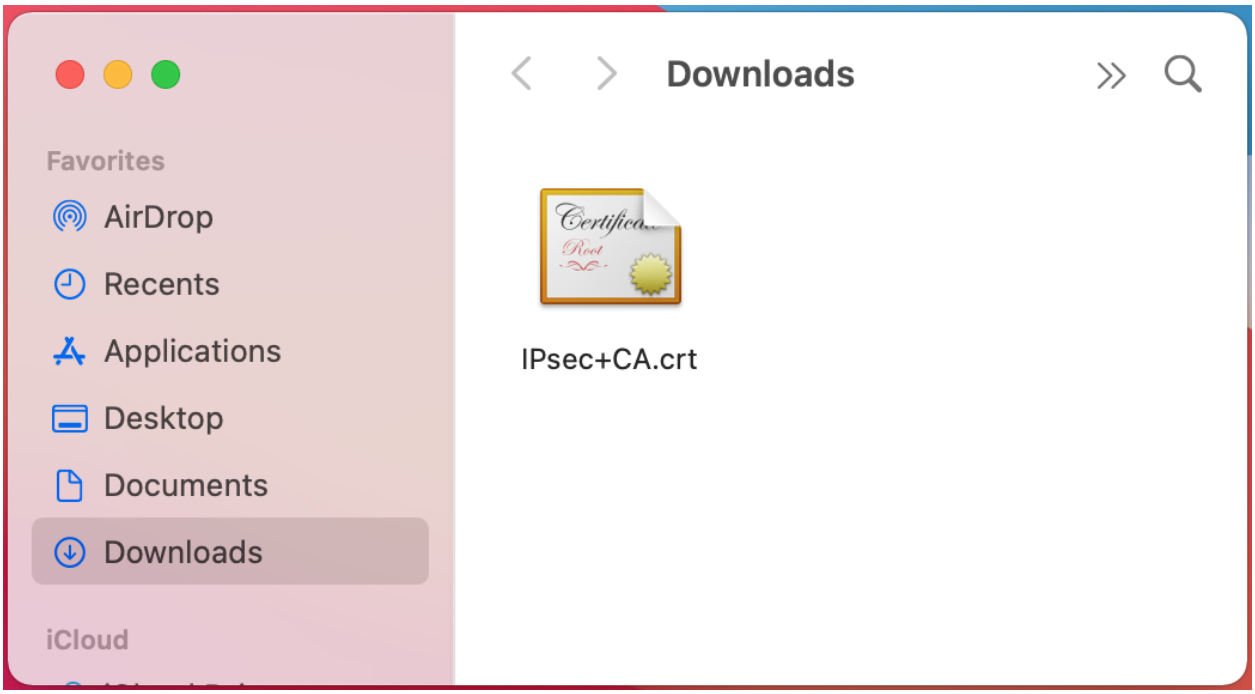


Fig. 21: macOS Certificate File in Finder

- Set **Keychain** to *System* on the **Add Certificates** dialog
- Click **Add**

Note: If this dialog does not appear and the CA is added into another keychain, locate it in Keychain Access and drag it into the **System** keychain.

- Locate the imported certificate under **System**, **Certificates** as shown in Figure *macOS Keychain Access System Certificate List*
- Click the Certificate

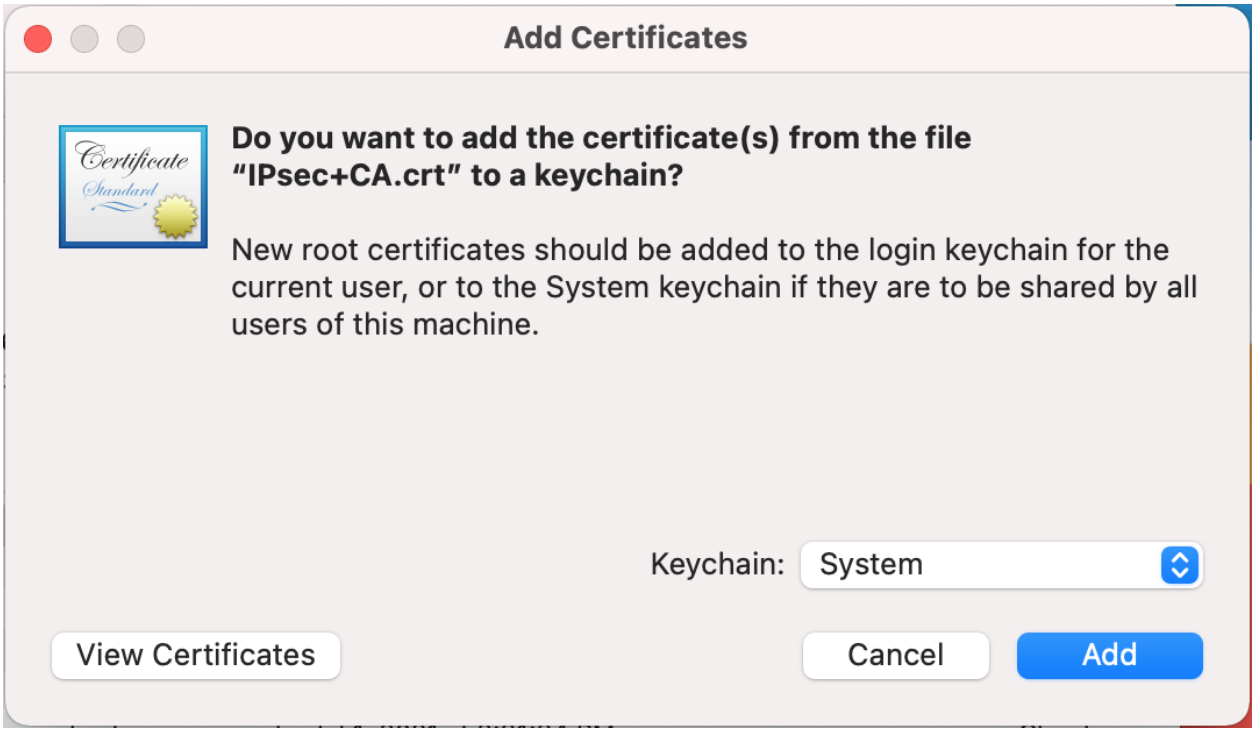


Fig. 22: macOS Keychain Access Add Certificate

- Click **File > Get Info**
- Expand **Trust**
- Set **When using this certificate** to *Always Trust* as shown in Figure *macOS Certificate Trust Settings*
- Click the red close button to close the certificate info window

  This triggers an authentication prompt to allow the change.
- Enter the login credentials and click **Update Settings**
- Quit Keychain Access

The certificate is now located in System Certificates and has been marked as trusted so it can be used for the VPN.

Set up the VPN Connection

- Open System Preferences
- Click **Network**
- Click the lock icon and enter credentials to make changes if the settings have not already been unlocked
- Click + to add a new VPN entry as shown in Figure *macOS Add Network Button*
- Select *VPN* for the **Interface**
- Select *IKEv2* for the **VPN Type** (default)
- Set **Service Name** to a description for the VPN (e.g. `ExampleCo Mobile VPN`) to complete the form, which will look similar to Figure *macOS Create VPN Prompt*
- Click **Create**

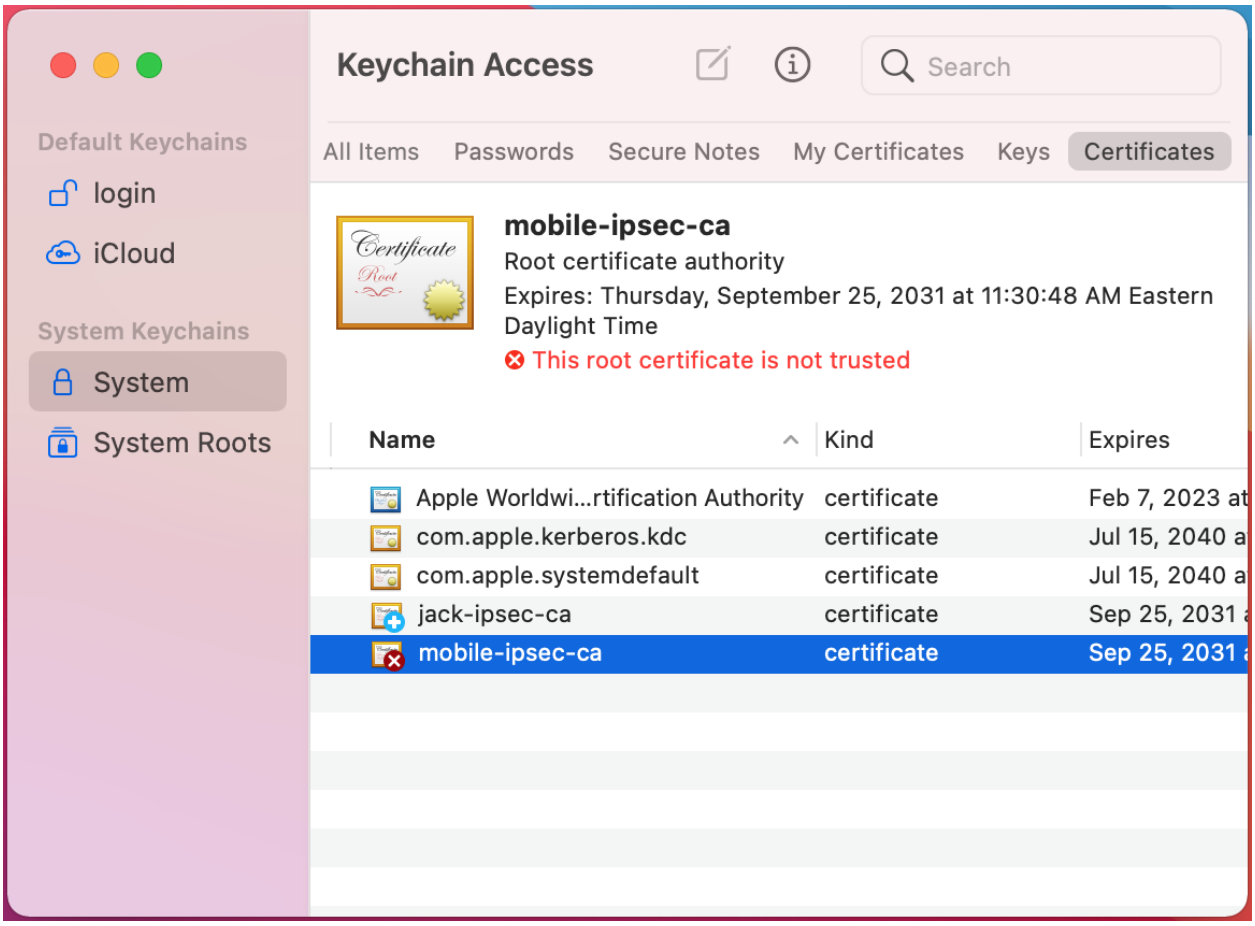


Fig. 23: macOS Keychain Access System Certificate List

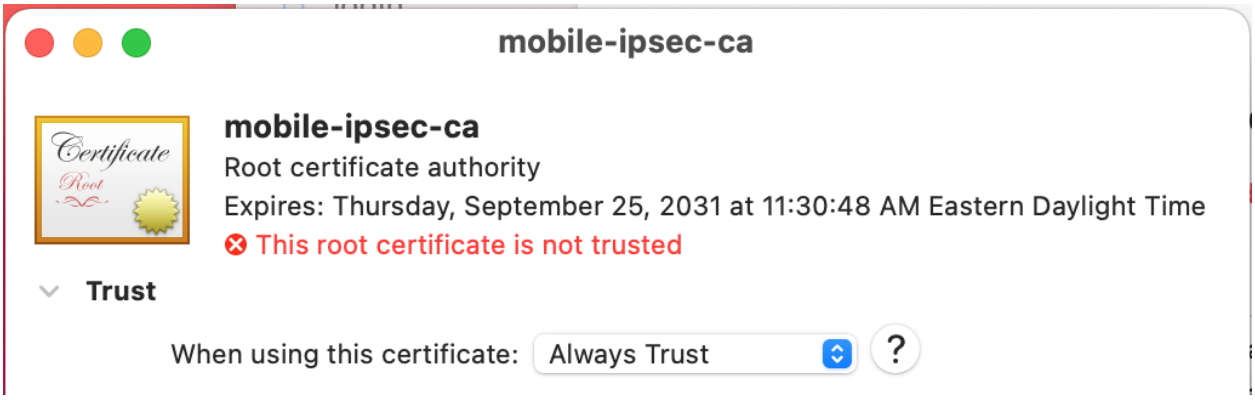


Fig. 24: macOS Certificate Trust Settings

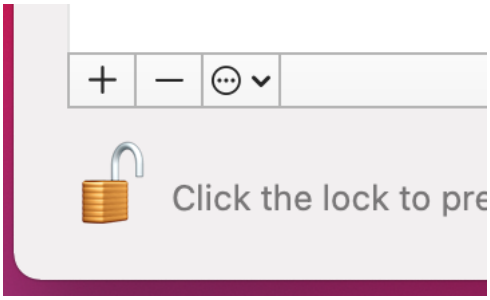


Fig. 25: macOS Add Network Button

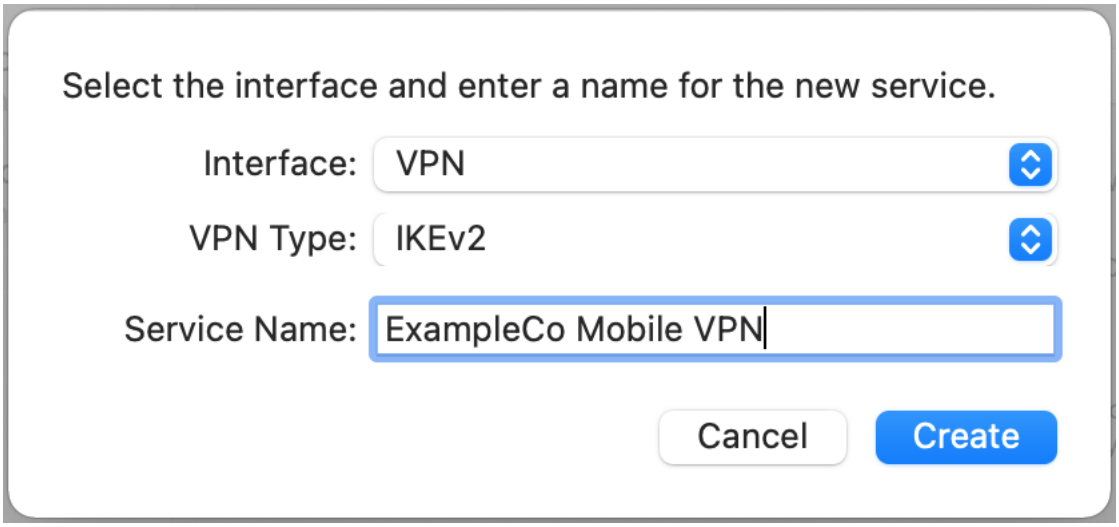


Fig. 26: macOS Create VPN Prompt

- Enter the hostname of the firewall in DNS as the **Server Address**
- Enter the hostname of the firewall again in **Remote ID**

Note: This must match the server certificate's Common Name and SAN entry.

- Leave **Local ID** blank, the settings will now look like Figure *macOS IKEv2 VPN Settings*

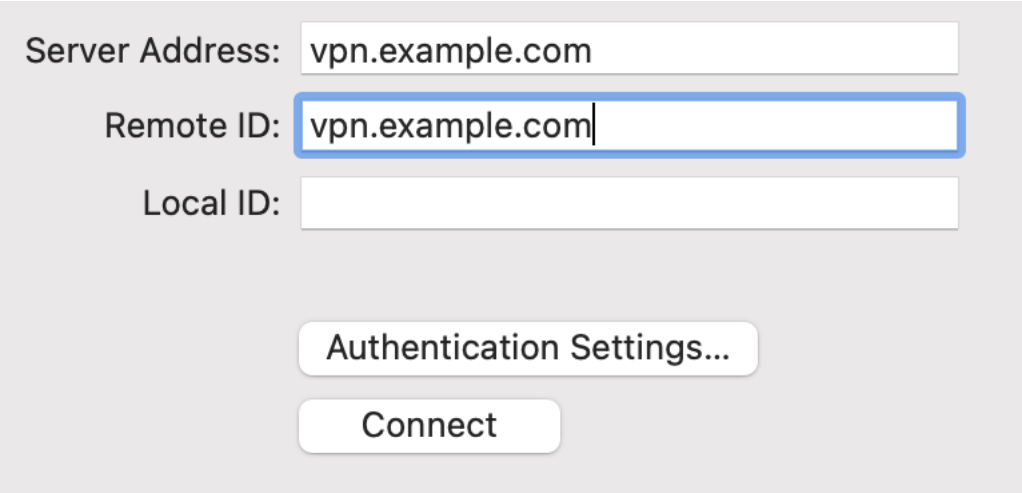


Fig. 27: macOS IKEv2 VPN Settings

- Click **Authentication Settings**
- Select **Username**
- Enter the **Username** and **Password** as shown in Figure *macOS IKEv2 VPN Authentication Settings*

Note: With EAP-MSCHAPv2 the **Username** is the **Identifier** configured for the user's entry on the **Pre-Shared Keys** tab under **VPN > IPsec**. With EAP-RADIUS this would be the username set on the RADIUS server.

- Check **Show VPN status in the menu bar** (if desired)
- Click **Apply**

Connecting and Disconnecting

Managing the connection can be done multiple ways. The first method is to click **Connect** or **Disconnect** on the VPN entry in Network settings. The second, easier method is to check **Show VPN Status in the menu bar** in the VPN settings and then manage the connection from that icon, as shown in Figure *macOS VPN Status Menu*.

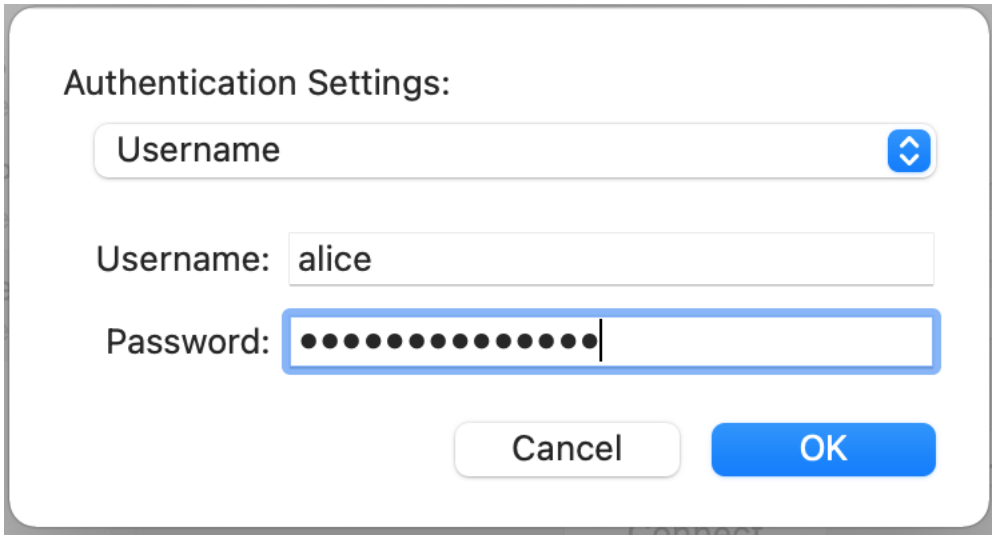


Fig. 28: macOS IKEv2 VPN Authentication Settings

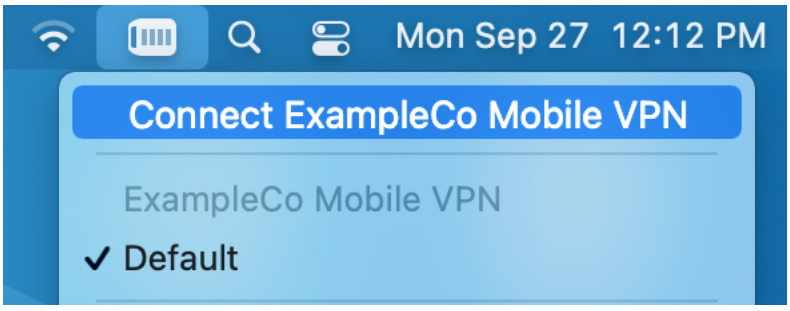


Fig. 29: macOS VPN Status Menu

35.25.4 Configuring IPsec IKEv2 Remote Access VPN Clients on iOS

iOS version 9 and later have built-in capability for configuring a basic IKEv2 connection without a VPN Profile.

Tip: The *ipsec-profile-wizard* package on pfSense Plus software generates a set of files which can automatically import VPN settings into Apple macOS and iOS (**VPN > IPsec Export: Apple Profile**) as well as Windows clients (**VPN > IPsec Export: Windows**).

This feature allows much greater flexibility in settings as it will configure clients to match what is set on the server specifically rather than making the server accommodate the default settings on various operating systems.

This package is exclusive to pfSense Plus software and is not available on the community edition.

If the package is not already installed, add it using the *Package Manager*.

Warning: The best practice is to use a VPN profile, such as from the Apple Configurator or *IPsec Export Package*. A profile requires less configuration on the client and can use options which are more secure and offer faster performance than the client attempts by default. Additionally, without using a profile it may not be possible to create a mobile IPsec configuration which can natively allow connections from different client types.

Import the CA to the iOS Device

As with other clients, the CA certificate must be installed on the client. Importing the CA Certificate to the client device is a relatively easy process. The first step is to get the CA Certificate to the client device. The easiest way to accomplish this is via e-mail as shown in Figure *iOS Mail Client Receiving CA Certificate*

To install the certificate from e-mail:

- Send the CA Certificate only (not the key) to an e-mail address reachable from the client device
- Open the Mail app on the client device
- Open the message containing the CA Certificate
- Tap the attachment to install the CA Certificate and a prompt labeled **Install Profile** will show as seen in *iOS CA Certificate Install Profile Prompt*

  Note: Newer versions of iOS may copy the CA to an entry under the Settings app for review before it can be installed. The device will instruct the user how to proceed when this happens. Typically, this involves opening the Settings app and tapping **Profile Downloaded**.

- Tap **Install** in the upper right, and a warning screen is presented as shown in *iOS CA Certificate Install Warning*
- Tap **Install** in the upper right once more to confirm and then one final prompt is presented as seen in *iOS CA Certificate Confirmation Prompt*
- Tap **Install** at the confirmation prompt and the CA Certificate is now stored as a trusted entry.

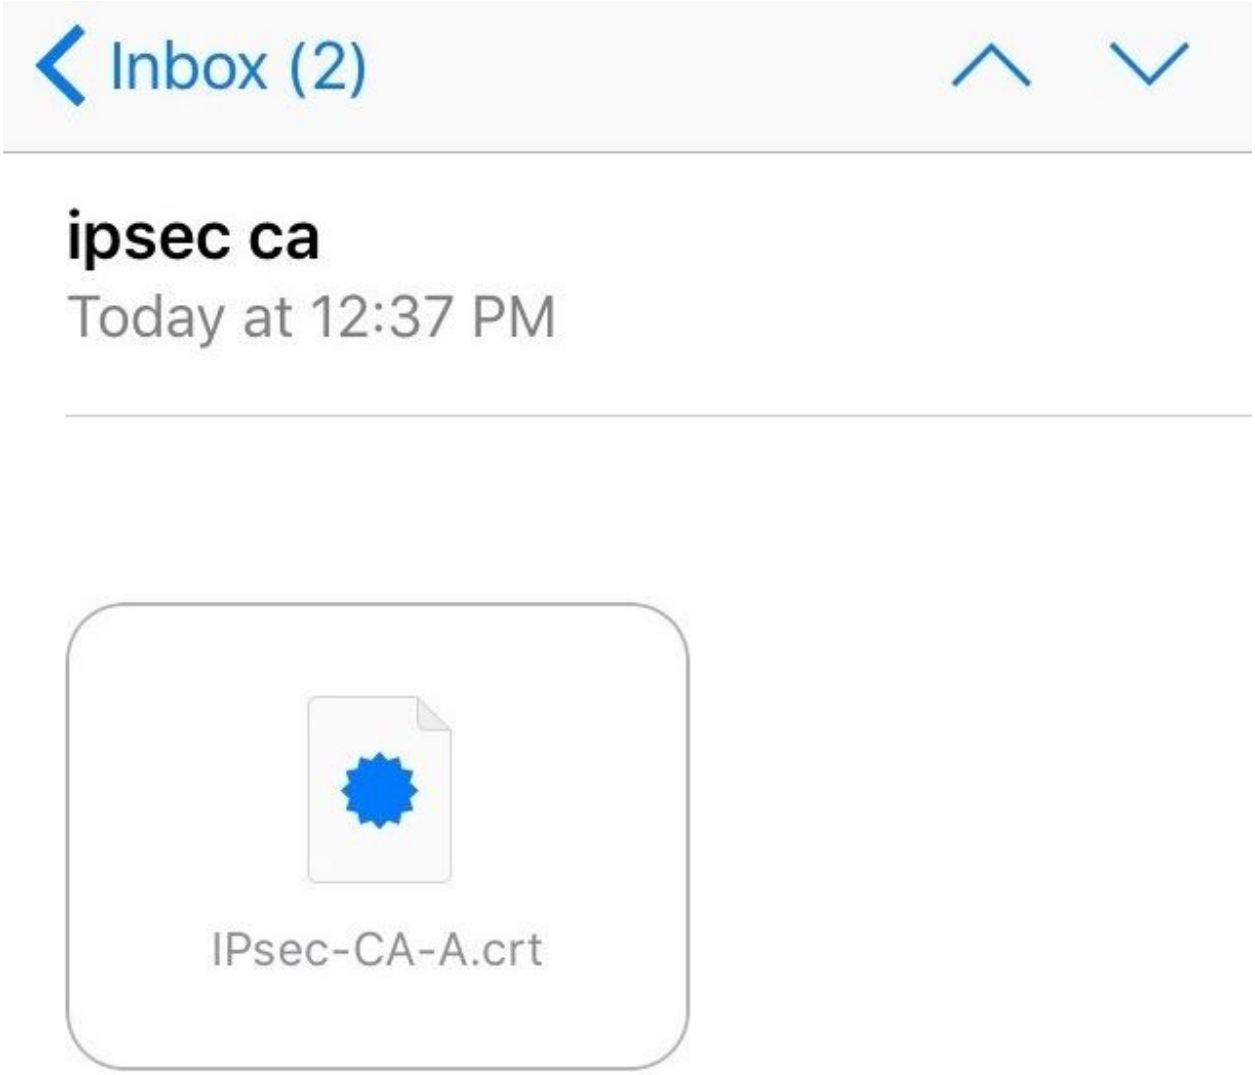


Fig. 30: iOS Mail Client Receiving CA Certificate

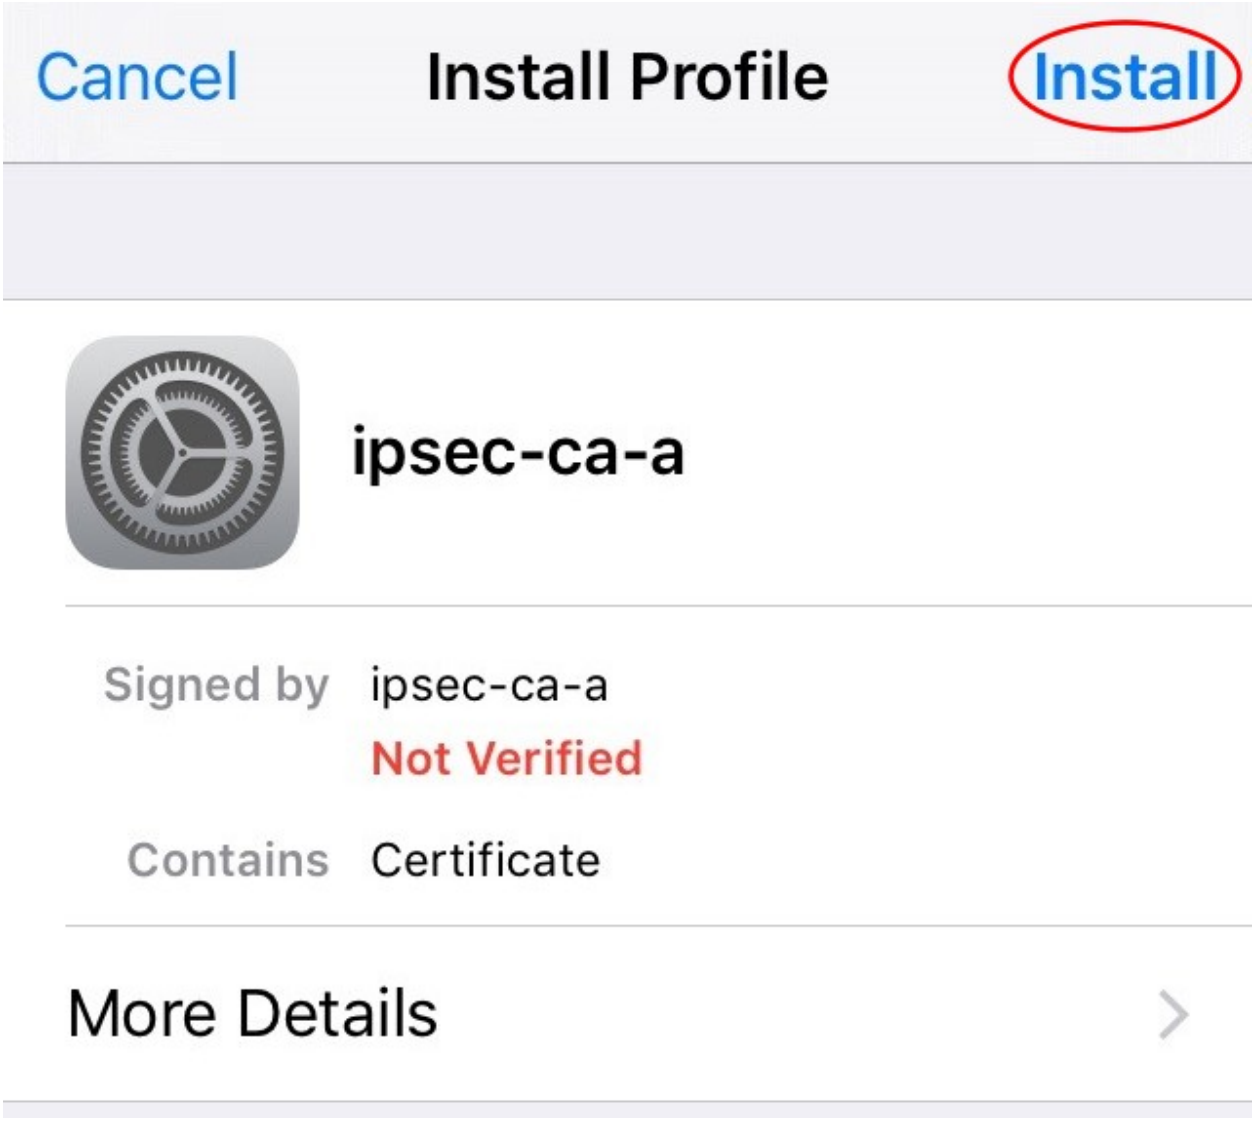


Fig. 31: iOS CA Certificate Install Profile Prompt

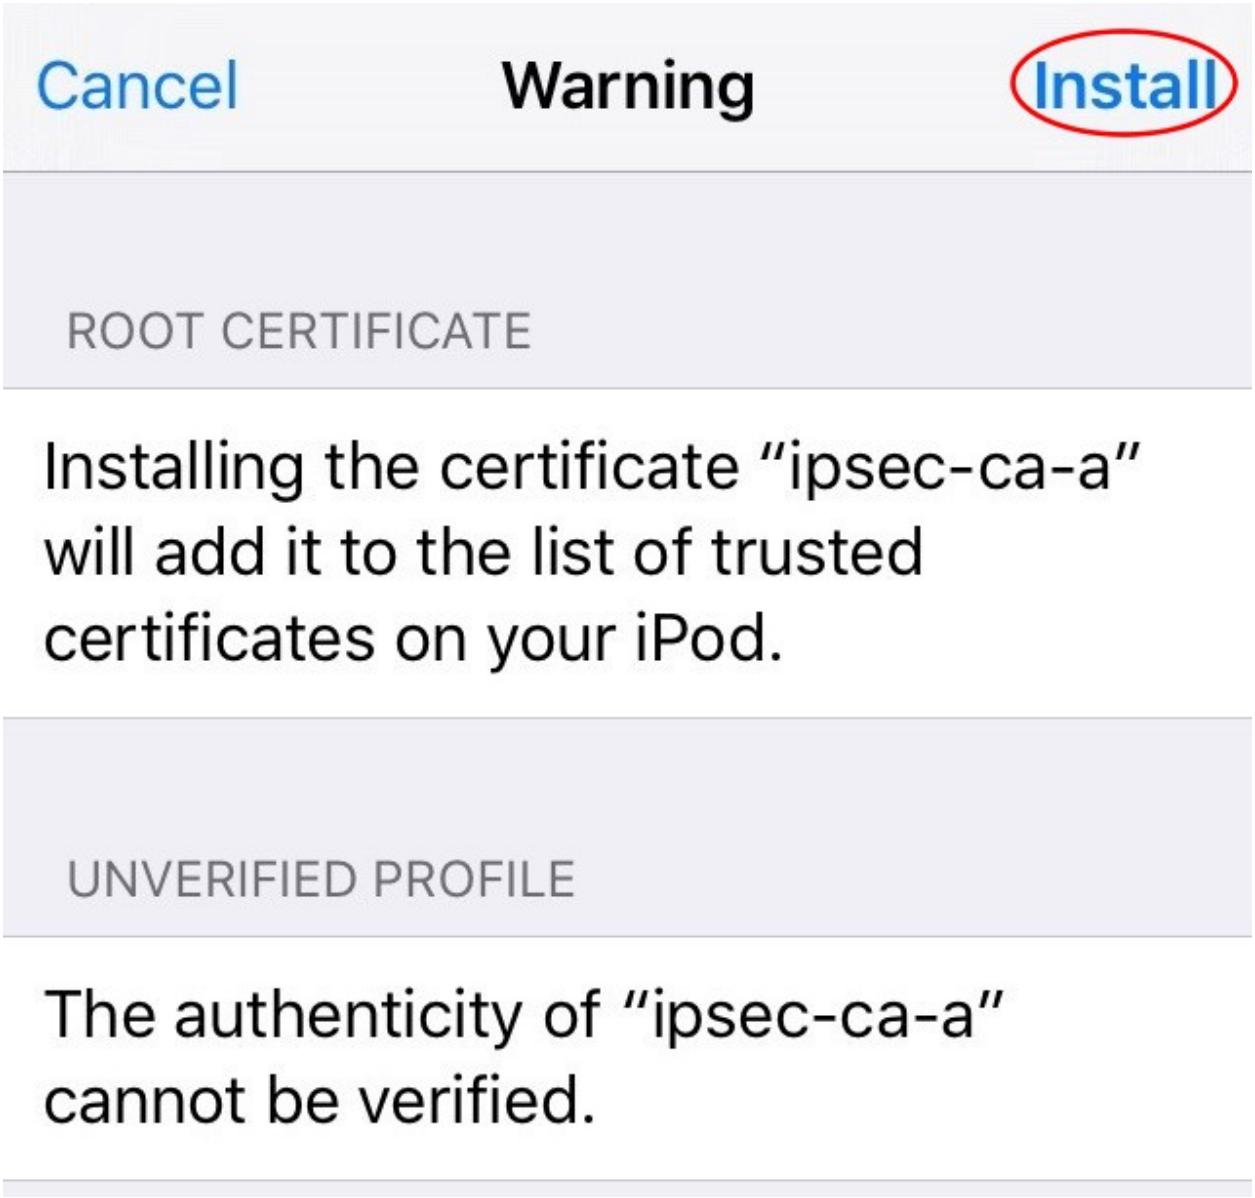


Fig. 32: iOS CA Certificate Install Warning

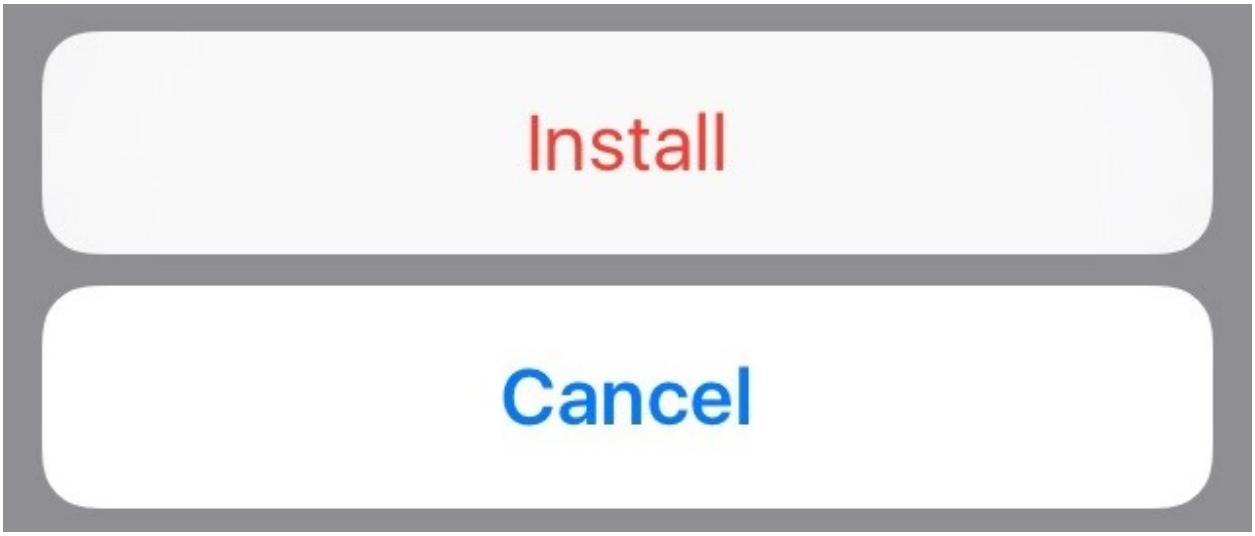


Fig. 33: iOS CA Certificate Confirmation Prompt

Set up the VPN Connection

Once the CA Certificate has been installed, a VPN entry must be configured:

- Open **Settings**
- Tap **VPN**
- Tap **Add VPN Configuration**
- Set **Type** to *IKEv2* (default)
- Fill in the settings as follows:

  Description
  A name for the VPN connection, `ExampleCo VPN`

  Server
  The hostname of the firewall in DNS

  Note: This must match a SAN value in the server certificate.

  Remote ID
  The hostname of the firewall again

  Note: This must match a SAN value in the server certificate.

  Local ID
  Leave blank

  User Authentication
  Username

  Username
  The username for this user.

Note: With EAP-MSCHAPv2 the **Username** is the **Identifier** configured for the user entry on the **Pre-Shared Keys** tab under **VPN > IPsec**. With EAP-RADIUS this would be the username set on the RADIUS server.

Password
The password for this user.

- Tap **Done** to complete the VPN entry. When complete, it looks similar to *iOS IKEv2 Client Settings*

Connecting and Disconnecting

The VPN may be connected or disconnected by visiting the VPN entries under **Settings**. This varies a bit but typically shows in at least two places:

- Settings > VPN
- Settings > General > VPN

The entry directly under **Settings** appears near the top of the list with the other Network entries (Airplane mode, Wi-Fi, and Bluetooth) once there is at least one VPN connection present.

Once in the VPN list, the VPN entry must be selected (shows a checkmark next to its entry) and then the slider may be moved to the "On" position to connect.

35.25.5 Configuring IPsec IKEv2 Remote Access VPN Clients on Ubuntu

This document demonstrates how to configure an IKEv2 EAP-MSCHAPv2 or EAP-RADIUS connection on Ubuntu. This procedure was performed on Linux Mint 20.2, but the procedure is identical on most recent similar distributions.

Before starting, install `network-manager-strongswan` and `strongswan-plugin-eap-mschapv2` using `apt-get` or a similar mechanism.

Set up the VPN Connection

- Copy the CA Certificate for the VPN from the firewall to the workstation
- Click the **Network Manager** icon in the notification tray by the clock

Note: The icon varies depending on the type of network in use.

- Click **Network Connections**
- Click **Add**
- Select *IPsec/IKEv2 (strongswan)* under **VPN** as shown in *Adding an IKEv2 VPN on Ubuntu*

Note: If the option is not present, double check that `network-manager-strongswan` is installed.

- Click **Create**
- Select the **VPN** Tab
- Set the fields as follows:

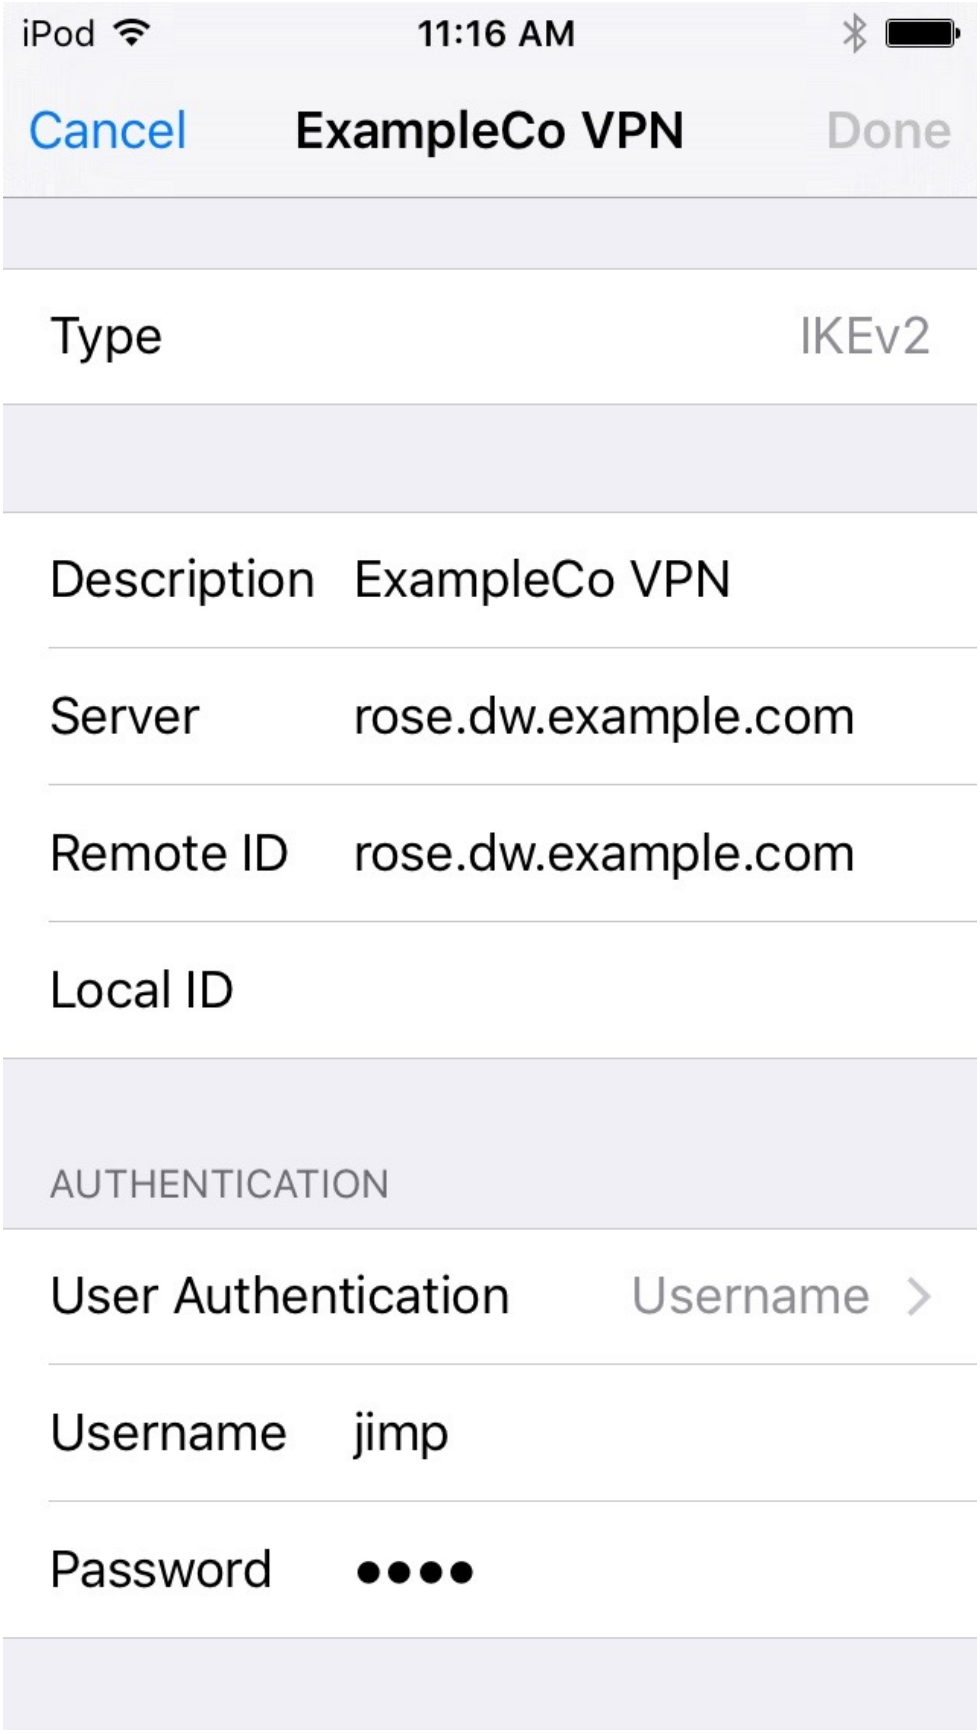


Fig. 34: iOS IKEv2 Client Settings



Fig. 35: iOS VPN List

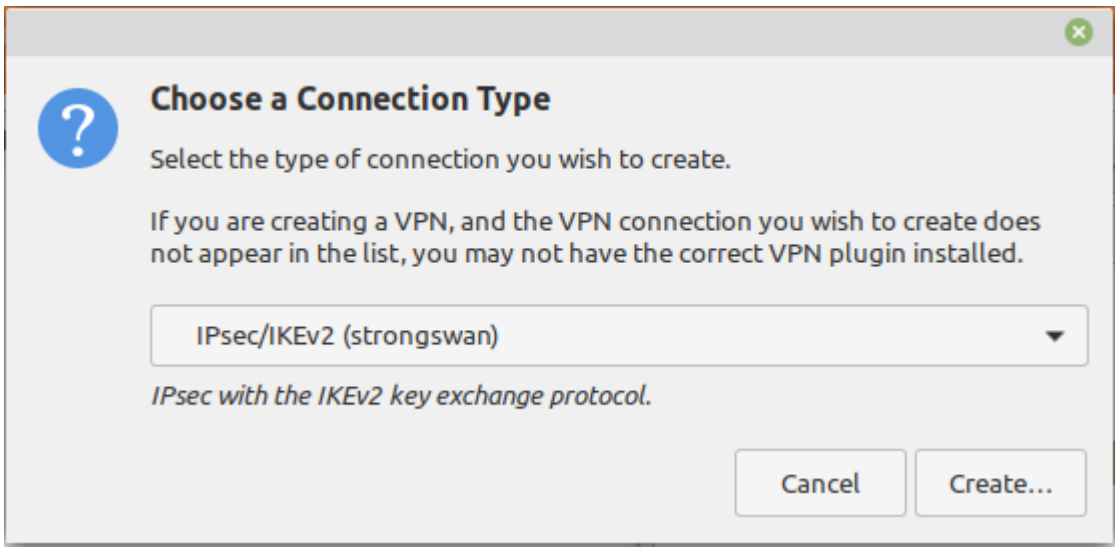


Fig. 36: Adding an IKEv2 VPN on Ubuntu

Connection Name
A name for this connection, `ExampleCo Mobile VPN`.

Address
The **Address** of the firewall, `vpn.example.com`.

Certificate
Click the field and browse to find the downloaded CA Certificate file.

Authentication
EAP

Username
The username to be used for this connection, `alice`.

Password
Click the icon in the **Password** field and select the desired action. The default behavior is to ask for the password on every connection.

To store the password, pick one of the options which allow storing the value then set it in this field.

Request an inner IP address
Checked

- Compare the settings to those shown in figure *Ubuntu VPN Client Settings*
- Click **Save**
- Click **Close**

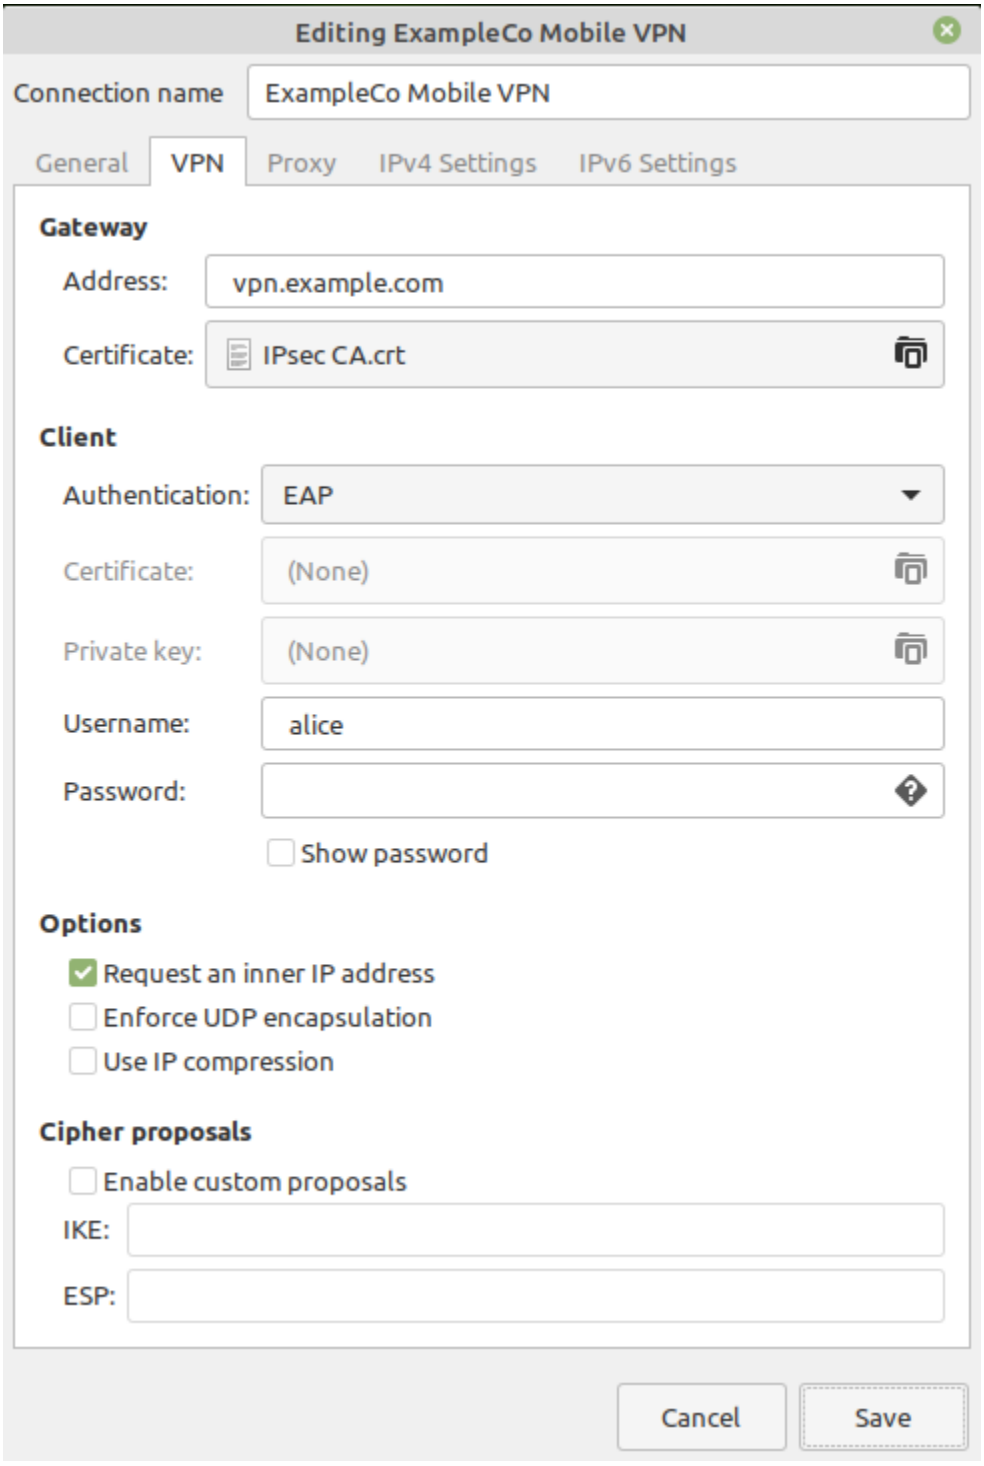


Fig. 37: Ubuntu VPN Client Settings

Connecting and Disconnecting

To Connect:

- Click the **Network Manager** icon
- Click the VPN Name or click **VPN Connections** to move the slider to the *On (1)* position

---

Note: If a password prompt does not appear, the network manager service may need restarted or a reboot of the workstation may be necessary.

To Disconnect:

- Click the **Network Manager** icon
- Click **VPN Connections** to move the slider to the *Off (0)* position

35.26 IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2

- *IKEv2 Server Configuration*
  - *IKEv2 Certificate Structure*
  - *Mobile Client Settings*
  - *Mobile IPsec User Creation*
  - *Firewall Rules*
- *Client Configuration*

35.26.1 IKEv2 Server Configuration

There are several components to the server configuration for mobile clients:

- Creating a certificate structure for the VPN
- Configuring the IPsec **Mobile Client** settings
- Creating the phase 1 and phase 2 for the client connection
- Adding IPsec firewall rules.
- Create user credentials for the VPN

IKEv2 Certificate Structure

Tip: A certificate created by the ACME package (*ACME package*) will be natively trusted by many clients and can be used in place of a manually created private CA and server certificate.

Windows clients may not accept an ACME certificate as it may lack a property for IKE in the extended key usage attributes. This check can be disabled, however. See *Disable EKU Check*.

Create a Certificate Authority

If a suitable Certificate Authority (CA) is not present in the certificate manager, creating one is the first task:

- Navigate to **System > Certificates**
- Click  **Add** to create a new certificate authority
- Set the options as follows:

  Descriptive Name
  `Mobile IPsec CA`

  Method
  Create an internal Certificate Authority

  Randomize Serial
  Checked

  Common Name
  `mobile-ipsec-ca`

- Leave the rest of the fields at their default values or adjust to suit local preferences
- Click **Save**

Create a Server Certificate

Warning: Follow these directions exactly, paying close attention to how the server certificate is created at each step. If any one part is incorrect, some or all clients may fail to connect.

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add** to create a new certificate
- Set the options as follows:

  Method
  Create an internal Certificate

  Descriptive Name
  `IKEv2 Server`

  Certificate Authority
  Mobile IPsec CA

Common Name
The hostname of the firewall as it exists in DNS, e.g. `vpn.example.com`.

If clients connect by IP address, place the IP address here instead.

Certificate Type
Server Certificate

Alternative Names

Type
IP Address

Value
The WAN IP address of the firewall, e.g. `198.51.100.3`

Add more **Alternative Names** as needed for additional hostnames or IP addresses on the firewall that clients may use to connect.

- Click **Save**

Mobile Client Settings

The next step is to choose an IP address range to use for mobile clients. Ensure that IP addresses do not overlap any existing network.

> **Warning:** The IP addresses used by mobile clients must differ from those in use at the site hosting the mobile tunnel as well as the LAN from which the client will be connecting.

This example uses `10.3.200.0/24`, but it can be any unused subnet.

- Navigate to **VPN > IPsec**, **Mobile Clients** tab
- Enable IPsec:

Enable IPsec Mobile Client Support
Checked



Fig. 38: Enable Mobile IPsec Clients

- Set the authentication options as follows:

User Authentication
Local Database as seen in Figure *Mobile Clients Authentication*.

This setting is not needed for EAP-MSCHAPv2, but it must have something selected.

> **Tip:** RADIUS servers defined in the User Manager (*User Management and Authentication*) can be selected here for authenticating users when using EAP-RADIUS.

- Set the **Client Configuration** options

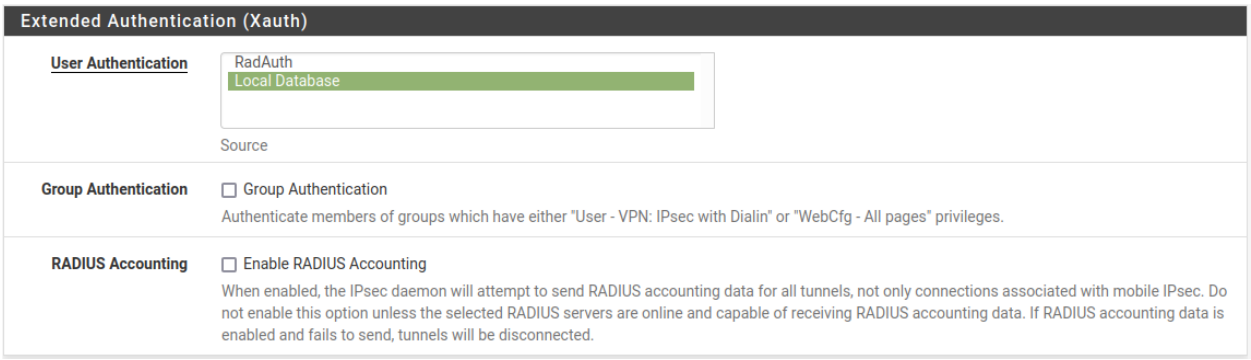


Fig. 39: Mobile Clients Authentication

These settings may be pushed to the client, such as the client IP address and DNS servers. These options are shown in Figure *Mobile Clients Pushed Settings*. Compatibility with these options varies between clients, but is common in most current operating systems.

Virtual Address Pool
Checked, `10.3.200.0/24`

Virtual IPv6 Address Pool
Checked, `2001:db8:1:df01::/64`

Network List
Checked

DNS Default Domain
Checked, `example.com`

Split DNS
Checked, `example.com example.org`

DNS Servers
Checked, `10.3.0.1`

See also:

For additional information these options work, see *Client Configuration*.

- Click **Save**
- Click  **Create Phase 1** at the top of the screen if it appears

Phase 1

The phase 1 configuration for mobile clients must be configured as follows:

Description
`Mobile IPsec` or another suitable description

Key Exchange Version
IKEv2

Internet Protocol
IPv4 for this example as it only uses an IPv4 WAN

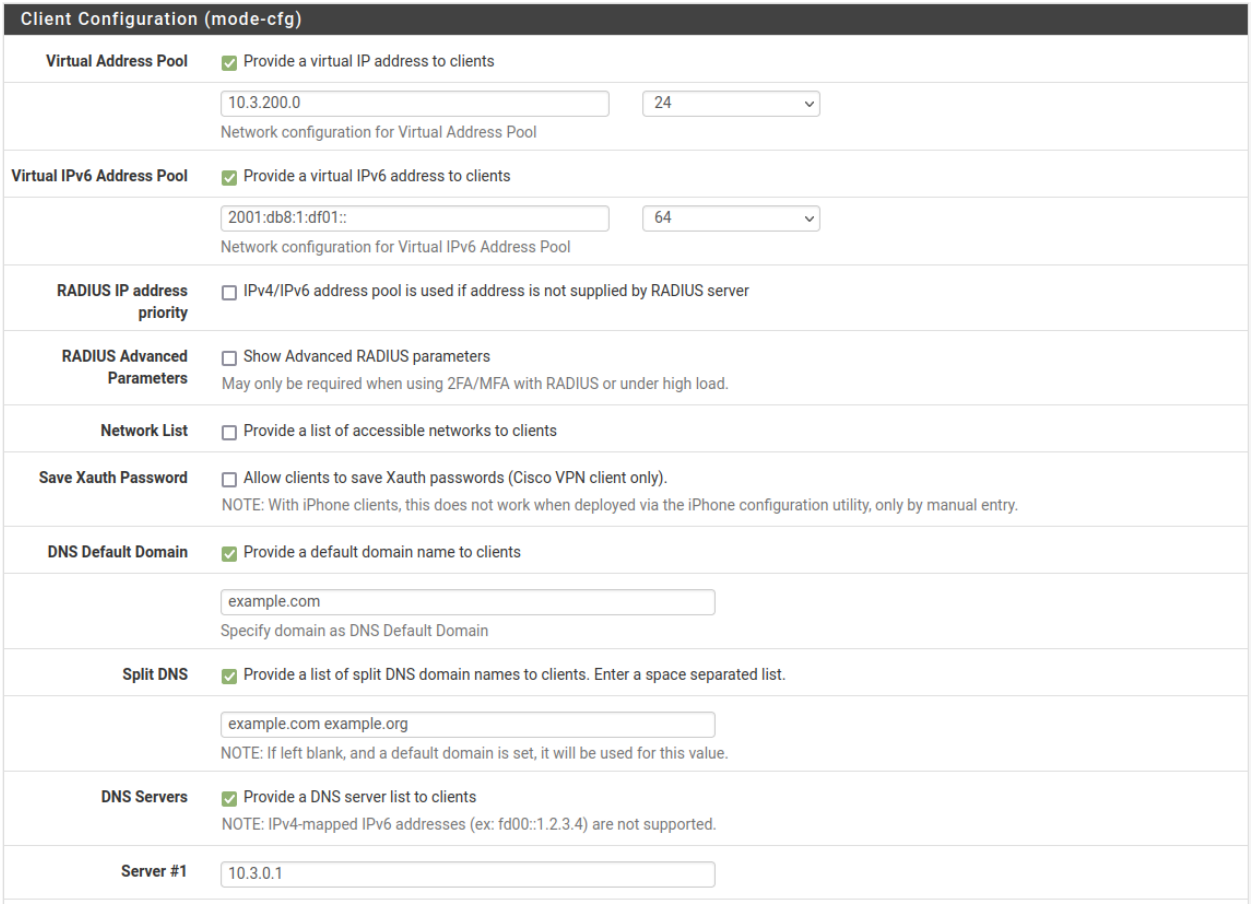


Fig. 40: Mobile Clients Pushed Settings



Fig. 41: Mobile Clients Phase 1 Creation Prompt

Interface
WAN

Authentication Method
EAP-MSCHAPv2

My identifier

Choose *Fully Qualified Domain Name* from the drop-down list and then
enter the hostname of the firewall, `vpn.example.com`.

> **Warning:** This must match the value in the server certificate.

Peer Identifier
Any

My Certificate
Choose the IPsec server certificate created earlier

Encryption Algorithm
Add multiple combinations of encryption, hash, and DH options to accommodate various clients with different requirements. Click  **Add Algorithm** to add more entries. They should be added with the most secure and preferred options first.

A good starting set of options is:

- **Algorithm** *AES256-GCM*, **Hash** *SHA256*, **DH Group** *16*
- **Algorithm** *AES256-GCM*, **Hash** *SHA256*, **DH Group** *2*
- **Algorithm** *AES 256*, **Hash** *SHA256*, **DH Group** *14*
- **Algorithm** *AES 256*, **Hash** *SHA1*, **DH Group** *14*

Life Time
`28800`

MOBIKE
Set to *Enable* to allow clients to roam between IP addresses, otherwise set to *Disable*.

- Click **Save**

Phase 2

- Click  **Show Phase 2 Entries** to expand the list of mobile phase 2 entries
- Click  **Add P2** to add a new mobile phase 2.

Description
`Mobile IPv4`

Mode
Tunnel IPv4

Local Network
Set to *LAN subnet* or another local network.

> **Tip:** To tunnel all traffic over the VPN, use *Network* and enter `0.0.0.0` with a mask of *0*

Protocol
ESP

Encryption algorithms
Set to combinations of options which will be accepted by all possible clients. A good starting set is:

- *AES*, *128*
- *AES128-GCM*, *128*
- *AES256-GCM*, *128*

Hash algorithms
Similar to the algorithm, select all options required by clients:

SHA256, *SHA384*, *SHA512*

PFS
14 (2048 bit)

> **Warning:** When configuring a VPN on Apple iOS manually, it is not compatible with PFS in phase 2 as demonstrated in *Configuring IPsec IKEv2 Remote Access VPN Clients on iOS*. Most other clients require PFS. Rather than creating a VPN connection on iOS manually, the best practice is to use a VPN profile from the Apple Configurator or *IPsec Export Package* which allows the use of PFS.

Lifetime
`3600`

- Add another phase 2 entry but this time for IPv6. The settings are identical to the previous entry with the following changes:

  Description
  `Mobile IPv6`

  Mode
  Tunnel IPv6

  Local Network
  LAN subnet as before or use the network value `::` with a mask of *0* to tunnel all traffic.

- Click **Save**
- Click Apply changes

The tunnel setup for mobile clients is complete.

Mobile IPsec User Creation

The next step is to add users for use by EAP-MSCHAPv2.

- Navigate to **VPN > IPsec**, **Pre-Shared Keys** tab
- Click  **Add** to add a new key
- Configure the options as follows:

  Identifier
  The username for the client.

  This can be expressed in multiple ways, such as an e-mail address like `jimp@example.com`

  Secret Type
  EAP

  Pre-Shared Key
  The password for the client, for example `abc123`

- Click **Save**
- Repeat as many times as needed for additional VPN users.

A complete user is shown in Figure *Mobile IPsec User*.

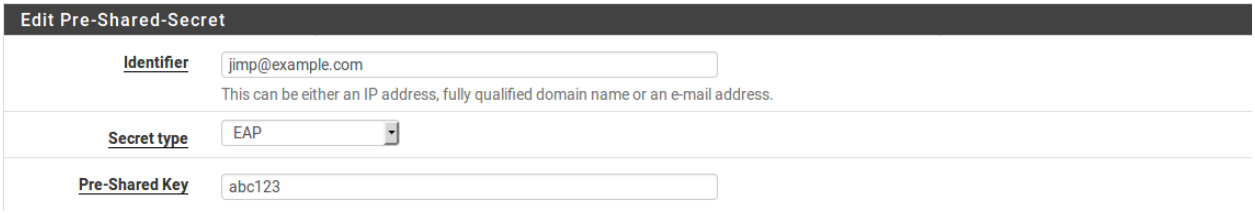


Fig. 42: Mobile IPsec User

Firewall Rules

As with the static site-to-site tunnels, mobile tunnels also require firewall rules at **Firewall > Rules** on the **IPsec** tab. In this instance the source of the traffic would be the subnet chosen for the mobile clients and the destination will be the LAN network, or *any* if tunneling all traffic.

See also:

For more details, *IPsec and firewall rules*.

35.26.2 Client Configuration

Each mobile client device needs a VPN instance or client configured. In some cases a third-party IPsec client may be required. There are many different IPsec clients available for use, some free, and some commercial applications. With IKEv2, as used in this example, many operating systems have native VPN clients and do not need extra software.

Common clients are covered at *Configuring IPsec IKEv2 Remote Access VPN Clients*.

35.27 IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS

To set up IKEv2 with EAP-RADIUS, follow the directions for *IKEv2 with EAP-MSCHAPv2* with a slight variation:

- Define a RADIUS server under **System > User Manager**, **Servers** tab before starting
- Select the RADIUS server on **VPN > IPsec**, **Mobile Clients** tab
- Check **Group Authentication** and select **Authentication Groups** list entries to optionally filter access based on *RADIUS group* membership
- Select *EAP-RADIUS* for the **Authentication method** on the Mobile IPsec phase 1 entry

35.27.1 EAP-RADIUS with FreeRADIUS

The default settings are OK for this use case. If the defaults do not work, see *Using EAP and PEAP with FreeRADIUS*

35.27.2 EAP-RADIUS with Windows Network Policy Server (NPS)

To allow strongSwan to authenticate against NPS using EAP-MSCHAPv2, alter the NPS policy as follows:

- Open **Network Policy Server (NPS)**
- Expand **Policies**
- Click **Network Policies**
- Edit the policy currently in use
- Click on the **Constraints** tab
- Click **Authentication Methods**
- Click **Add**
- Select **Microsoft: Secured Password (EAP-MSCHAP v2)**
- Click **OK**
- Click **Apply** (To restart NPS)
- Click **OK**

35.28 IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS

Mobile IPsec using IKEv2 with EAP-TLS enables per-user certificate authentication. To authenticate against the VPN, a user must have a valid certificate signed by a specific certificate authority (CA).

The basic setup is similar to *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2*, this document will focus on the differences.

35.28.1 Setup Certificates

Per-user certificate authentication requires a certificate for the server and a set of certificates the clients.

Note: While these do not have to share the same SA, it makes the process easier.

Create a Certificate Authority

If one is not already available, then the first task is to create a new Certificate Authority as described in *Create a Certificate Authority*.

Create a Server Certificate

Create a server certificate as described in *Create a Server Certificate*.

Create Client Certificates

- Navigate to **System > Certificates**, **Certificates** tab
- Click  to create a new certificate
- Set the options as follows:

  Method
  Create an internal Certificate

  Descriptive Name
  A name associated with the client, for example `client1`.

  This is cosmetic only, so it does not affect values placed in the certificate data.

  Certificate Authority
  Mobile IPsec CA

  Common Name
  The username associated with this user, for example `client1`.

  Note: The best practice is to use identifiers in username, hostname, or FQDN formats for this field.

  Certificate Type
  User Certificate
- Change the other fields if desired to make the information more specific to the user.
- Click **Save**

Repeat as needed for additional clients.

35.28.2 Set up Mobile IPsec for IKEv2+EAP-TLS

With the certificate structure prepared, the next task is to configure the necessary IPsec settings.

Most of this configuration is identical to *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2* and only the differences will be called out.

Mobile Clients

Configure as described in *Mobile Client Settings*.

Phase 1

Configure as described in *Phase 1* but with the following changes:

Authentication method
EAP-TLS

Peer Identifier
Any

Peer Certificate Authority
Select the CA created previously for this purpose.

Phase 2

Configure as described in *Phase 2*.

35.28.3 Add Firewall Rules for IPsec

Add firewall rules to pass traffic from clients as described in *Firewall Rules*.

35.28.4 Configure the Client

The server setup is complete, but the certificates must be imported to the client.

Client configuration for a variety of operating systems is covered in *Configuring IPsec IKEv2 Remote Access VPN Clients*. (e.g. *Configuring IPsec IKEv2 Remote Access VPN Clients on Windows*).

35.29 IPsec Site-to-Site VPN Example with Pre-Shared Keys

A site-to-site IPsec tunnel interconnects two networks as if they were directly connected by a router. Systems at Site A can reach servers or other systems at Site B, and vice versa. This traffic may also be regulated via firewall rules, as with any other network interface. If more than one client will be connecting to another site from the same controlled location, a site-to-site tunnel will likely be more efficient, not to mention more convenient and easier to maintain.

With a site-to-site tunnel the devices on either local network need not have any knowledge that a VPN exists. No client software is required and all the work is handled by the tunnel endpoints. This is also a good solution for devices that have network connectivity but do not handle VPN connections, such as printers, cameras, HVAC systems, and other embedded hardware.

See also:

- *IPsec*
- *IPsec Configuration*
- *Troubleshooting Duplicate IPsec SA Entries*
- *Client Routing and Gateway Considerations*
- *Accessing Firewall Services over IPsec*

35.29.1 Site-to-site example configuration

The key to making a working IPsec tunnel is to ensure that both sides have matching settings for authentication, encryption, and so on. Before starting make a note of the local and remote WAN IP addresses as well as the local and remote internal subnets that will be carried across the tunnel. Aside from the cosmetic tunnel **Description** and these pieces of information the other connection settings will be identical.

The following settings are assumed by this example and some of the subsequent examples in the IPsec recipes:

Table 11: IPsec Endpoint Settings

Site A		Site B	
Name	Austin Office	Name	London Office
WAN IP	198.51.100.3	WAN IP	203.0.113.5
LAN Subnet	10.3.0.0/24	LAN Subnet	10.5.0.0/24
LAN IP	10.3.0.1	LAN IP	10.5.0.1

Note: To avoid issues with security association duplication, this example uses settings described in *Troubleshooting Duplicate IPsec SA Entries*.

Figure *Site-to-Site IPsec* shows the general layout of this VPN.

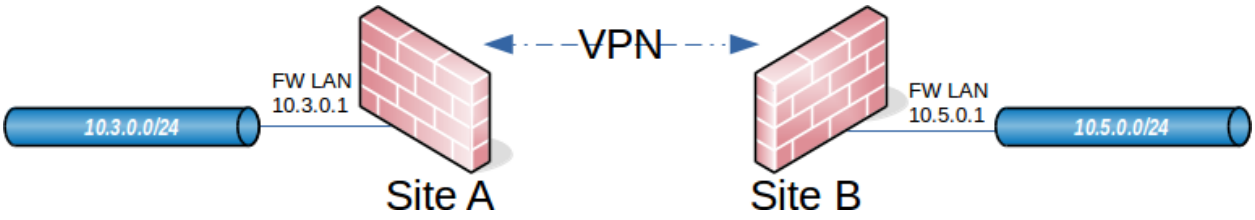


Fig. 43: Site-to-Site IPsec

35.29.2 Site A

Start with configuring the tunnel and related settings on the firewall at Site A.

Phase 1

To add a new IPsec phase 1:

- Navigate to **VPN > IPsec**
- Click  **Add P1**
- Fill in the settings as described below
- Click **Save** when complete

Use the following settings for the phase 1 configuration. Many of these settings may be left at their default values unless otherwise noted.

See also:

For comprehensive coverage of all IPsec phase 1 settings, see *Phase 1 Settings*.

First fill in the top section that holds the general phase 1 information and IKE endpoint configuration, as shown in Figure *figure-vpn-tunnel-settings*. Items in bold are required. Fill in the settings as described:

Description
Text describing the purpose or identity of the tunnel. The best practice is to put the name of Site B in this box, and brief detail about the purpose of the tunnel to help with future administration.

For this example `ExampleCo London Office` is used for the **Description** to identify where this example tunnel terminates.

Disabled
Uncheck this box so that the tunnel will be operational.

Key Exchange version
Specifies whether to use IKEv2 or IKEv1. IKEv2 is the best practice when both endpoints are compatible. If one side cannot use IKEv2, use IKEv1 instead.

Internet Protocol
IPv4 in most cases unless both WANs have IPv6, in which case either type may be used.

Interface
Most likely set to *WAN*, but see the note at *Interface Selection* on selecting the proper interface when unsure.

Remote Gateway
The WAN address at Site B, `203.0.113.5` in this example.

The next section controls IPsec phase 1 proposals for authentication. The defaults are desirable for most of these settings which simplifies the process.

Authentication Method
The default, *Mutual PSK*, is used for this example.

My Identifier
The default, *My IP Address*, is kept for this example.

Peer Identifier
The default, *Peer IP Address*, is kept for this example.

Pre-Shared Key
Use a strong key, at least 10 characters in length containing a mix of upper and lowercase letters,

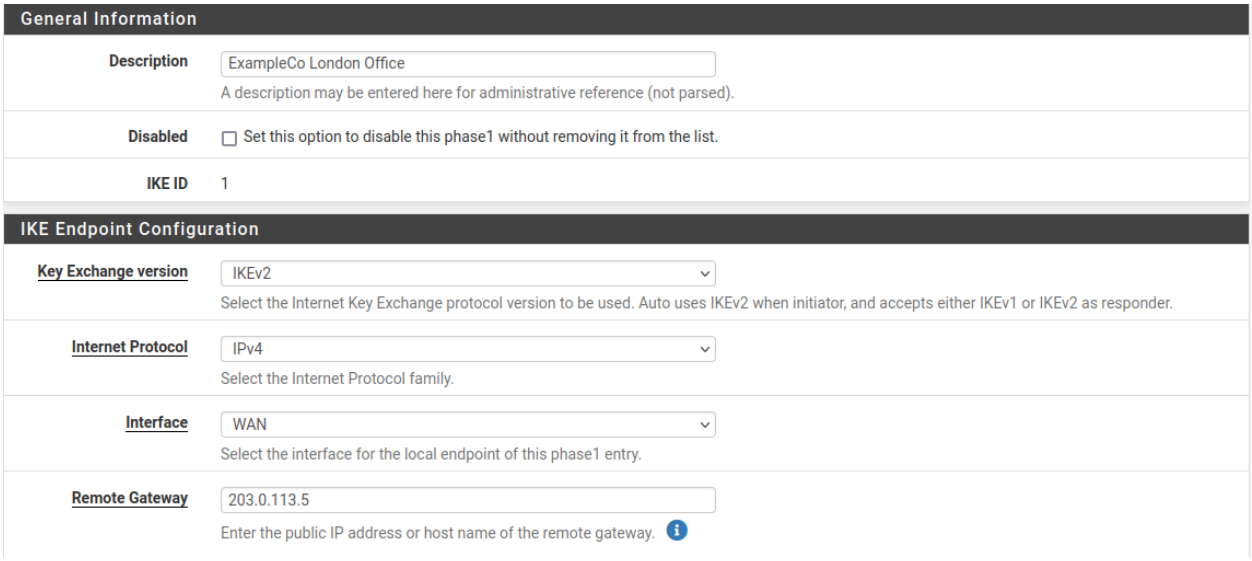


Fig. 44: Site A IPsec Phase 1 General Information and IKE Endpoint Configuration

numbers and symbols. Enter a custom key or click **Generate new Pre-Shared Key** to automatically populate the field with a random long string suitable for use as a Pre-Shared Key.

> **Warning:** This is the most important setting to get correct. As mentioned in the VPN overview, IPsec using pre-shared keys can be broken if the tunnel uses a weak key.

The **exact** same key must be entered into the tunnel configuration for Site B later, so note it down or copy and paste it elsewhere. Copy and paste may come in handy, especially with a complex key.

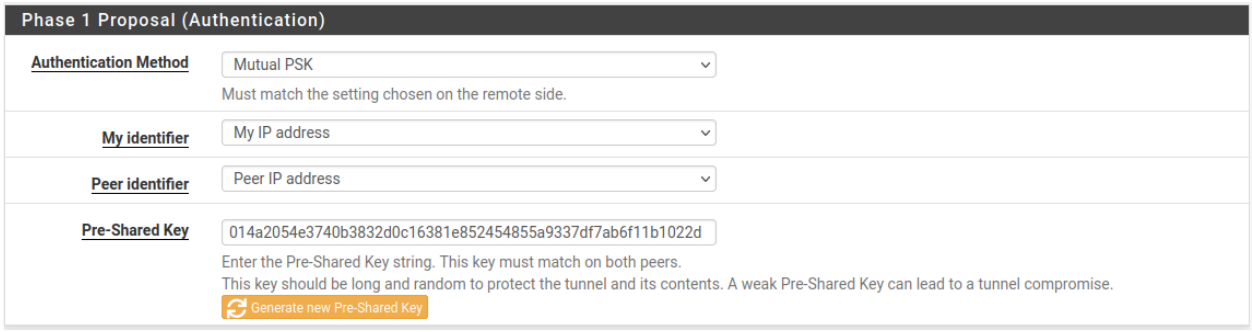


Fig. 45: Site A Phase 1 Authentication Settings

The next section controls IPsec phase 1 proposals for encryption.

Encryption Algorithm
Use *AES* with a **Key Length** of *256 bits*.

Hash Algorithm
Use *SHA256* at a minimum if both sides are compatible, otherwise use the strongest hash available on both endpoints.

DH Group
The default of *14 (2048 bit)* is OK, higher values are more secure but may use more CPU.

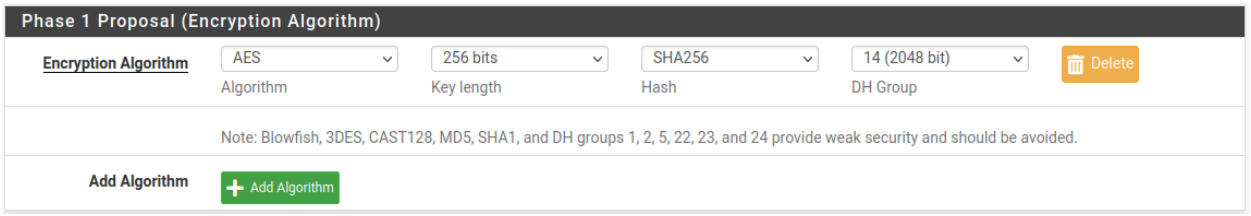


Fig. 46: Site A Phase 1 Encryption Settings

The **Expiration and Replacement** section controls the timing and method by which the phase 1 will be renegotiated.

Life Time
The default `28800` is OK for this endpoint.

See also:

See *Troubleshooting Duplicate IPsec SA Entries* for best practices in choosing lifetime values.

The other lifetime-related values (**Rekey Time**, **Reauth Time**, **Rand Time**) should be left at their defaults on this endpoint as they are automatically calculated as the correct values.

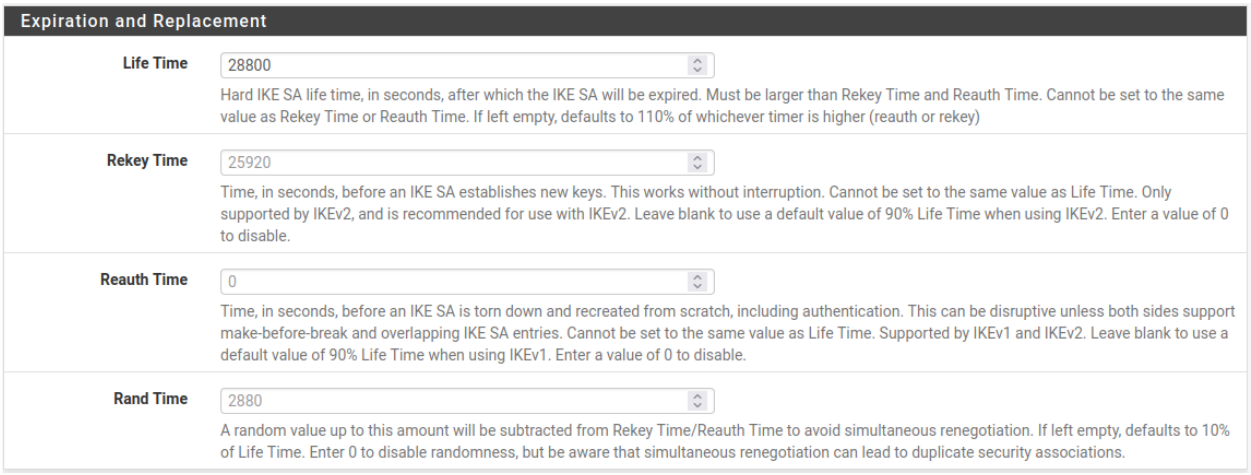


Fig. 47: Site A Phase 1 Lifetime Settings

Finally, the **Advanced** section contains a couple settings to check as well:

Child SA Close Action
Set this endpoint to **Restart/Reconnect** so that the phase 2 entries will be reconnected if they get disconnected.

Dead Peer Detection
Leave checked and at the default values.

Click **Save** to complete the phase 1 setup.

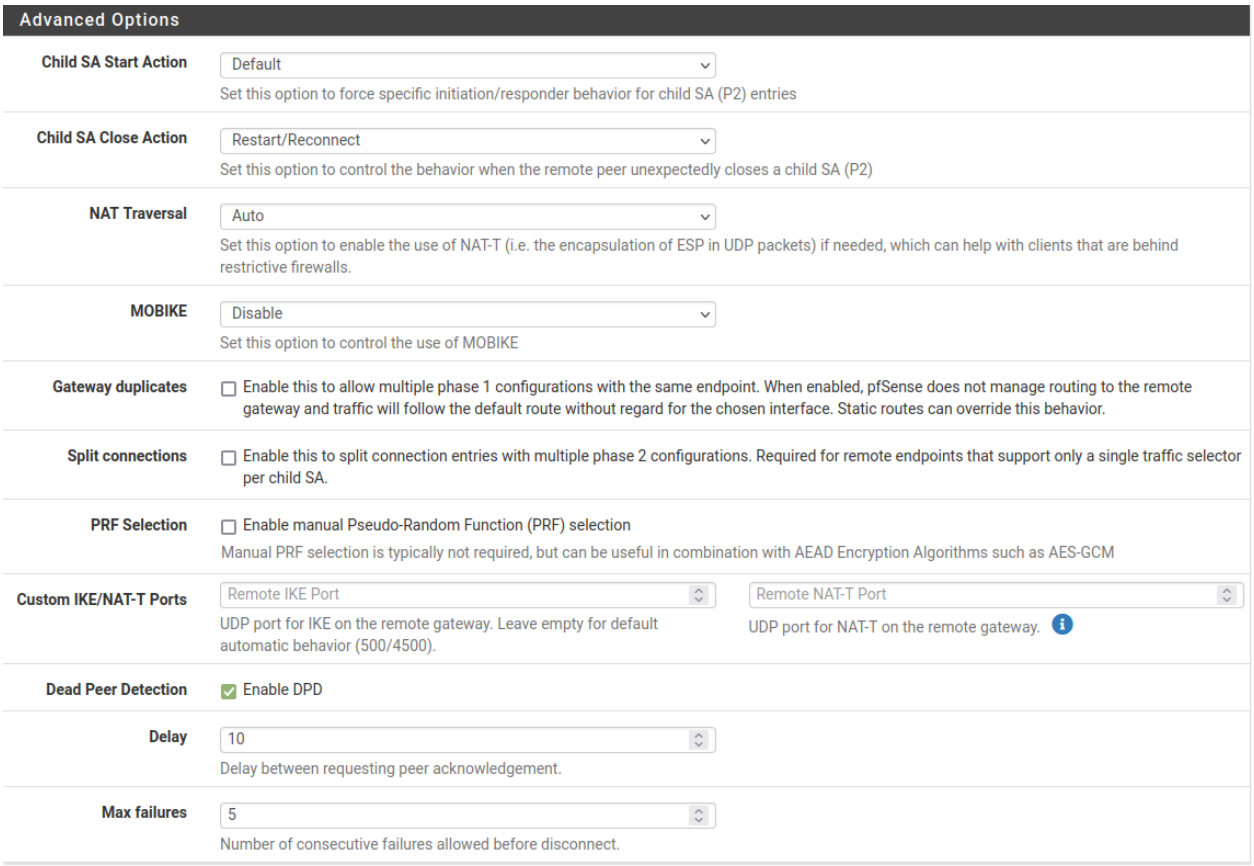


Fig. 48: Site A Phase 1 Advanced Settings

Phase 2

With the phase 1 entry complete, now a new phase 2 definition to the VPN:

- Click  **Show Phase 2 Entries** as seen in Figure *Site A Phase 2 List (Empty)* to expand the phase 2 list for this VPN.
- Click  **Add P2** to add a new phase 2 entry, as seen in Figure *Adding a Phase 2 entry to Site A*.

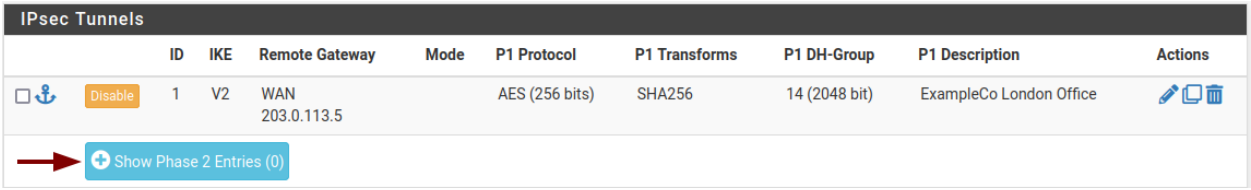


Fig. 49: Site A Phase 2 List (Empty)

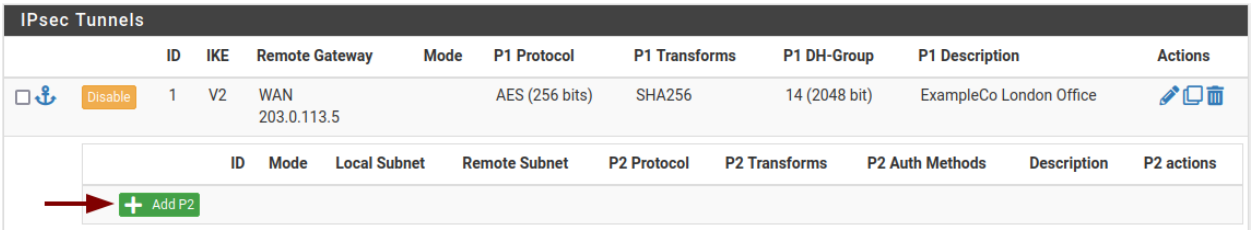


Fig. 50: Adding a Phase 2 entry to Site A

Now add settings for phase 2 on this VPN. The settings for phase 2 (Figure *Site A Phase 2 General Information and Networks*) can vary more than phase 1.

See also:

For comprehensive coverage of all IPsec phase 2 settings, see *Phase 2 Settings*.

Description
A brief description of the network(s) involved in this phase 2 entry.

Mode
Since this example is for a policy-based tunnel, select *Tunnel IPv4*

Local Network
In most cases the best practice is to leave this as *LAN Subnet*, but it can be changed to *Network* with the proper subnet value filled in. In this case that would be `10.3.0.0/24`. Leaving it as *LAN Subnet* will ensure that if the network is renumbered in the future, this end of the tunnel will follow. If that does happen, the other end must be changed manually.

NAT/BINAT
Set to *None*.

Remote Network
Set to the network at Site B, in this case `10.5.0.0/24`.

The next section of the phase 2 settings covers traffic encryption. **Encryption algorithms** and **Hash algorithms** can both be set to allow multiple options in phase 2, and both sides will negotiate and agree upon the settings so long as

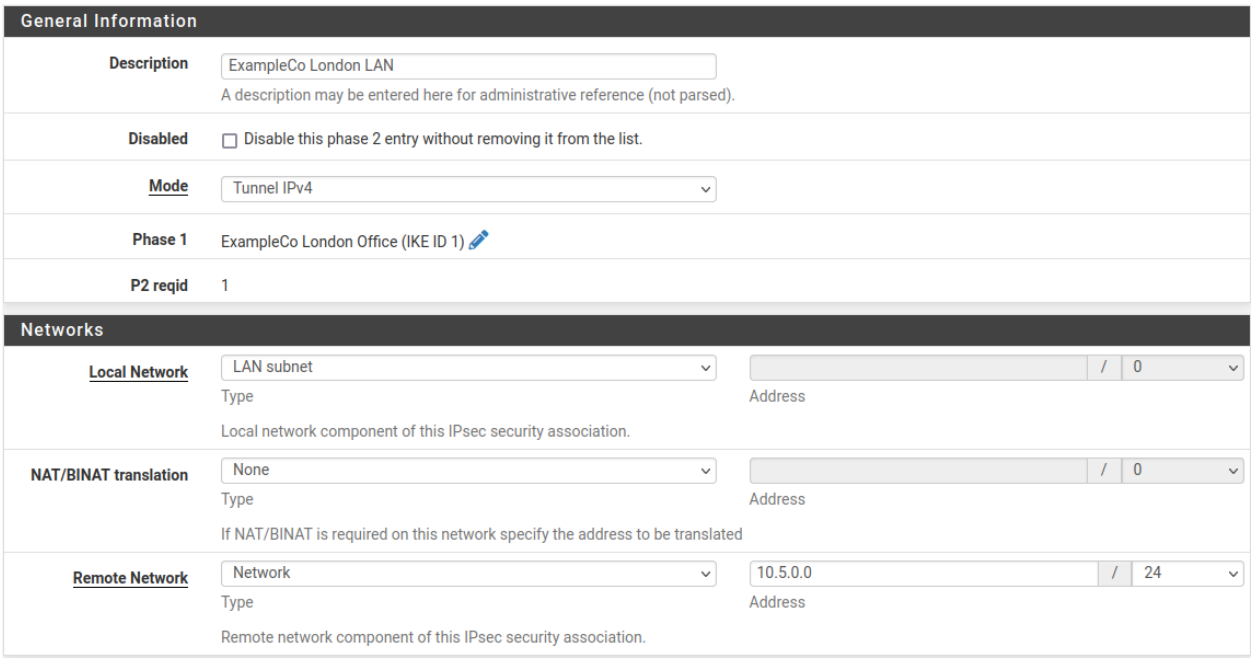


Fig. 51: Site A Phase 2 General Information and Networks

each side has at least one of each in common. In some cases that may be a good thing, but it is usually better to restrict this to the single specific options desired on both sides.

Protocol
Set to *ESP* for encryption.

Encryption algorithm
The best practice is to use an AEAD cipher such as AES-GCM if it is compatible with both endpoints.

Select *AES256-GCM* with a *128-bit* key length. Otherwise, use *AES 256*, or the highest strength cipher available on both endpoints.

Hash algorithm
If AES-GCM is selected for **Encryption Algorithm** do not select any hashes.

Otherwise, use *SHA256* or the strongest hash available on both endpoints.

PFS
Perfect Forward Secrecy (PFS) is optional but can help protect against certain key attacks. This example uses *14 (2048 bit)*.

Next are the settings which govern timing and methods for renewing phase 2 keys.

Life Time
Use `3600` for this example, and leave **Rekey Time** and **Rand Time** at their default calculated placeholder values.

To finalize the settings and put them into action:

- Click *Save*
- Click *Apply Changes* on the IPsec Tunnels screen, as seen in Figure *Apply IPsec Settings*.

The tunnel configuration for Site A is now complete.

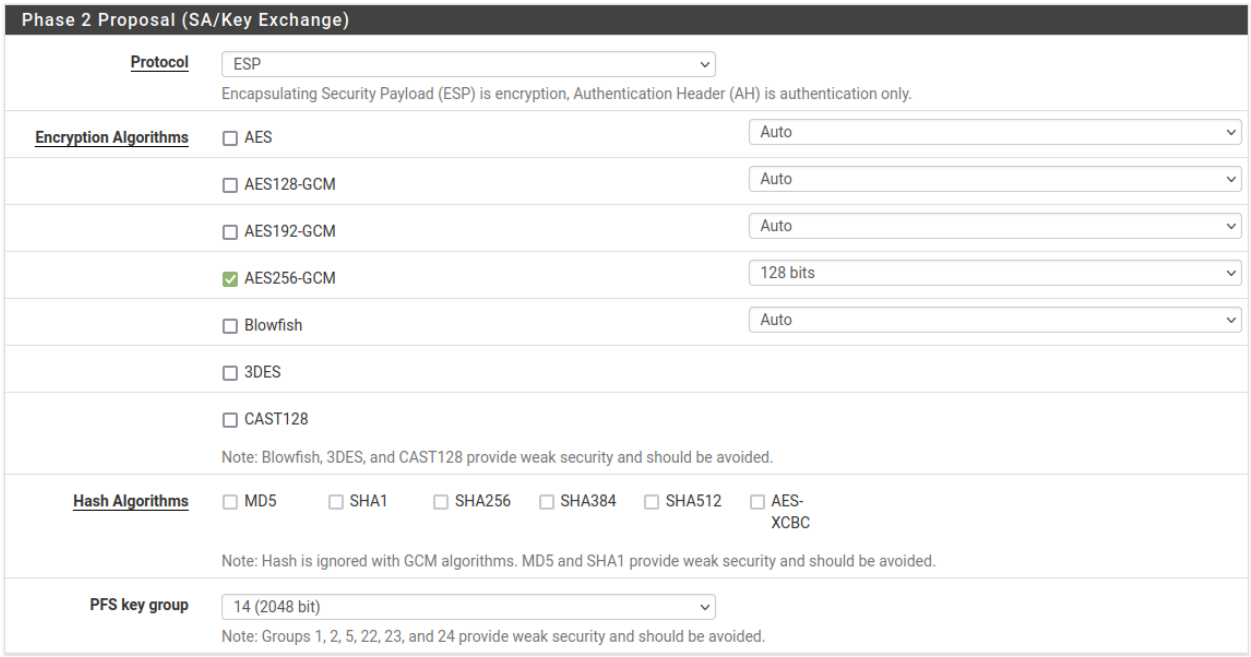


Fig. 52: Site A Phase 2 Proposal Settings

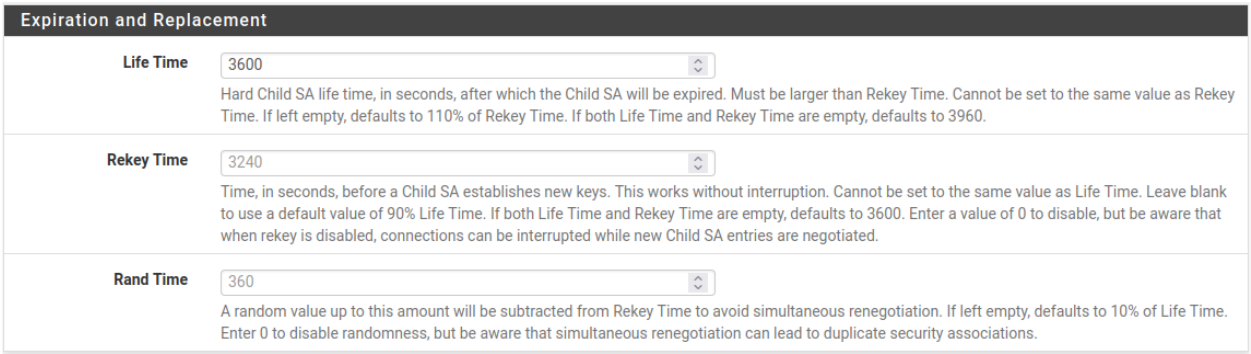


Fig. 53: Site A Phase 2 Expiration and Replacement Settings



Fig. 54: Apply IPsec Settings

Firewall Rules

Firewall rules are necessary to allow traffic from the network at Site B to enter through the IPsec tunnel.

Navigate to **Firewall > Rules** on the **IPsec** tab and add rules there to pass traffic from the remote side of the VPN.

See also:

See *Firewall* for specifics on adding rules, and *IPsec and firewall rules* for firewall rule advice specific to IPsec.

Rules may be as permissive or restrictive as desired. For example, they can allow any protocol from anywhere to anywhere or only allow TCP from a certain host on Site B to a certain host at Site A on a certain port.

As with other firewall rules the connections are checked on the way *into* the firewall; the source of all traffic on the IPsec tab rules will be remote VPN networks, such as those at Site B.

Make sure the source addresses on the firewall rules match Site B addresses, such as `10.5.0.0/24`. The destination addresses will be on Site A, such as `10.3.0.0/24`.

35.29.3 Site B

Now that Site A is configured, it is time to tackle Site B. Repeat the process on the Site B endpoint to add a tunnel.

Only a few parts of this setup will differ from Site A as shown in Figure *Site B Phase 1 General Settings* and Figure *Site B Phase 2 General Settings*:

- The phase 1 settings for Description, WAN address, Life Time, Child SA Start Action, and Child SA Close Action
- The phase 2 tunnel networks and Life Time

Add a phase 1 entry to the Site B firewall using similar settings to Site A, but with the following differences:

Description
`ExampleCo Austin Office`.

Remote Gateway
The WAN address at Site A, `198.51.100.3`.

Life Time
At least 10% higher than Site A, `31680`

Child SA Start Action
Set to **None (Responder Only)** so that this endpoint will not initiate on its own, but will wait for Site A to initiate.

Child SA Close Action
Set this endpoint to **Close Connection and clear SA** so that the phase 2 will not automatically reconnect, since Site A will be managing that.

- Click **Save**

Add a phase 2 entry to the Site B firewall using similar settings to Site A, but with the following differences:

Description
`ExampleCo Austin LAN`.

Remote Subnet
The network at Site A, in this case `10.3.0.0/24`.

Life Time
At least 10% higher than Site A, `5400`

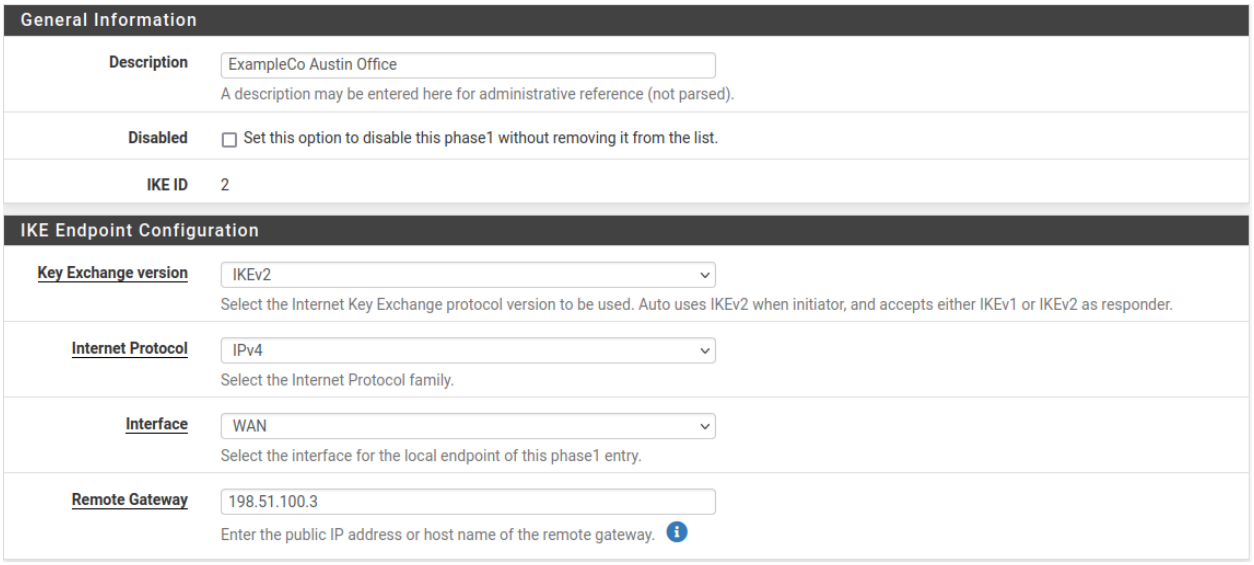


Fig. 55: Site B Phase 1 General Settings

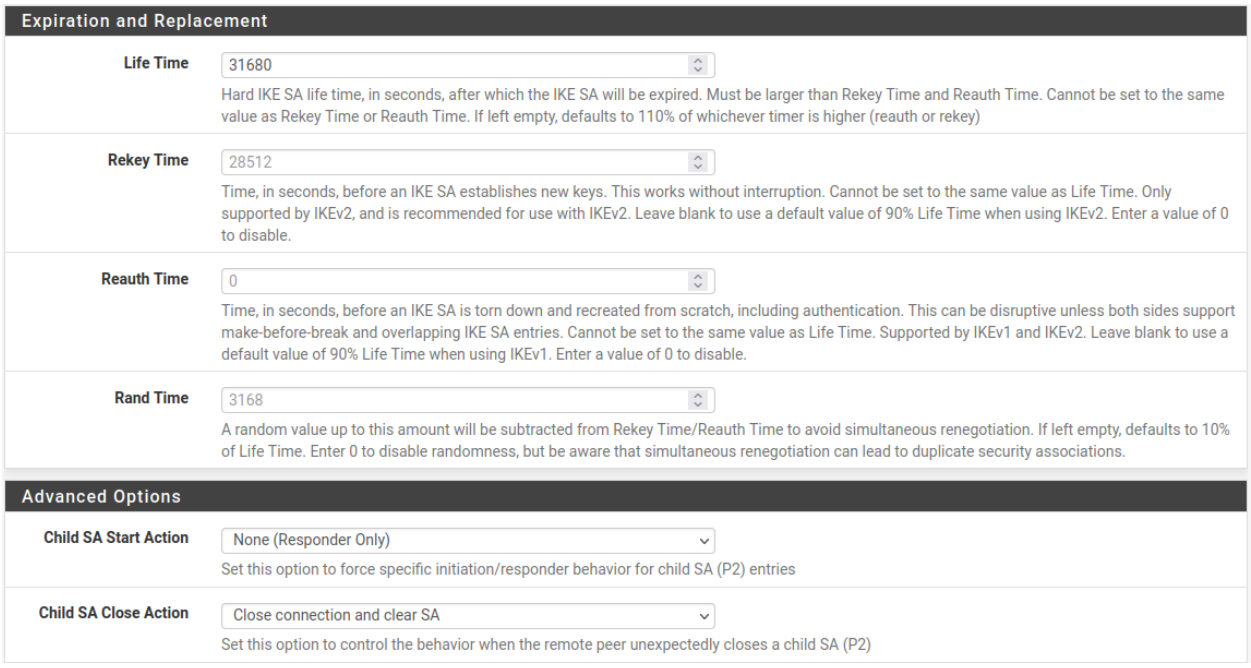


Fig. 56: Site B Phase 1 Other Settings

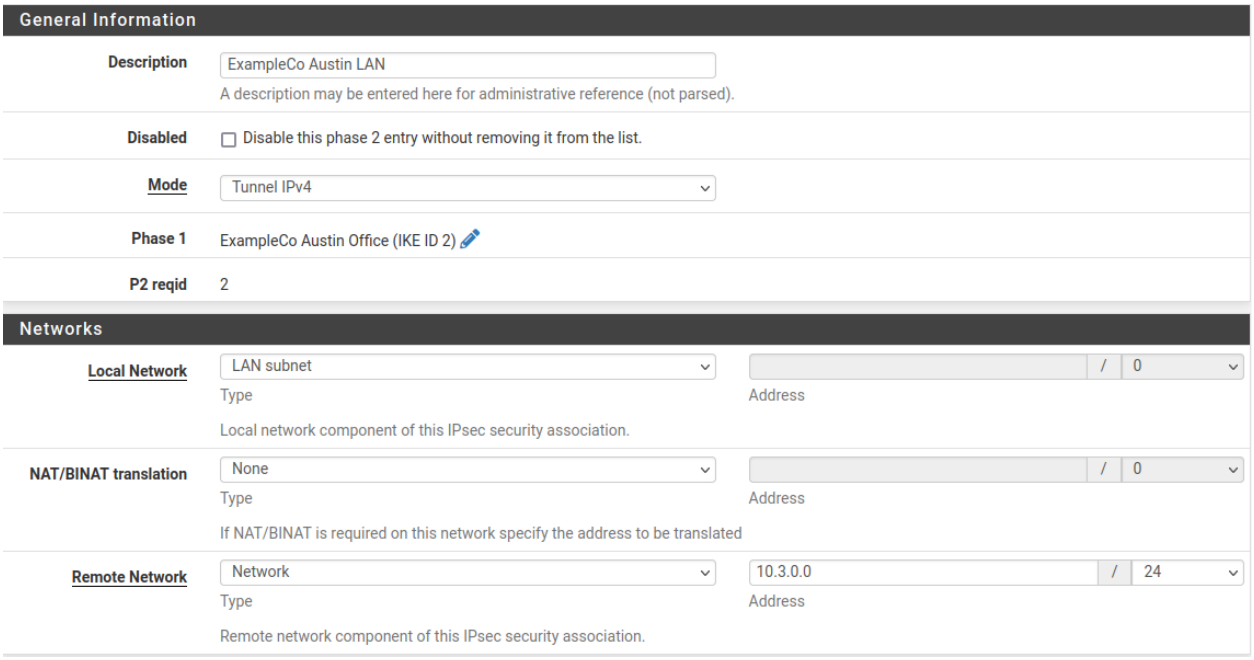


Fig. 57: Site B Phase 2 General Settings

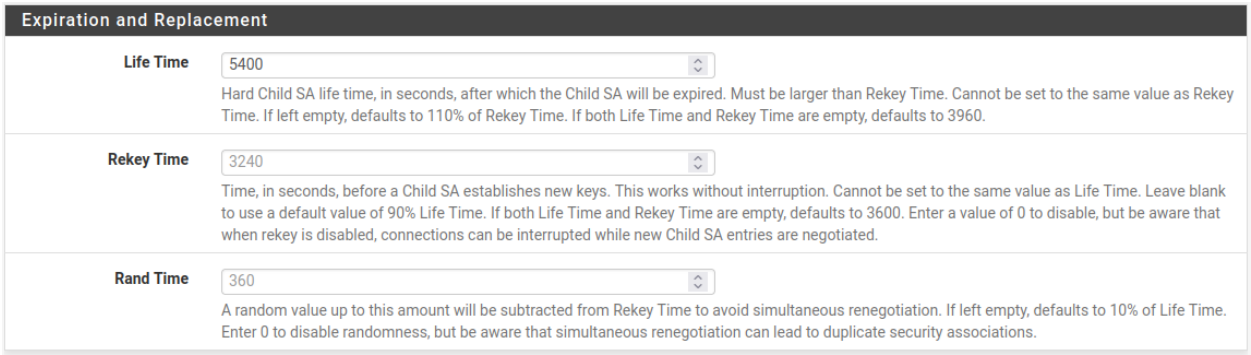


Fig. 58: Site B Phase 2 Lifetime Settings

- Click **Save**
- Click **Apply changes** on the IPsec Tunnels screen.

As with Site A, firewall rules must also be added to allow traffic on the tunnel to cross from Site A to Site B. Add these rules to the **IPsec** tab under **Firewall > Rules**. For more details, see *IPsec and firewall rules*. This time, the source of the traffic would be Site A, destination Site B.

35.29.4 Check Status

Both tunnels are now configured and active. Check the IPsec status by visiting **Status > IPsec**. A description of the tunnel is shown along with its status.

If the tunnel is not listed as **Established**, there may be a problem establishing the tunnel. At this point in time, the most likely reason is that no traffic has yet attempted to cross the tunnel.

A connect button is offered on this screen that will attempt to initiate the tunnel. Click the  **Connect VPN** button to attempt to bring up the tunnel as seen in Figure *Site A IPsec Status*.

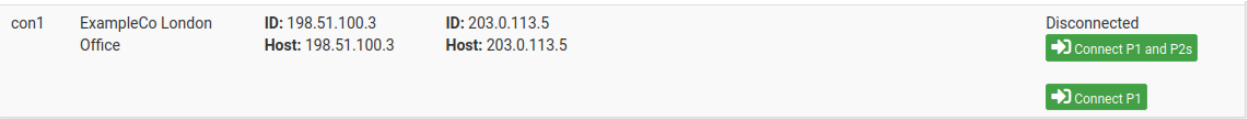


Fig. 59: Site A IPsec Status

If the connect button does not appear, try sending an ICMP ping to a system in the remote subnet at Site B from a device inside the phase 2 local network at Site A (or vice versa) to see if the tunnel establishes. Look at *Testing IPsec Connectivity* for other means of testing a tunnel.

Failing that, the IPsec logs will typically offer an explanation. They are located under **Status > System Logs** on the **IPsec** tab. Be sure to check the status and logs at both sites. For more troubleshooting information, check the *Troubleshooting IPsec VPNs* section later in this chapter.

When the tunnel is connected the status will look like Figure *Site A IPsec Status while Connected*.

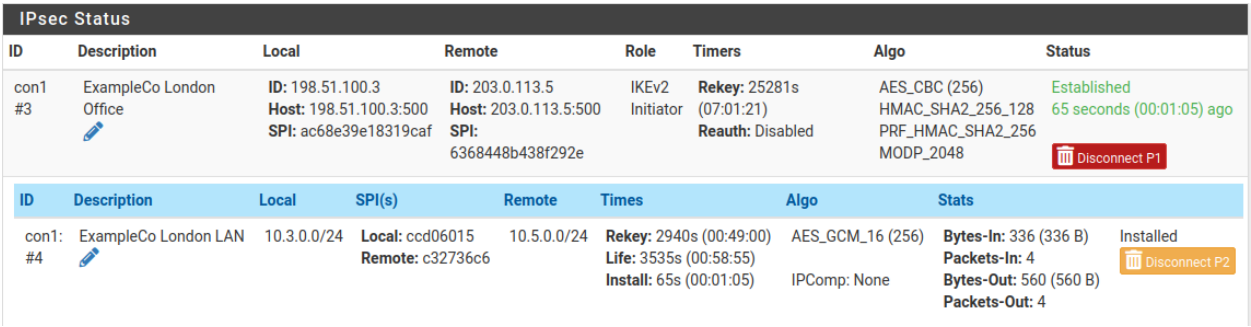


Fig. 60: Site A IPsec Status while Connected

35.30 Routing Internet Traffic Through a Site-to-Site IPsec Tunnel

It is possible to use IPsec on a firewall running pfSense® software to send Internet traffic from a remote site such that it appears to be coming from another location. This may be needed if a vendor requires that connections originate from a specific address.

The basis of this tunnel is a working site-to-site IPsec VPN as described in *IPsec Site-to-Site VPN Example with Pre-Shared Keys*. Refer to that recipe for detailed instructions. Only the differences from that recipe will be mentioned here.

As a reminder, this example uses two sites:

- **Site A** is the main site. The Internet traffic will *exit* this location.
- **Site B** is a remote office with LAN subnet `10.5.0.0/24`. This is the source of local traffic which will traverse the tunnel and reach the Internet through site A.

The only differences from tunnel in *IPsec Site-to-Site VPN Example with Pre-Shared Keys* are:

Site A, phase 2

Local Network
`0.0.0.0/0`

Site B, phase 2

Remote Network
`0.0.0.0/0`

This will cause the firewall to send **all** traffic from the LAN through the IPsec tunnel to the remote end of the tunnel.

35.30.1 Allow IPsec traffic through the firewall

Since this tunnel must pass traffic from the Internet, the firewall rules must be fairly lenient. The rules on site A will need to pass traffic from a source of the site B LAN (`10.5.0.0/24`) to a destination of *any*.

Tip: To prevent site B from reaching sensitive local resources at site A or sites connected to additional VPNs, place block rules above the rule passing the Internet traffic.

The rules at site B do not necessarily have to allow much traffic back through unless there are public resources at site B which will be reached across the tunnel (e.g. 1:1 NAT, port forwards).

35.30.2 Configure outbound NAT

For site B to reach the Internet, site A must perform outbound NAT on the traffic from the site B LAN (`10.5.0.0/24`) as it leaves the WAN.

To do this, first change the outbound NAT mode on the site A firewall:

- Navigate to **Firewall > NAT**, **Outbound** tab
- Set the **Outbound NAT Mode** to **Hybrid Outbound NAT**

Note: If site A is already on this mode or set to **Manual**, then do not change the mode.

- Click **Save**

Using this mode will allow the default automatic NAT rules to continue working without needing a full manual ruleset. Now add a custom rule to the top of the list which will match site B:

- Click  **Add**
- Set the following values:

  Source
  Network, `10.5.0.0/24`

  Destination
  Any

  Translation Address
  Interface Address

  Description
  `NAT for IPsec tunnel Site B`

- Click **Save**
- Click **Apply changes**.

The new entry is now in the outbound NAT rule list.

At this point site B will have a working Internet connection through the IPsec tunnel and the Internet provider at site A. Any Internet traffic from site B will look as if it were coming from site A.

35.31 IPsec Site-to-Site VPN Example with Certificate Authentication

Using certificate-based authentication for identification of VPN tunnel peers is much stronger than using a simple Pre-Shared Key, but it is more difficult to configure and manage.

Certificate authentication requires a PKI structure. Depending on the setup, each side may utilize its own certificate authority (CA) or they may share a common CA. This example utilizes a different CA on each node to more closely resemble connecting to third parties.

See also:

CA and certificate entries can be created and imported in the GUI by the *Certificate Manager*.

35.31.1 Required Information

Endpoint A:

Item	Value
Hostname	office.vpn.example.com
WAN IP Address	`198.51.100.16`

Endpoint B:

Item	Value
Hostname	home.vpn.example.com
WAN IP Address	`198.51.100.17`

35.31.2 Create CA

First, create a Certificate Authority (CA) on each side:

On Endpoint A:

- Navigate to **System > Certificates**, **CAs** tab
- Click  **Add**
- Set the options as follows:

  Descriptive Name
  `Office VPN CA`

  Method
  Create an internal Certificate Authority

  Randomize Serial
  Checked

  Common Name
  `office-vpn-ca`
- Leave the rest of the fields at their default values or adjust to suit local preferences
- Click **Save**
- Click  to export this CA as a file in the browser

On Endpoint B:

- Navigate to **System > Certificates**, **CAs** tab
- Click  **Add**
- Set the options as follows:

  Descriptive Name
  `Home VPN CA`

  Method
  Create an internal Certificate Authority

  Randomize Serial
  Checked

  Common Name
  `home-vpn-ca`
- Leave the rest of the fields at their default values or adjust to suit local preferences
- Click **Save**

- Click  to export this CA as a file in the browser

35.31.3 Import Peer CAs

Next, import the new CA entries into the peer. For example, import the Home CA to the Office side, and vice versa.

Note: This step only requires the certificate data, **not** the key. The key belonging to the CA should not be copied off the firewall where it was created.

On Endpoint A:

- Navigate to **System > Certificates**, **CAs** tab
- Click  **Add**
- Set the options as follows:

  Descriptive Name
  `Home VPN CA`

  Method
  Import an existing Certificate Authority

  Certificate Data
  Paste the contents of the exported `Home VPN CA.crt` file.

- Click **Save**

On Endpoint B:

- Navigate to **System > Certificates**, **CAs** tab
- Click  **Add**
- Set the options as follows:

  Descriptive Name
  `Office VPN CA`

  Method
  Import an existing Certificate Authority

  Certificate Data
  Paste the contents of the exported `Office VPN CA.crt` file.

- Click **Save**

35.31.4 Create Endpoint Certificates

On Endpoint A:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add**
- Set the options as follows:

  Method
  Create an internal Certificate

  Descriptive Name
  `Office VPN Certificate`

  Certificate Authority
  Office VPN CA

  Common Name
  `office-vpn-cert`

  Certificate Type
  User Certificate

  Alternative Names

  Type
  FQDN or Hostname

  Value
  `office.vpn.example.com`

- Click  **Add**
- Set the new row options to:

  Alternative Names

  Type
  IP Address

  Value
  `198.51.100.16`

  Note: If the IP address is dynamic, skip this step or use the LAN IP address.

- Leave the rest of the fields at their default values or adjust to suit local preferences
- Click **Save**

On Endpoint B:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add**
- Set the options as follows:

Method
Create an internal Certificate

Descriptive Name
`Home VPN Certificate`

Certificate Authority
Home VPN CA

Common Name
`home-vpn-cert`

Certificate Type
User Certificate

Alternative Names

Type
FQDN or Hostname

Value
`home.vpn.example.com`

- Click  **Add**
- Set the new row options to:

Alternative Names

Type
IP Address

Value
`198.51.100.17`

Note: If the IP address is dynamic, skip this step or use the LAN IP address.

- Leave the rest of the fields at their default values or adjust to suit local preferences
- Click **Save**

35.31.5 Setup IPsec VPN

On both firewalls, configure the IPsec tunnel as described in *IPsec Site-to-Site VPN Example with Pre-Shared Keys*, with the following exceptions:

Endpoint A:

Authentication method
Mutual Certificate

My Identifier
Set appropriately to *match the certificate* for **this** endpoint

Peer Identifier
Set appropriately to *match the certificate* of the **peer**

My Certificate
Office VPN Certificate

Peer Certificate Authority
Home VPN CA

Endpoint B:

Authentication method
Mutual Certificate

My Identifier
Set appropriately to *match the certificate* for **this** endpoint

Peer Identifier
Set appropriately to *match the certificate* of the **peer**

My Certificate
Home VPN Certificate

Peer Certificate Authority
Office VPN CA

35.31.6 Matching Certificate and Identifiers

In order for the IPsec daemon to properly match a certificate and its keys to a peer, the local and peer identifier must match data in the certificate **exactly**.

> **Warning:** Do not place quotes (single or double) around the identifier values.

There are several ways to accomplish this matching. The key factors are:

- The IPsec daemon must be able to confirm that an endpoint matches the expected identifier, which matches a peer to a specific tunnel.
- The IPsec daemon must be able to match that identifier to a certificate and validate its trust, which confirms the identity and authenticates the tunnel peer.

The following identifier types are the best practices to use with certificate authentication:

Fully Qualified Domain Name
This choice can work with fully qualified domain names or short hostnames. If the certificates were created as specified in *Create Endpoint Certificates*, use the full hostname such as `office.vpn.example.com` or `home.vpn.example.com`. This is the easiest choice and most likely to succeed, assuming the SAN value is present in the certificate.

Modern certificates typically include the certificate CN as a SAN entry, so the CN may also be used if it resembles a hostname (e.g. `office-vpn-cert`). Check the certificate properties to ensure it is present as an FQDN SAN entry.

> **Warning:** This mode will not work if the CN contains spaces or other characters not compatible with hostnames.

ASN.1 Distinguished Name
The full ASN.1 Distinguished Name of the certificate. This is similar to the certificate subject but has stricter rules about its order.

This can be formatted in several ways so long as it matches the data in the certificate exactly, for example:

- `/CN=host.example.com/C=US/ST=Texas/L=Austin/O=Example Co`
- `CN=host.example.com, C=US, ST=Texas, L=Austin, O=Example Co`
- `CN = host.example.com, C = US, ST = Texas, L = Austin, O = Example Co`

Note: The type, number, and order of fields will vary depending on how the certificate was made.

To find this string, inspect the certificate in one of the following ways:

- From the **Certificate Manager**, **Certificates** tab, find the entry and click the  icon to expand the certificate details. In the details, copy the contents of the **DN:** field exactly.

```
DN: /CN=host.example.com/C=US/ST=Texas/L=Austin/O=Example Co
```

- Use OpenSSL on a copy of the certificate contents and look for the **Subject** contents:

```
$ openssl x509 -text -noout -in mycert.crt | grep Subject:
        Subject: CN = host.example.com, C = US, ST = Texas, L = Austin,␣
↪O = Example Co
```

- If the certificate is configured in IPsec already, look at how `strongSwan` reports the certificate subject:

```
$ swanctl --list-certs | grep subject
  subject:  "CN=host.example.com, C=US, ST=Texas, L=Austin, O=Example Co
↪"
```

Warning: When copying these values remember that they must be entered **exactly** as shown but **without** any single or double quotes around the string. Only include the DN contents and not any headers or labels such as `DN:` or `Subject:`.

My IP Address / Peer IP Address

These choices are viable if all the following items are true:

- Both endpoints have static IP addresses
- These static IP addresses are used as the **Remote Gateway** address on each side of the IPsec tunnel
- The static IP address of an endpoint is present in its certificate as a SAN

IP Address

Similar to the **My IP Address / Peer IP Address** case above, but instead of using endpoint static IP addresses, uses a pre-determined local addresses instead. This could be the LAN IP address or another agreed upon address which does not change. This value does not need to match the **Remote Gateway** address in this case.

- The value must be present as an IP address type SAN in the certificate

In most cases, this is not ideal, and the hostname is easier to use instead.

35.31.7 Troubleshooting

If the IPsec daemon cannot match an identifier to a known certificate, the following error is logged on one or both of the peers:

```
charon[5319]: 08[IKE] <con100000|1> no trusted RSA public key found for '<identifier>'
```

In that case:

- Check over all the identifier data again to ensure that the values **exactly** match an appropriate certificate field (DN, SAN, etc.)
- If using an ASN.1 DN, ensure the order of DN/subject components **exactly** matches the order reported by the DN field in the Certificate Manager, `strongSwan`, or `openssl`
- Ensure there are no single or double quotes around the identifier value in the GUI
- Ensure the correct **Peer Certificate Authority** is imported and selected

Attempt to initiate the tunnel in both directions manually and compare output (*Manually connect IPsec from the shell*).

35.32 Configuring IPv6 Through A Tunnel Broker Service

A location that does not have access to native IPv6 connectivity may obtain it using a tunnel broker service such as Hurricane Electric. Similarly, a core site with IPv6 can deliver IPv6 connectivity to a remote site by using a VPN or GIF tunnel.

This section provides the process for connecting pfSense® software with Hurricane Electric (Often abbreviated to HE.net or HE) for IPv6 transit. Using HE.net is simple and easy. It allows for multi-tunnel setup, each with a transport `/64` and a routed `/64`. Also included is a routed `/48` to be used with one the tunnels. It is a great way to get a lot of routed IPv6 space to experiment with and learn, all for free.

35.32.1 Sign Up for Service

ICMP echo requests must be allowed to the WAN from the tunnel broker server, or it cannot function. A rule to pass ICMP echo requests from a source of *any* is an acceptable temporary measure. Once the tunnel endpoint for HE.net has been chosen, the rule can be made more specific.

To get started on HE.net, sign up at www.tunnelbroker.net. HE.net will allocate `/64` networks after registering and selecting a regional IPv6 tunnel server.

A summary of the tunnel configuration can be viewed on HE.net's website as seen in Figure *HE.net Tunnel Config Summary*. It contains important information such as:

Tunnel ID
: A number to uniquely identify this tunnel.

Server IPv4 Address
: IP address of the HE.net tunnel server

Client IPv4 Address
: The external IP address of the firewall

Server IPv6 Addresses
: The IPv6 address used inside the tunnel for the remote endpoint.

Client IPv6 Addresses
: The IPv6 address used inside the tunnel for this firewall.

Routed IPv6 Prefixes
: The IPv6 prefixes routed to the firewall over this tunnel. By default, there is at least a `/64` prefix listed, but HE.net can also allocate a `/48` upon request.

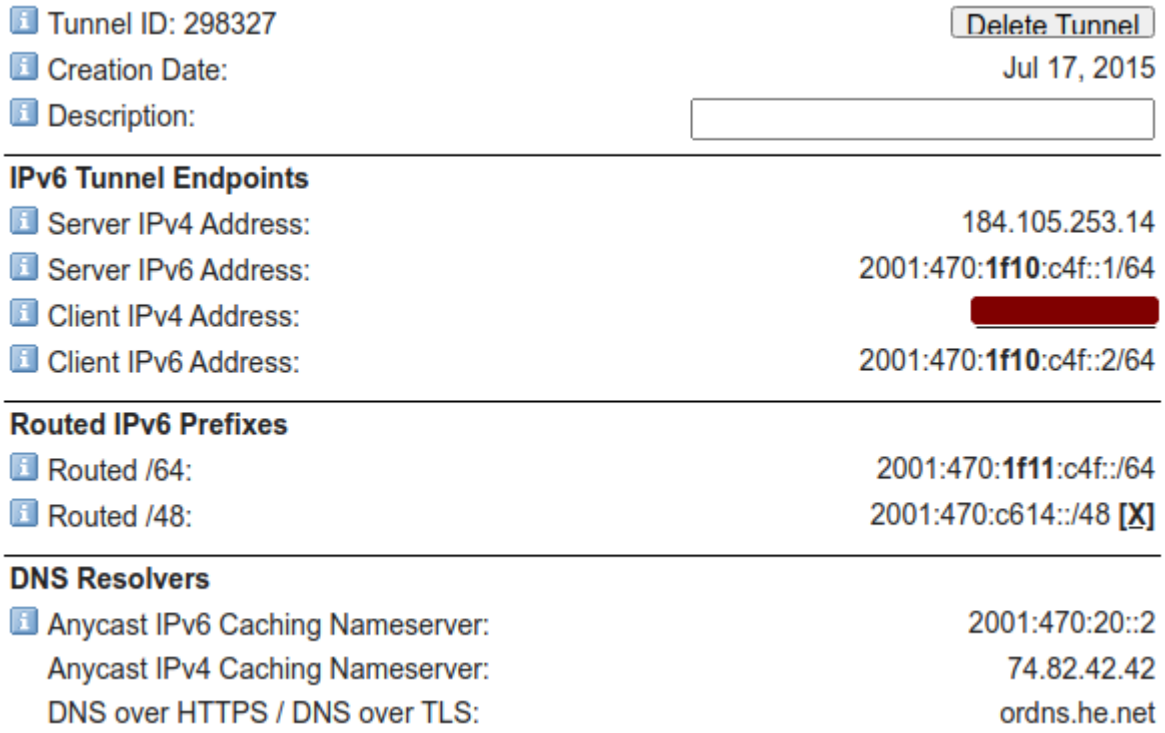


Fig. 61: HE.net Tunnel Config Summary

The **Advanced** tab on the tunnel broker site has two additional notable options:

MTU
: The MTU for packets sent by HE.net over the tunnel.

  If the WAN used for terminating the GIF tunnel is PPPoE or another WAN type with a low MTU, move the slider down as needed. For example, a common MTU for PPPoE lines with a tunnel broker is `1452`.

Update Key
: A key for updating the tunnel address using dynamic DNS mechanisms.

  If the WAN has a dynamic IP address (e.g. DHCP, PPPoE), note this key for later use.

Once the initial setup for the tunnel service is complete, configure the firewall to use the tunnel.

35.32.2 Allow IPv6 Traffic

New installations of pfSense software allow IPv6 traffic by default. Configurations upgraded from older versions may still be set to block IPv6.

To enable IPv6 traffic, perform the following:

- Navigate to **System > Advanced** on the **Networking** tab
- Check **Allow IPv6** if not already checked
- Click **Save**

35.32.3 Allow ICMP

The firewall must allow ICMP echo requests on the WAN address that is terminating the tunnel. This allows HE.net to ensure that the firewall is online and reachable. If the firewall blocks ICMP the tunnel broker may refuse to set up the tunnel to the IPv4 address.

Edit the ICMP rule created earlier, or create a new rule to allow ICMP echo requests from a source IP address of the **Server IPv4 Address** in the tunnel configuration as shown in Figure *Example ICMP Rule*.

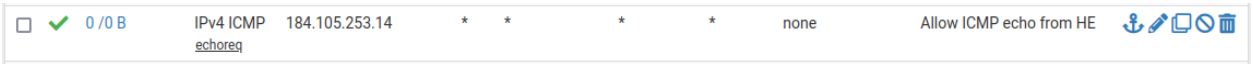


Fig. 62: Example ICMP Rule

35.32.4 Create and Assign the GIF Interface

Next, create the interface for the GIF tunnel. Complete the fields with the corresponding information from the tunnel broker configuration summary.

- Navigate to **Interfaces > Assignments** on the **GIF** tab
- Click  **Add** to add a new entry
- Configure the settings as follows:

  Parent Interface
  The WAN where the tunnel terminates. This would be the WAN which has the **Client IPv4 Address** on the tunnel broker.

  GIF Remote Address
  The **Server IPv4 Address** on the summary.

  GIF Tunnel Local Address
  The **Client IPv6 Address** on the summary.

  GIF Tunnel Remote Address
  The **Server IPv6 Address** on the summary, along the with prefix length (typically `/64`).

  Description
  Text describing the tunnel, such as `HE Tunnel Broker`

- Leave remaining options blank or unchecked
- Click **Save**

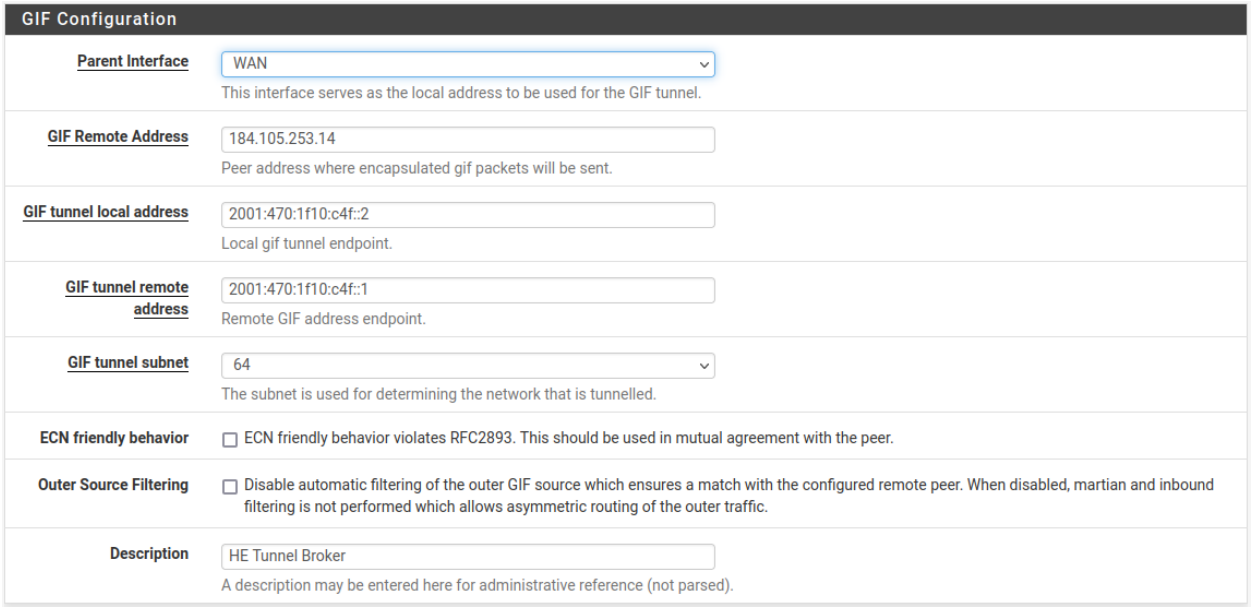


Fig. 63: Example GIF Tunnel

See Figure *Example GIF Tunnel*.

Note: If the WAN containing this tunnel uses a dynamic IP address, see *Updating the Tunnel Endpoint* for information on how to keep the tunnel endpoint IP address updated with HE.net.

Now assign the GIF tunnel as an interface:

- Navigate to **Interfaces > Assignments**, **Interface Assignments** tab
- Select the newly created GIF under **Available Network Ports**
- Click  **Add** to add it as a new interface

35.32.5 Configure the New OPT Interface

The new interface is accessible at **Interfaces > OPTx**, where `x` is a sequential number assigned to the interface.

- Navigate to the new interface configuration page. (**Interfaces > OPTx**)
- Check **Enable Interface**
- Enter a name for the interface in the **Description** field, e.g. `WANv6`
- Click **Save**
- Click **Apply Changes**

Warning: After applying the interface changes the firewall may need to be restarted before the interface configuration will be fully operational. Check **Status > Interfaces** and if the **IPv6 Address** field is missing or empty for the assigned GIF interface, reboot the firewall.

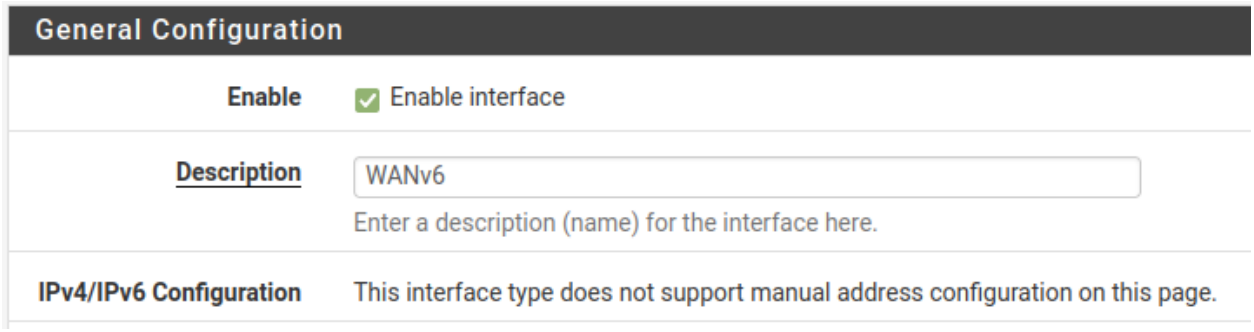


Fig. 64: Example Tunnel Interface

35.32.6 Set up the IPv6 Gateway

The firewall automatically creates a dynamic IPv6 gateway for the assigned GIF interface, but it is not yet marked as default.

- Navigate to **System > Routing**
- Set **Default Gateway IPv6** to the dynamic IPv6 gateway with the same name as the IPv6 WAN created above (e.g. `WANV6_TUNNELV6`)
- Click **Save**
- Click **Apply Changes**

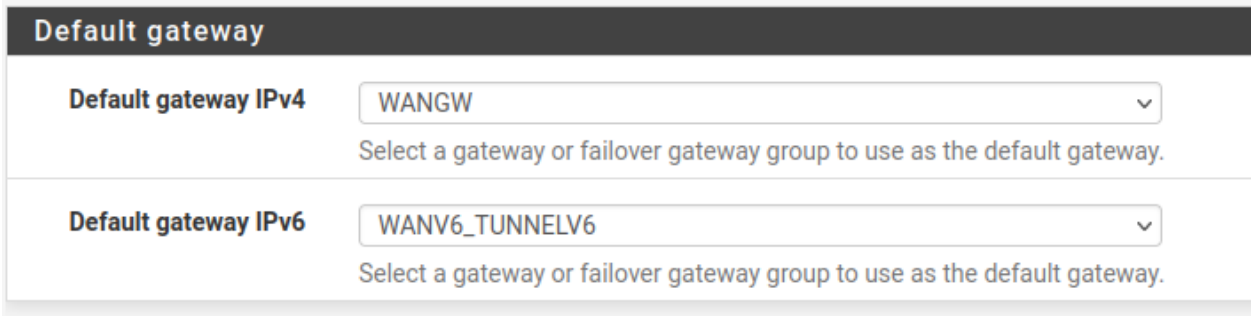


Fig. 65: Example Tunnel Gateway

Navigate to **Status > Gateways** to view the gateway status. The gateway will show as "Online" if the tunnel is operational, as seen in Figure *Example Tunnel Gateway Status*.

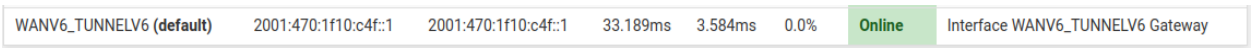


Fig. 66: Example Tunnel Gateway Status

35.32.7 Setup IPv6 DNS

The firewall DNS configuration likely already properly handles DNS queries for `AAAA` records already. If the firewall is configured to use the *DNS Resolver* in resolver mode, which is the default, then nothing needs to be done.

If the firewall is configured to use the DNS Resolver in forwarding mode, or it uses the *DNS Forwarder*, then the best practice is to add the tunnel broker *DNS Servers* under **System > General Setup**.

Enter at least one IPv6 DNS server or use a public DNS service such as Google public IPv6 DNS servers (`2001:4860:4860::8888`, `2001:4860:4860::8844`), Quad9, or CloudFlare.

At this point the firewall itself should have full working IPv6 connectivity.

35.32.8 Setup LAN for IPv6

For clients on LAN to access the internet using IPv6, the LAN must also be configured for IPv6.

The most common method is to set LAN as dual stack IPv4 and IPv6.

- Navigate to **Interfaces > LAN**
- Change the configuration as follows:

  IPv6 Configuration Type
  Static IPv6

  Static IPv6 Configuration
  Enter an IPv6 address from the **Routed /64** in the tunnel broker configuration with a prefix length of `64`.

  For example, use `2001:db8:1111:2222::1` for the LAN IPv6 address if the **Routed /64** is `2001:db8:1111:2222::/64`.
- Click **Save**
- Click **Apply Changes**

Alternately, use a `/64` from within the Routed `/48` prefix.

Setup DHCPv6 and/or Router Advertisements

Router Advertisements and/or DHCPv6 can assign IPv6 addresses to clients automatically. This is covered in detail in *IPv6 Router Advertisements*.

A brief overview is as follows:

- Navigate to **Services > DHCPv6 Server**
- Change to the local interface where IPv6 clients are located
- Check **Enable**
- Enter a range of IPv6 IP addresses inside the new LAN IPv6 prefix
- Click **Save**
- Navigate to **Services > **Router Advertisement**
- Change to the local interface where IPv6 clients are located
- Set the **Mode** to *Managed* (DHCPv6 only) or *Assisted* (DHCPv6+SLAAC)
- Click **Save**

Modes are described in greater detail at *Router Advertisements (Or: "Where is the DHCPv6 gateway option?")*.

To assign IPv6 addresses to LAN clients manually, use the firewall LAN IPv6 address as the gateway with a proper matching prefix length, and pick addresses from within the LAN prefix.

35.32.9 Add Firewall Rules

Now add firewall rules which allow IPv6 traffic from hosts on LAN.

Note: The default LAN ruleset on current installations already contains a rule to pass IPv6, but the best practice is to check and confirm it is present and configured appropriately.

- Navigate to **Firewall > Rules**, **LAN** tab.
- Check the list for an existing IPv6 rule

  If a rule to pass appropriate IPv6 traffic already exists, then no additional action is necessary.

- Click  **Add** to add a new rule to the bottom of the list
- Configure the rule as follows:

  Address Family
  IPv6

  Source
  LAN Net

  Destination
  Any

- Click **Save**
- Click **Apply Changes**

If a local interface contains servers which need to handle public IPv6 requests, add firewall rules on the tab for the IPv6 WAN (the assigned GIF interface) to allow IPv6 traffic to reach the servers on required ports.

35.32.10 Reboot

The best practice is to restart the firewall and then the clients before testing connectivity.

Reboot the firewall first using **Diagnostics > Reboot**. Monitor the boot process for errors and check the interface and gateway status once it is back online. This not only ensures that the firewall is configured properly but will also be configured correctly on subsequent reboots.

Next, reboot a client to test. Some clients may automatically obtain an IPv6 address while they are up and running, some may need their networking services restarted, and others will only check at boot time. Thus, the best practice is to reboot the client to ensure it obtains IPv6 configuration parameters from the firewall.

Note: If a client does not obtain an IPv6 address, check its network settings to see if IPv6 connectivity is enabled and active. Additionally, some clients do not implement certain types of IPv6 configuration. For example, Android clients are not compatible with DHCPv6, but they can use SLAAC.

35.32.11 Try It!

Finally, check for IPv6 connectivity using a site such as test-ipv6.com. An example of the output results of a successful configuration from a client on LAN is in Figure *IPv6 Test Results*.

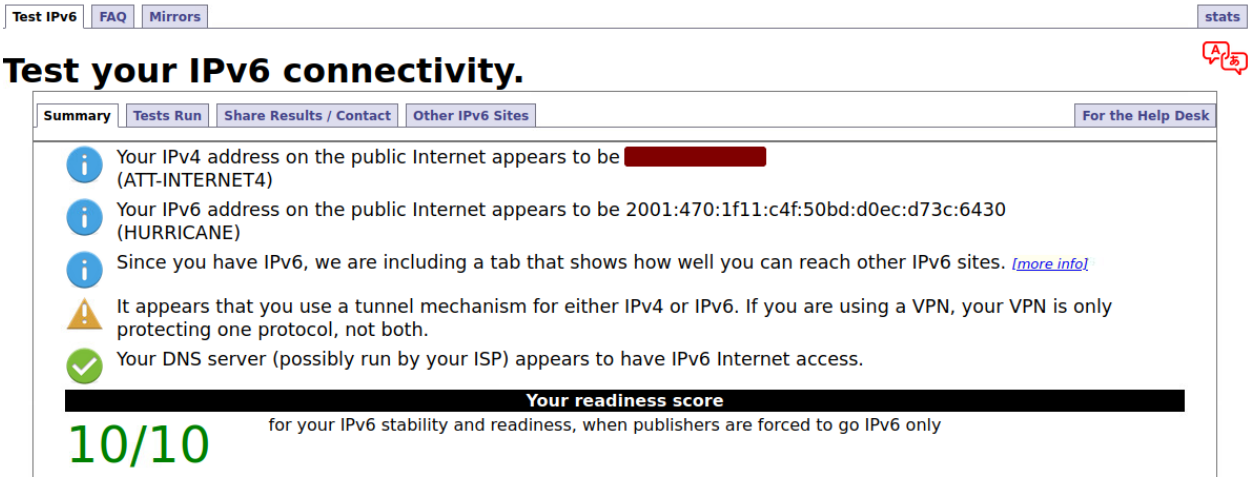


Fig. 67: IPv6 Test Results

35.32.12 Updating the Tunnel Endpoint

The firewall can still use HE.net as a tunnel broker on dynamic WAN types such as DHCP or PPPoE. pfSense software includes a Dynamic DNS type which updates the tunnel endpoint IP address whenever the WAN interface IP changes.

If necessary, configure Dynamic DNS as follows:

- Navigate to **Services > Dynamic DNS**
- Click **Add** to add a new entry
- Configure the entry as follows:

  Service Type
  HE.net Tunnelbroker

  Interface to Monitor
  WAN

  Hostname
  Enter the **Tunnel ID** from the tunnel broker configuration.

  Username
  The **Username** for the tunnel broker site.

  Password
  Enter either the **Password** or **Update Key** for the tunnel broker site.

  Description
  Text describing the entry, e.g. `HE Tunnel`

- Click **Save and Force Update**

If and when the WAN IP address changes, the firewall will automatically update the tunnel broker configuration.

35.33 L2TP/IPsec Remote Access VPN Configuration Example

On current versions of pfSense® software, L2TP/IPsec may be configured for mobile clients, though it is not a desirable configuration.

> **Warning:** Users have reported issues with Windows L2TP/IPsec clients behind NAT. If the clients will be behind NAT, Windows clients will most likely not function.
>
> Consider an IKEv2 implementation instead.

As warned at the start of the chapter, the Windows client, among others, and the strongSwan IPsec daemon are not always compatible, leading to failure in many cases. The best practice is to use another solution such as IKEv2 instead of L2TP/IPsec.

See also:

IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2 contains a walkthrough for configuring IKEv2.

Before configuring the IPsec portion, set up the L2TP server as described in *L2TP Server Configuration* and add users, firewall rules, etc., as covered there.

35.33.1 Setup IPsec

These settings have been tested and found to work with some clients, but other similar settings may function as well. Feel free to try other encryption algorithms, hashes, etc.

Mobile Clients Tab

- Navigate to **VPN > IPsec**, **Mobile Clients** tab in the pfSense software GUI
- Configure the settings as follows:

  Enable IPsec Mobile Client Support
  Checked

  User Authentication
  Local Database (Not used, but the option must have something selected)

  Provide a virtual IP address to clients
  Unchecked

  Provide a list of accessible networks to clients
  Unchecked

- Click **Save**

Phase 1

- Click the **Create Phase1** button at the top if it appears, or edit the existing Mobile IPsec Phase 1
  - If there is no Phase 1, and the **Create Phase1** button does not appear, navigate back to the **Mobile Clients** tab and click it there.
- Configure the settings as follows:

  Key Exchange version
  v1 or *Auto*

  Description
  Text describing the tunnel

  Authentication method
  Mutual PSK

  Negotiation Mode
  Main

  My Identifier
  My IP address

  Encryption algorithm
  AES 256

  Hash algorithm
  SHA1

  DH key group
  14 (2048 bit)

  Note: iOS and other platforms may work with a **DH key group** of *2* instead.

  Lifetime
  `28800`

  Disable Rekey
  Unchecked

  NAT Traversal
  Auto

  Enable DPD
  Checked, set for `10` seconds and `5` retries
- Click **Save**

Phase 2

- Click  **Show Phase 2 Entries** to show the Mobile IPsec Phase 2 list
- Click  **Add P2** to add a new Phase 2 entry if one does not exist, or click  to edit an existing entry
- Configure the settings as follows:

Mode
Transport

Description
Text describing the tunnel

Protocol
ESP

Encryption algorithms
ONLY *AES 128*

Hash algorithms
ONLY *SHA1*

PFS Key Group
off

Lifetime
`3600`

- Click **Save**

Pre-Shared Key

The Pre-Shared Key for the connection, which is common for all clients, must be configured in a special way.

- Navigate to **VPN > IPsec**, **Pre-Shared Keys** tab on pfSense software
- Click  **Add** to add a new PSK
- Configure the settings as follows:

Identifier
`allusers`

Note: The `allusers` name is a special keyword used by pfSense software to configure a wildcard PSK, which is necessary for L2TP/IPsec to function. Do not use any other **Identifier** for this PSK!

Secret Type
PSK

Pre-Shared Key
A password for the user, such as `aaabbbccc` – ideally one a lot longer, more random, and secure!

- Click **Save**
- Click **Apply Changes**

35.33.2 IPsec Firewall Rules

Firewall rules are necessary to pass traffic from the client host over IPsec to establish the L2TP tunnel, and inside L2TP to pass the actual tunneled VPN traffic to systems across the VPN. Adding the L2TP rules was covered in the previous section. To add IPsec rules:

- Navigate to **Firewall > Rules**, **IPsec** tab
- Review the current rules. If there is an "allow all" style rule, then there is no need to add another. Continue to the next task.
- Click  **Add** to add a new rule to the top of the list
- Configure the options as follows:

  Protocol
  any

  Source
  any

  Destination
  any

Note: This does not have to pass all traffic, but must at least pass L2TP (UDP port `1701`) to the WAN IP address of the firewall.

- Click **Save**
- Click **Apply Changes**

35.33.3 DNS Configuration

If DNS servers are supplied to the clients and the Unbound **DNS Resolver** is used, then the subnet chosen for the L2TP clients must be added to its access list.

- Navigate to **Services > DNS Resolver**, **Access Lists** tab
- Click  **Add** to add a new access list
- Enter an **Access List Name**, such as *VPN Users*
- Set **Action** to *Allow*
- Click  **Add Network** under **Networks** to add a new network
- Enter the VPN client subnet into the **Network** box, e.g. `10.3.177.128`
- Choose the proper **CIDR**, e.g. `25`
- Click **Save**
- Click **Apply Changes**

35.33.4 Client Setup

When configuring clients, there are a few points to look for:

- Ensure that the client operating system configuration is set to connect to the proper external address for the VPN.
- It may be necessary to force the VPN type to **L2TP/IPsec** on the client if it has an automatic mode.
- The client authentication type must match what is configured on the L2TP server (e.g. *CHAP*)

35.34 Connecting to L2TP/IPsec from Android

The L2TP/IPsec client on Android has the ability to set a custom identifier, which allows L2TP/IPsec to function with the server on pfSense® software using Pre-Shared Keys. Clients on other operating systems do not allow for this, which makes them incompatible with current versions of pfSense software.

35.34.1 IPsec Setup

The setup is similar to a standard *IPsec Remote Access VPN Example Using IKEv1 with Xauth* setup except that xauth is not used, but rather "**Mutual PSK**", and Phase 2 uses **Transport** mode rather than Tunnel.

Pre-Shared Keys

After the tunnel has been configured, click to the "Pre-Shared Keys" tab in the IPsec settings, and add IPsec keys. A single group key may be used if desired, or make many keys for different users.

That's it for IPsec!

35.34.2 L2TP Setup

To set up L2TP:

- Navigate to VPN > L2TP
- Configure the settings as follows:

  Enable L2TP Server
  Checked

  Interface
  WAN (or the same chosen for IPsec)

  Server Address
  An **unused IP address** in a new subnet, e.g. `x.x.x.2`.

  > **Warning:** This **MUST NOT** overlap any IP address in use on the firewall.

  Remote Address Range
  The starting IP of the clients, e.g. `x.x.x.128`

  Subnet netmask
  The netmask for the client connection, the server IP address should be included in this subnet, e.g. `/24`

Secret
blank

This does not appear to work, at least with the Android version tested.

Encryption Type
CHAP is recommended

L2TP DNS Servers
The LAN IP address of the firewall or another internal DNS server

RADIUS settings
Configure if needed, otherwise leave them at defaults

- Save
- Navigate to the Users tab
- Add L2TP user accounts and passwords
- Navigate to **Firewall > Rules** on the **L2TP VPN** tab
- Add a firewall rule to pass traffic, e.g. from any to any or much more restrictive if preferred.

35.34.3 Android Client Setup

On the phone/tablet/device:

- Navigate to the system settings and VPN settings

  The exact location varies by device and Android version.

- Tap **Add VPN Profile**
- Configure the settings as follows:

  Name
  Enter a name

  Type
  Tap *L2TP/IPsec PSK*

  Server Address
  The WAN IP of the firewall (or the IP address of the interface chosen for IPsec and L2TP)

  L2TP Secret
  blank

  IPsec Identifier
  Enter the identifier for the PSK entered previously, either a per-user or common identifier

  IPsec Pre-Shared Key
  The PSK that goes with the identifier for this user/group

  Advanced Options
  May be used to control which networks will attempt to use the VPN, or specify custom DNS server and domains for this client.

- Tap **Save**
- Tap the newly created VPN entry in the VPN list
- Enter the **username** and **password** from the **L2TP Users**
- Check **Save account information** to save the VPN credentials (not recommended!),

- Tap **Connect**

The connection should then connect and function. If it does not work, check the IPsec logs and the **Status > System Logs**, **VPN**, **L2TP Raw log** to see more specific errors.

35.34.4 Other Thoughts

In theory, Mutual RSA should also work, but so far it has not succeeded in testing. In RSA mode, Phase 1 requires main mode, but otherwise should be OK.

35.35 Migrating an Assigned LAN to LAGG

Only unassigned physical ports can be added to a LAGG, so to move an assigned LAN interface to a LAGG requires some shuffling around.

In this example, the LAN of a device (`igc0`) will be moved into a LAGG with another port on the same device (`igc1`).

35.35.1 Warnings/Precautions

It is best to perform this change from an interface that is not involved, such as a secondary LAN, DMZ, OPT port, perhaps WAN or VPN. Though if the switch is properly configured there should be no loss of connectivity, at least with LACP. Other modes may vary.

35.35.2 Prerequisites/Assumptions

The switch must be properly configured to accommodate the LAGG. This typically means configuring an LACP group and setting ports to use that group. The NICs involved, in this example `igc1` and `igc0`, should be connected to properly configured ports on the switch before starting.

35.35.3 Migrate LAN to a LAGG

- Ensure the second NIC for the LAGG is not assigned (e.g. `igc1` mapped to *LAN2*, *OPTx*, etc.)
  - Check **Interfaces > Assignments** and remove its entry if present.
- Create a new LAGG including only the second NIC
  - Navigate to **Interfaces > Assignments** on the **LAGG** tab
  - Click  to create a new LAGG
  - Click to select the NIC to use with this LAGG (`igc1`)
  - Select the proper LAGG protocol, such as *LACP*
  - Enter a description
  - Click **Save**
- Navigate to **Interfaces > Assignments**, change the assignment of LAN to the newly created LAGG interface (`lagg0`)
- Click **Save**

- Navigate back to the **LAGG** tab
- Click  to edit the LAGG
- Ctrl-click to add in the port that was formerly assigned to LAN (`igc0`) so that both NICs are selected
- Click **Save**

Now both `igc1` and `igc0` are members of a LAGG and that LAGG is the LAN interface with all the existing configuration in place.

35.35.4 VLANs

If any VLANs were in use directly on the interfaces involved, migrate them as follows:

- Add new VLAN tags using the LAGG interface as the parent (**Interfaces > Assignments**, **VLAN** tab)
- Fix the assignments to use the LAGG versions of the tags (**Interfaces > Assignments**)
- Remove the old tags from the physical interface(s) (**Interfaces > Assignments**, **VLAN** tab)

> **Warning: Do not edit the existing tags and change the parent interface**
>
> Changing the parent in place will cause problems with the interface assignments. Always create new tags, switch the assignments, then remove the old tags.

35.36 Accessing a CPE/Modem from Inside the Firewall

Most end-user Customer Premise Equipment (CPE) devices like cable or DSL modems have a web interfaces on a private IP address. Since these sit outside the firewall and do not typically have a public IP address, accessing them isn't as straight forward as it might seem. The firewall is typically assigned a public IP, and sends all outbound traffic upstream to the ISP. The ISP won't route the private subnet back to the modem, leaving it unreachable. This page describes the workaround needed to access the management interface on the modem from the inside of the network.

Note: The CPE management IP address **must** be on a different IP subnet than the internal network. If it is not, attempts to connect to it will never go to the firewall to be routed out to the modem, as hosts on the internal network would try to connect to it on the local network and fail.

35.36.1 Configure a new Interface

A PPPoE WAN is actually assigned to a virtual PPPoE adapter, not the physical port.

- Navigate to **Interfaces > Assignments**
- Set **Available network ports:** to the physical network card for the PPPoE WAN

  For example, if the WAN is *PPPOE0(ix3)*, choose *ix3*.

- Click  **Add** to assign this port as a new OPT interface
- Navigate to **Interfaces > (new OPT interface)**

- Configure the settings as follows:

  Enable
  Checked

  Description
  `ModemAccess` or a similar useful name.

  IPv4 Configuration Type
  Static

  IPv4 Address
  Configure an IP address in the same subnet as the modem, such as `192.168.1.5/24`.

  IPv4 Upstream Gateway
  None

  Do not set a gateway.

- Click **Save**
- Click **Apply Changes**

35.36.2 Configure NAT

Now NAT needs to be configured to translate traffic destined to the modem to the new interface. This is necessary for the modem to see the traffic sourced from an IP address on its local subnet. Without this NAT, it would be necessary to configure a route on the modem so it knows how to reach the internal subnet. With some modems this is not possible, and in most cases it's easier to NAT the traffic so that routing isn't a concern.

To add the NAT:

- Navigate to **Firewall > NAT**, **Outbound** tab.
- Switch to **Hybrid Outbound NAT** and click **Save**
- Click  to add a new Outbound NAT rule
- Configure the settings as follows:

  Interface
  ModemAccess

  Source
  Network, enter the LAN subnet

  Destination
  The IP subnet of the modem

  Translation
  Interface Address

- Click **Save**
- Click **Apply changes**

It should now be possible to access the modem from LAN.

35.37 Configuring Multi-WAN for IPv6

Multi-WAN can be utilized with IPv6 provided that the firewall is connected to multiple ISPs or tunnels with static addresses.

See also:

See *Configuring IPv6 Through A Tunnel Broker Service* for help setting up a tunnel.

Gateway Groups work the same for IPv6 as they do for IPv4, but address families cannot be mixed within a group. A group must contain either *only* IPv4 gateways, or *only* IPv6 gateways.

Throughout this section "Second WAN" refers to the second or additional interface with IPv6 connectivity. It can be an actual interface that has native connectivity, or a tunnel interface when using a tunnel broker.

35.37.1 Caveats

In most cases, NAT is not used with IPv6 in any capacity as everything is routed. That is great for connectivity and for businesses or locations that can afford Provider Independent (PI) address space and a BGP peering, but it doesn't work in practice for small business and home users.

Network Prefix Translation (NPt) allows one subnet to be used for LAN which has full connectivity via its native WAN, but also has translated connectivity on the additional WANs so it appears to originate there. While not true connectivity for the LAN subnet via the alternate paths, it is better than no connectivity at all if the primary WAN is down.

Warning: This does not work for dynamic IPv6 types where the subnet is not static, such as DHCP6-PD.

35.37.2 Requirements

To set up Multi-WAN for IPv6 the firewall must have:

- IPv6 connectivity with static addresses on two or more WANs
- Gateways added to **System > Routing** for both IPv6 WANs, and confirmed connectivity on both.
- A routed `/64` from each provider/path
- LAN using a static routed `/64` or similar

35.37.3 Setup

The setup for IPv6 Multi-WAN is very close to the setup for IPv4. The main difference is that it uses NPt instead of NAT.

First, under **System > Routing** on the **Gateway Groups** tab, add Gateway Groups for the IPv6 gateways, with the tier members setup as desired. This works identically to IPv4.

Next, navigate to **System > General** and set one IPv6 DNS server set for each IPv6 WAN, also identically to IPv4.

Now add an NPt entry under **Firewall > NAT** on the **NPt** tab, using the following settings:

Interface
: Secondary WAN (or tunnel if using a broker)

Internal IPv6 Prefix
: The *LAN* IPv6 subnet

Destination IPv6 Prefix
The **routed IPv6 subnet** for the *secondary* WAN or tunnel

> **Note:** This is **not** the `/64` of the WAN interface itself – it is the `/64` routed to the firewall on that WAN by the upstream.

What this does is akin to 1:1 NAT for IPv4, but for the entire subnet. As traffic leaves the second WAN, if it is coming from the LAN subnet, it will be translated to the equivalent IP address in the other subnet.

For example if the firewall has `2001:xxx:yyy::/64` on LAN, and `2001:aaa:bbb::/64` on the second WAN, then `2001:xxx:yyy::5` would appear as `2001:aaa:bbb::5` if the traffic goes out the second WAN. For more information on NPt, see *IPv6 Network Prefix Translation (NPt)*.

As with IPv4, the Gateway Groups must be used on LAN firewall rules. Edit the LAN rules for IPv6 traffic and set them use the gateway group, making sure to have rules for directly connected subnets/VPNs without a gateway set so they are not policy routed.

35.37.4 Alternate Tactics

Some users prefer to configure LAN with a "private" IPv6 subnet from the `fc00::/7` space and setup NPt for both WANs.

35.38 Configuring NAT for a VoIP PBX

For fully functional VoIP with inbound and outbound audio from a local PBX, many use cases require the configuration of several components:

1. 1:1 NAT or port forward entries with firewall rules to direct inbound traffic
2. 1:1 NAT or Manual Outbound NAT with a rule at the top set to perform static port NAT on outgoing traffic from the PBX
3. Ensure the PBX itself is set properly for NAT with the correct external IP address and local subnets defined.

35.38.1 Aliases to make it easy

The best practice for ease of administration is to create firewall aliases which make the rules easier to create and maintain:

Start by navigating to **Firewall > Aliases** and then create the following aliases:

- Host alias for the PBX itself, named `PBX`, containing the local IP address of the PBX.
- Network or Host alias named `SIP_Trunks` for the upstream SIP trunk addresses, if known.

  If the `SIP_Trunk` address/network is dynamic or not known, do not make an alias and leave these values set to *any*.
- Port alias named `PBX_Ports` containing all port numbers the trunk requires for SIP, RTP, and other control ports.

  It many cases, the list of ports is `5060` and `10000:20000`, but it can vary between different providers and PBX implementations.

35.38.2 Port Forwards

Create a port forward:

- Navigate to **Firewall > NAT**, **Port Forwards** tab
- Create a new entry and configure it as follows:

  Interface
  WAN

  Protocol
  UDP (or *TCP/UDP* if needed)

  Source
  Type *Address or Alias*: *SIP_Trunks*

  If the SIP trunk IP addresses are not known, use *Any*.

  Source Port
  Any/Any

  Destination
  WAN address or external VIP for the PBX

  Destination Port
  PBX_Ports

  Redirect target IP
  PBX

  Redirect target port
  PBX_Ports

- Click **Save**
- Click **Apply Changes**

35.38.3 Outbound NAT

Setup *Hybrid Outbound NAT*.

- Navigate to **Firewall > NAT**, **Outbound** tab
- Select *Hybrid Outbound NAT*
- Click **Save**
- Click  to create a new rule at the top of the list.
- Configure the rule as follows:

  Interface
  WAN

  Protocol
  UDP

  Source
  Network, *PBX*

  Source Port
  blank

Destination
Network or Alias, *SIP_Trunks* – Or *Any* for the type if the SIP trunk IP addresses are not known

Destination Port
PBX_Ports (or leave blank)

Translation
WAN address if using the WAN IP address, or the external VIP for the PBX

Port
blank

Static Port
CHECKED

- Click **Save**
- Click **Apply Changes**

35.38.4 Reset States

After making the changes to NAT rules, the states for the PBX must be reset.

- Navigate to **Diagnostics > States**
- Enter the IP address of the PBX
- Click **Filter**
- Click **Kill States**
- Click **OK** to confirm killing the states

Once the PBX re-registers, test inbound and outbound calls to confirm both inbound and outbound audio work as expected.

35.39 Configuring NAT for VoIP Phones

The default settings for pfSense® software may not be optimal for certain VoIP environments. The default settings handle the majority of scenarios, but depending on the specifics of a particular VoIP configuration, changes may be necessary to make successful VoIP calls.

The following sections guide through common settings which can help local handsets to work with a remote PBX.

See also:

If the PBX is local and trying to communicate with a remote SIP trunk, see *Configuring NAT for a VoIP PBX* instead.

35.39.1 Disable source port rewriting

By default, pfSense software rewrites the source port on all outbound traffic. This is necessary for proper NAT in some environments, such as having multiple SIP phones behind a single public IP address which all register to a single external PBX. With some providers, rewriting the source port of RTP can cause one way audio. In that case, setup manual outbound NAT and *Static Port* on all UDP traffic potentially with the exclusion of UDP `5060`.

Performing static port NAT on UDP `5060` traffic by default is not desirable because it breaks more scenarios than it helps in current environments. However, in cases where a PBX requires static port on UDP `5060`, configuring outbound NAT to perform static port NAT for `udp/5060` will allow it to function. This can be done using Hybrid outbound NAT and a phone-specific rule or by using manual outbound NAT.

35.39.2 Set Conservative state table optimization

The default UDP timeouts in PF are too low for some VoIP services. If phones work most of the time, but randomly disconnect, set **Firewall Optimization Options** to *Conservative* under **System > Advanced**, **Firewall/NAT** tab.

A keep-alive or re-registration timer on the phone set for 20-30 seconds can also help, and is often a better solution.

35.39.3 Use the siproxd package

The *Siproxd package* is only a good fit for deployments with local phones and a remote PBX where phones cannot register or make calls with rewritten source ports. In this very specific circumstance, the siproxd package enables multiple phones to connect to a single outside server with a static source port of `5060`.

> **Warning:** Do not use siproxd if the PBX is local. Only use this package if the upstream PBX strictly requires all phones to have a source port of `5060`.

35.39.4 Disable scrub

Certain very rare circumstances may require disabling scrubbing under **System > Advanced**, **Firewall/NAT** tab. In most cases this should be left at the default setting (*unchecked*). Only change this setting when necessary. Some phones send malformed packets that will be silently dropped without scrub active. For example, unfragmented packets which claim to be fragmented.

35.40 Configuring NAT64 for IPv6-only Clients

This recipe covers the configuration of *NAT64* to allow local IPv6-only clients to contact remote resources only reachable over IPv4.

This entire process is transparent to the client software, such as a web browser. It will act as though it is communicating directly to IPv6 hosts even when using hostnames for sites with no IPv6 connectivity.

See also:

- *NAT64*
- *NAT64 Firewall Rule Configuration*
- *PREF64*
- *DNS64*

35.40.1 Base Configuration

This guide assumes the firewall already has working IPv6 and IPv4 connectivity and at least one local interface which is configured for IPv6 clients.

Consult the *NAT64 Requirements* for other prerequisite configuration items.

35.40.2 Configure NAT64 Firewall Rules

The first step is to configure a *NAT64 firewall rule* to perform the IPv6 to IPv4 translation:

- Navigate to **Firewall > Rules**
- Navigate to the tab for the IPv6 only client interface
- Click  to add a new firewall rule
- Configure the rule as follows

  Action
  Pass

  Address Family
  IPv6

  Enable NAT64
  Checked

  Protocol
  Any

  Address Family Translation Source
  Automatic (default)

  Destination
  Leave at the current value which should be the default NAT64 prefix, `69:ff9b::/96`.

  Description
  `Default Allow IPv6 to IPv4 via NAT64`

- Click **Save**
- Drag and move the rule so it is *above* any other default IPv6 allow rule

  If this rule is below another rule which passes IPv6 to any destination, it would never be matched.

- Click **Save**
- Click **Apply Changes**

The rule should look similar to the following:

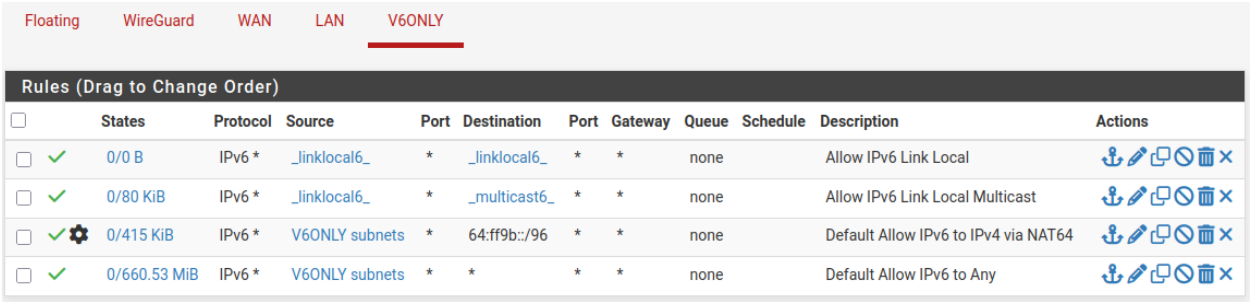


Fig. 68: IPv6 Client Interface Firewall Rules with NAT64

35.40.3 Configure DHCPv6 for DNS

NAT64 requires working DNS, so clients must have a way to be informed about DNS server addresses. It is possible to deliver this via router advertisements (RDNSS/DNSSL) but not all clients implement those methods, so configuring DHCPv6 for DNS is a safe backup.

- Navigate to **Services > DHCPv6 Server**
- Navigate to the tab for the IPv6 only client interface
- Ensure the service is **Enabled** and has **Enable DNS** checked
- Manually change the **DNS Servers** list if the default entries are not acceptable.
- Click **Save**
- Click **Apply Changes**

35.40.4 Configure PREF64

Router advertisements can use *PREF64* to announce the NAT64 prefix to IPv6 clients for automatic discovery:

- Navigate to **Services > Router Advertisement**
- Navigate to the tab for the IPv6 only client interface
- Ensure the **Router Mode** is set to one of *Managed*, *Assisted*, or *Stateless DHCP*.
- Set **NAT64 Prefix** to `64:ff9b::/96`
- Check **Enable DNS** to enable RDNSS/DNSSL
- Check **Mirror DHCPv6** or manually enter DNS servers
- Click **Save**

35.40.5 Configure DNS64

DNS64 includes IPv6-mapped IPv4 addresses in DNS responses so IPv6-only clients can contact these hosts via NAT64. This behavior is not active by default and must be enabled separately.

- Navigate to **Services > DNS Resolver**
- Navigate to the **Advanced Settings** tab
- Check **Enable DNS64 (RFC 6147)**
- Set **DNS64 Prefix** to `64:ff9b::/96`
- Click **Save**
- Click **Apply Changes**

35.40.6 Configure IPv4 DHCP (Optional)

If the segment does not have IPv4 DHCP enabled, it can remain disabled.

However, if the segment has local/private IPv4 configured, such as for a VPN, but no external IPv4 connectivity, then clients can be sent a DHCP option which informs them to prefer IPv6 for external communication.

Note: This guide assumes Kea is in use for DHCP service

- Navigate to **Services > DHCP Server**
- Navigate to the tab for the IPv6 only client interface
- Locate the **Custom Configuration** section near the bottom of the page
- Set **JSON Configuration** to the following:

```
{
  "option-data": [
    {
      "name": "v6-only-preferred",
      "data": "3600"
    }
  ]
}
```

This enables DHCP option `108` for IPv4 DHCP clients on that segment and clients which respect this option should only contact IPv6 servers when possible.

35.40.7 Finish Up / Testing

IPv6 clients on the configured interface should now be able to contact IPv4 only Internet sites and have their traffic translated automatically.

Connectivity can be tested manually as described in *Contacting Remote IPv4 Hosts*.

A popular testing technique is to attempt loading a well-known Internet site which only has IPv4 connectivity, such as `https://github.com`.

Contacting IPv4 addresses directly using IPv4 notation may be possible depending on the presence of a customer-side translator (CLAT), but this is completely dependent on the client and is not handled at the firewall level.

35.41 Exporting NetFlow with softflowd

softflowd is a NetFlow collector that can be deployed on pfSense® software.

Tip: This recipe requires an add-on package. pfSense Plus software contains a native solution which is easier to configure and more efficient: *Firewall Packet Flow Data*.

35.41.1 Installing softflowd



There is a package available under **System > Packages** on the **Available Packages** tab. Find it in the list, click at the end of its row, and confirm the installation.

35.41.2 Configuring and Launching softflowd

Once the package has been installed, visit **Services > softflowd** to configure the service.

Interface
Ctrl-click to select all the interfaces upon which the daemon will gather NetFlow data.

Host
The target NetFlow server which will receive flow data.

Port
The port on the **Host** which is listening for NetFlow data.

Max Flows
The number of flows to track before older flows expire.

NetFlow Version
The desired version of the NetFlow protocol.

See also:

See NetFlow Versions on Wikipedia for more information.

35.41.3 Controlling softflowd from the Command Line

To view statistics about the running softflowd process, run the following command, replacing `igc0` with the actual network interface to query:

```
: softflowctl -c /var/run/softflowd.igc0.ctl statistics
```

To expire all flows and force an update to be sent to the NetFlow server, run the following command, replacing `igc0` with the actual network interface to control:

```
: softflowctl -c /var/run/softflowd.igc0.ctl expire-all
```

35.41.4 Known issues

See also:

The pfSense software issue tracker contains a list of known issues with this package.

35.41.5 Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

35.42 OpenVPN Site-to-Site Configuration Example with SSL/TLS

A site-to-site connection using **SSL/TLS** in client/server mode works for connecting one or more remote sites and is especially convenient for managing numerous remote sites connecting back to a central site in a hub-and-spoke fashion.

35.42.1 Example Configuration Overview

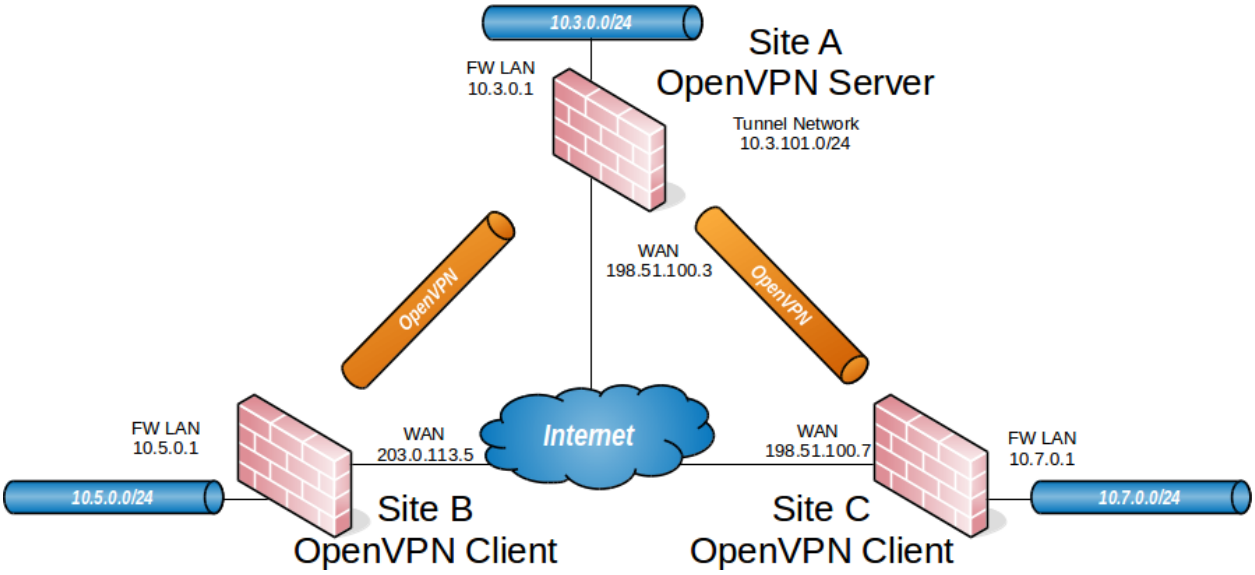


Fig. 69: OpenVPN Example Site-to-Site SSL/TLS Network

When configuring a site-to-site OpenVPN connection using SSL/TLS one firewall will be the server and the others will be clients.

Tip: Usually the main location will be the server and the remote offices will act as clients, though if one location has a static IP address and more bandwidth than the main office that may be a more desirable location for the server.

This style of VPN requires a dedicated subnet for the OpenVPN interconnection between networks in addition to the subnets on both ends. Figure *OpenVPN Example Site-to-Site SSL/TLS Network* shows a depiction of this layout, using `10.3.101.0/24` as the IPv4 Tunnel Network for the VPN. This can be any valid IPv4 subnet so long as it does not overlap another subnet currently in use on any of the connected networks.

OpenVPN allocates IP addresses the same way it does for remote access clients. When using a **Topology** style of *subnet*, each client obtains one IP address in a common subnet. When using a **Topology** style of *net30*, each connecting client gets a /30 subnet to interconnect itself with the server.

See also:

The *subnet* topology style uses address space efficiently and has very few quirks with its behavior compared to the alternatives. See *Topology* for more details.

The following sections describe how to configure the server and client sides of the connection.

35.42.2 Summary of Requirements

This style of VPN requires several items, all of which are covered in this recipe:

On the Server at the main site:

- A certificate structure including a Certificate authority, server certificate, and one or more client certificate(s)
- An OpenVPN Server instance
- OpenVPN Client Specific Override entries for **all** clients

On each remote site:

- Imported CA and client certificate
- OpenVPN Client instance

35.42.3 Example Configuration Settings

Table 12: OpenVPN Endpoint Settings - Site A - Server

Site A - Server	
Name	Austin Office
WAN Address	198.51.100.3
LAN Subnet	10.3.0.0/24
LAN Address	10.3.0.1
CA Name	S2SCA
Cert CN	serverA
Tunnel Net	10.3.101.0/24

Table 13: OpenVPN Endpoint Settings - Site B - Client

Site B - Client	
Name	London Office
Cert CN	clientB
WAN Address	203.0.113.5
LAN Subnet	10.5.0.0/24
LAN Address	10.5.0.1

Table 14: OpenVPN Endpoint Settings - Site C - Client

Site C - Client	
Name	Colorado Office
Cert CN	clientC
WAN Address	198.51.100.7
LAN Subnet	10.7.0.0/24
LAN Address	10.7.0.1

35.42.4 Configuring SSL/TLS Server Side

The server **requires** two items for **each** network reachable through an OpenVPN client:

- An entry in the **IPv4 Remote network(s)** field on the **server instance**.

  This configures a `route` to tell the operating system that OpenVPN knows about a remote network. This makes the OS hand packets for the destination network over to OpenVPN.

- An entry in the **IPv4 Remote network(s)** field of an **OpenVPN Client Specific Override**.

  This configures an internal route (`iroute`) to tell OpenVPN how to route a subnet to a specific client certificate. This allows OpenVPN to deliver packets for the destination network to the appropriate connected client without hardcoding its address anywhere.

More detail on these will follow in the example.

> **Warning:** This setup assumes the server will not be using DCO, as that requires a different routing configuration. For a similar recipe that works with DCO, see *OpenVPN Site-to-Site Configuration Example with SSL/TLS and DCO*.

See also:

- *Client Specific Overrides*
- *Troubleshooting OpenVPN Internal Routing (iroute)*
- *Tunnel Settings*

Create Certificate Structure

The first step is to create a certificate structure for this VPN.

This example uses the names listed in *Example Configuration Settings* – The CA is named `S2SCA`, the Server CN is named `serverA`, and the clients are `clientB` and `clientC`.

See also:

Certificate Management

Create a Certificate Authority

Create a CA unique to this VPN:

- Navigate to **System > Certificates**, **Authorities** tab
- Click  **Add** to create a new a CA
- Enter the settings as follows:

  Descriptive Name
  `S2SCA`

  Method
  Create an internal Certificate Authority

  Randomize Serial
  Checked

Key Type
RSA, *2048* (or higher)

Digest Algorithm
sha256 (or higher)

Lifetime (days)
`3650`

Common Name
`S2SCA`

Subject Component Fields
The remaining fields are optional, but can be set to reflect the location of the CA.

- Click **Save**

Create a Server Certificate

Create a server certificate signed by the VPN CA:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add/Sign** to create a new certificate
- Enter the settings as follows:

Method
Create an internal Certificate

Descriptive Name
`serverA`

Certificate Authority
S2SCA

Key Type
RSA, *2048* (or higher)

Digest Algorithm
sha256 (or higher)

Lifetime (days)
`398`

Note: Some current operating systems and software limit server certificates to a maximum lifetime of `398` days for security reasons. Clients on these platforms may reject a server certificate with a longer lifetime.

Common Name
`serverA`

Subject Component Fields
The fields contain data copied from the CA and are optional, but can be set to reflect the location of the server.

Certificate Type
Server Certificate

> **Warning:** This setting is critical, **do not forget to set this value**!

Alternative Names
Optional extra entries, if needed, which specify alternate ways to identify the server. This can be left blank if the certificate will only be used by OpenVPN. Otherwise, add fields with additional information such as alternate hostnames, static IP addresses, and so on which are relevant to this server.

- Click **Save**

Create User Certificates

Create user certificates for each remote site signed by the VPN CA.

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add/Sign** to create a new certificate
- Enter the settings as follows:

Method
Create an internal Certificate

Descriptive Name
`clientB`

Certificate Authority
S2SCA

Key Type
RSA, *2048* (or higher)

Digest Algorithm
sha256 (or higher)

Lifetime (days)
`3650`

Common Name
`clientB`

Subject Component Fields
The fields contain data copied from the CA and are optional, but can be set to reflect the location of the client.

Certificate Type
User Certificate

> **Warning:** This setting is critical, **do not forget to set this value**!

Alternative Names
Optional extra entries which specify alternate ways to identify the client. These can be left blank if the certificate will only be used by OpenVPN. Otherwise, add fields with additional information such as alternate hostnames, static IP addresses, and so on which are relevant to this client.

- Click **Save**

Repeat this process for every client (e.g. `clientC` and any future clients).

Export Certificates

The next task is to export the certificates and keys which the client requires when connecting to the OpenVPN server.

- Navigate to **System > Certificates**, **Authorities** tab
- Click  on the row for the CA to export its certificate
- Navigate to **System > Certificates**, **Certificates** tab
- Click  on the row for each client certificate to export the certificates
- Click  on the row for each client certificate to export the private key for the client certificates.

Warning: **Do not** export the CA key, server certificate, or server key! The client does not need these files and copying them unnecessarily significantly weakens the security of the VPN.

Configure the OpenVPN Server Instance

- Navigate to **VPN > OpenVPN**, **Servers** tab
- Click  **Add** to create a new server
- Fill in the fields as described below, with everything else left at defaults.

  Use values appropriate for this network, or the defaults if unsure.

  See also:

  See *Server Configuration Options* for details on each of these options.

  Description
  Enter text to describe the connection, e.g. `Site-to-Site VPN`.

  Server Mode
  Peer to Peer (SSL/TLS)

  DCO (Plus Only)
  This example uses routing that only functions when DCO is **disabled**. To use DCO on the server, check this box to activate the feature and read the DCO section of the documentation for specifics on how routing differs.

  See also:

  - For additional information on DCO in general, see *OpenVPN Data Channel Offload (DCO)*.
  - For a similar recipe that works with DCO, see *OpenVPN Site-to-Site Configuration Example with SSL/TLS and DCO*.

  Device Mode
  tun

Protocol
UDP on IPv4 only

Interface
WAN

Local Port
`1194`

TLS Configuration
Check the **Use a TLS Key** box to enable TLS authentication which provides protection for the tunnel control channel.

Leave **Automatically generate a TLS Key** checked, so the firewall will generate a new key automatically the first time this entry is saved.

Peer Certificate Authority
Select the CA created at the beginning of this process (`S2SCA`)

Peer Certificate Revocation List
Select a CRL for the CA, if one exists.

Server Certificate
Select the server certificate created at the beginning of this process (`serverA`)

IPv4 Tunnel Network
Enter the chosen tunnel network, `10.3.101.0/24`

IPv4 Local Network(s)
Enter the LAN subnets for all sites including the server: `10.3.0.0/24, 10.5.0.0/24, 10.7.0.0/24`

Note: If there are more networks on the server side that clients need to reach, such as networks reachable via static routes, other VPNs, and so on, add them as additional entries in the **IPv4 Local Network** box.

IPv4 Remote Network(s)
Enter **only** the client LAN subnets: `10.5.0.0/24, 10.7.0.0/24`

Inter-client communication
Check if the client sites will communicate between each other. Leave unchecked if the remote clients only communicate with the server network(s).

Inactive
`0` to disable disconnecting idle clients, so that site-to-site connections can stay up indefinitely.

- Click Save.
- Click  to edit the new server instance
- Find the **TLS Key** box
- Select all the text inside
- Copy the text to the clipboard
- Save this to a file or paste it into a text editor such as Notepad temporarily

Create Client-Specific Overrides

Now add **Client Specific Overrides** for each client site. These tie a client subnet to a particular certificate so that OpenVPN can properly route a subnet to the correct site.

- Navigate to **VPN > OpenVPN**, **Client Specific Overrides** tab
- Click  to add a new override
- Fill in the fields on this screen as follows:

  Description
  A name for the override, such as the common name of the client (e.g. `clientB`).

  Common Name
  Enter the CN of the first client site. In this example, that is `clientB`.

  Server List
  Select the server instance configured previously.

  IPv4 Remote Network/s
  The clientB LAN subnet, `10.5.0.0/24`.

  Note: This field sets up the internal route (`iroute`) for OpenVPN.

- Click **Save**

Add an override for the second site, adjusting the **Description**. **Common Name** and **IPv4 Remote Network** to match. In the example for site C, these values would be `clientC` and `10.7.0.0/24` respectively.

See also:

- *Client Specific Overrides*
- *Troubleshooting OpenVPN Internal Routing (iroute)*
- *Tunnel Settings*

Firewall Rules

External Traffic (WAN)

Next, add a firewall rule for the WAN interface which allows access to the OpenVPN server.

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  UDP

  Source
  any (since multiple sites must connect)

> **Tip:** For extra security, if the clients have static IP addresses, create an alias containing these addresses, then set it as the source on this rule.

Destination
WAN Address

Destination port
`1194`

Description
`OpenVPN Multi-Site VPN`

- Click **Save**
- Click **Apply Changes**

Tunneled Traffic

Now add a rule to the **OpenVPN** tab to pass traffic over the VPN from the Client-side LAN to the Server-side LAN. This can be an "Allow all" style rule or a set of stricter rules. This example allows all traffic using this rule:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

Protocol
any

Source
any

> **Tip:** For extra security, create an alias containing only the remote hosts or subnets which must initiate contact with hosts on the sever LAN, then use that alias as the source on this rule.

Destination
any

> **Tip:** For extra security, create an alias containing only the local hosts or subnets on the server LAN which must accept connections from remote hosts across the VPN, then use that alias as the destination on this rule.

Description
`Allow all on OpenVPN`

- Click **Save**
- Click **Apply Changes**

That completes the server setup, next, now move on to configure the clients.

35.42.5 Configuring SSL/TLS Client Side

Import CA and Certificate

On the client, import the CA certificate along with the client certificate and key for that site. This is the same CA and client certificate created earlier in this document.

See also:

Certificate Management

Import these items at **System > Certificates**.

First import the CA:

- Navigate to **System > Certificates**, **Authorities** tab
- Click  **Add** to create a new certificate authority
- Enter the settings as follows:

  Descriptive Name
  `S2SCA`

  Method
  Import an existing Certificate Authority

  Certificate Data
  Open the CA certificate file in a text editor on the client PC, select all the text, and copy it to the clipboard. Then paste it into this field.

- Click **Save**

Next, import the client certificate:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add** to create a new certificate
- Enter the settings as follows:

  Method
  Import an existing Certificate

  Descriptive Name
  `clientB VPN Certificate`

  Certificate Type
  X.509 (PEM)

  Certificate Data
  Open the client certificate file in a text editor on the client PC, select all the text, and copy it to the clipboard. Then paste it into this field.

  Private Key Data
  Open the client certificate private key in a text editor on the client PC, select all the text, and copy it to the clipboard. Then paste it into this field.

- Click **Save**

Repeat these steps on each client firewall.

Configure the OpenVPN Client Instance

After importing the certificates, create the OpenVPN client:

- Navigate to **VPN > OpenVPN**, **Clients** tab
- Click  **Add** to create a new client
- Fill in the fields as follows, with everything else left at defaults:

  See also:

  See *Client Configuration Options* for details on each of these options.

  Description
  Text to describe the connection (e.g. `Site A VPN`)

  Server Mode
  Peer to Peer (SSL/TLS)

  DCO (Plus Only)
  Check this box to activate the *OpenVPN Data Channel Offload (DCO)* feature for the client if desired.

  See also:

  See *OpenVPN Data Channel Offload (DCO)* for additional information.

  Device Mode
  tun

  Protocol
  UDP on IPv4 only

  Interface
  WAN

  Server host or address
  The public IP address or hostname of the OpenVPN server (`198.51.100.3` in this example)

  Server Port
  `1194`

  Use a TLS Key
  Checked

  Automatically generate a TLS key
  Unchecked

  TLS Key
  Paste in the TLS key copied from the server instance

  Peer Certificate Authority
  The CA imported at the beginning of this process

  Client Certificate
  The client certificate imported at the beginning of this process

- Click Save

Note: With SSL/TLS server/client configurations such as this example, routes and other configuration options are automatically pushed from the server and thus not present in the client configuration. If the client side must reach additional networks, configure them in the **server** settings or a client-specific override as **Local Networks**.

Firewall Rules

This next step is optional depending on whether hosts on the server network or other client sites need to initiate contact with hosts on the client network. If the other sites do not need to initiate contact with this client, then no action is necessary.

If the other sites will initiate contact, then this traffic requires a firewall rule on the **OpenVPN** tab on the *client* firewall to allow traffic from other VPN sites to reach the Client-side LAN. An "Allow all" style rule is OK in some cases, but a set of stricter rules is the best practice.

This example allows all traffic:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  any

  Source
  any

  Tip: For extra security, create an alias containing only the remote hosts or subnets which must initiate contact with hosts on the client LAN, then use that alias as the source on this rule.

  Destination
  any

  Tip: For extra security, create an alias containing only the local hosts or subnets on the client LAN which must accept connections from remote hosts across the VPN, then use that alias as the destination on this rule.

  Description
  `Allow all on OpenVPN`
- Click **Save**
- Click **Apply Changes**

35.42.6 Testing the Connection

The configuration is now complete. The OpenVPN client instance automatically starts when created, so it should already be attempting to connect at this point and if the configuration is correct, it will be connected.

Try to ping across to the remote end LAN to verify connectivity.

See also:

Troubleshooting OpenVPN.

35.43 OpenVPN Site-to-Site Configuration Example with SSL/TLS and DCO

A site-to-site server using OpenVPN and Data Channel Offload (DCO) in Peer-to-Peer SSL/TLS mode can connect one remote site and can work with either static or dynamic routing. OpenVPN DCO allows for huge performance gains when processing encrypted OpenVPN data by reducing the amount of context switching that happens for each packet.

Note: DCO is available exclusively in pfSense Plus software.

See also:

- *OpenVPN Site-to-Site Configuration Example with SSL/TLS (without DCO)*, upon which this recipe is based.
- *OpenVPN Data Channel Offload (DCO)*
- *OpenVPN*

35.43.1 Warnings and Limitations

Note: Some OpenVPN features and use cases are not compatible with DCO. See *Limitations* for a list of known DCO limitations.

- This configuration does not require assigning the OpenVPN instance on either side as an interface in pfSense Plus software. While assigning the interface may work, this recipe assumes the interface is **not** assigned on the server or client.
- When DCO is enabled on pfSense Plus software, OpenVPN handles the routing in a much different manner at the operating system level. DCO is not compatible with internal routing in OpenVPN (`iroute`) which limits a server to one client. However, it is compatible with routing traffic on the interface using the operating system routing table, which allows OpenVPN with DCO to utilize dynamic routing protocols such as BGP or OSPF to exchange routes between VPN peers.
- Static routes must be configured through OpenVPN itself and not by using gateways and routing in pfSense Plus software.
- Each server can only have one client connection, however, separate servers can be created for each remote client by changing the port number and tunnel network to unique values for each client/server pair.
- OpenVPN DCO is **not compatible** with a `/30` or smaller tunnel network. There are problems with the code for this mode in OpenVPN which can lead to failed connections and instability. That style of connection is being phased out of OpenVPN itself, so the best practice is to avoid that type of configuration.

35.43.2 Example Configuration Overview

When configuring a site-to-site OpenVPN connection using SSL/TLS one node will be the server and the other will be a client.

Tip: Usually a primary site or data center will be the server and a smaller or remote site will act as a client. However, if a remote site has a static IP address and more bandwidth than the primary site that may be a more desirable location for the server.

This style of VPN requires a dedicated subnet for the OpenVPN interconnection between networks in addition to the subnets on both ends. This can be any valid IPv4 subnet so long as it is **larger** than a `/30` (e.g. `/29` or `/24`) and it does not overlap another subnet currently in use on any of the connected networks.

OpenVPN allocates IP addresses for this type of configuration the same way it does for remote access clients. Using a **Topology** style of *subnet*, the client obtains an IP address in a common subnet with the server.

35.43.3 Summary of Requirements

This style of VPN requires several items, all of which are covered in this recipe:

On the server:

- A certificate structure including a certificate authority, server certificate, and a client certificate
- An OpenVPN server instance
- An IPv4 subnet for the OpenVPN tunnel network, larger than a `/30`, such as `10.18.104.0/29`.
- Firewall rules on the WAN to allow a connection from the client node

On the client:

- Imported CA and client certificate
- OpenVPN client instance

On both nodes:

- pfSense Plus software version 22.05 or later for DCO capability.

  The best practice is to use the latest available version, pfSense Plus software version 25.07.1-RELEASE
- Firewall rules to pass traffic inside the tunnel
- An appropriate routing configuration (static or dynamic) – specifics vary by method, see *Configuring Routing for SSL/TLS with DCO*.

35.43.4 Example Configuration Settings

Table 15: OpenVPN Endpoint Settings - Site A - Server

Site A - Server	
Name	Austin
CA Name	S2SCA
Cert CN	serverA
WAN Address	198.51.100.18
LAN Subnet	10.18.0.0/24
LAN Address	10.18.0.1
Tunnel Net	10.18.104.0/29

Table 16: OpenVPN Endpoint Settings - Site B - Client

Site B - Client	
Name	Colorado
Cert CN	clientB
WAN Address	203.0.113.19
LAN Subnet	10.19.0.0/24
LAN Address	10.19.0.1

35.43.5 Configuring SSL/TLS with DCO Server Instance

The server configuration consists of a certificate structure, OpenVPN server instance, firewall rules, and a routing configuration.

Note: Parts of this configuration will vary depending on the routing style the tunnel will use. See *Configuring Routing for SSL/TLS with DCO* for details.

Create Certificate Structure

The first step is to create a certificate structure for this VPN.

This example uses the names listed in *Example Configuration Settings* – The CA is named `S2SCA`, the server certificate CN is `serverA`, and the client certificate CN is `clientB`.

See also:

Certificate Management

Create a Certificate Authority

Create a CA unique to this VPN:

- Navigate to **System > Certificates**, **Authorities** tab
- Click  **Add** to create a new a CA
- Enter the settings as follows:

  Descriptive Name
  `S2SCA`

  Method
  Create an internal Certificate Authority

  Randomize Serial
  Checked

  Key Type
  RSA, *2048* (or higher)

  Digest Algorithm
  sha256 (or higher)

  Lifetime (days)
  `3650`

  Common Name
  `S2SCA`

  Subject Component Fields
  The remaining fields are optional, but can be set to reflect the location of the CA.

- Click **Save**

Create a Server Certificate

Create a server certificate signed by the VPN CA:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add/Sign** to create a new certificate
- Enter the settings as follows:

  Method
  Create an internal Certificate

  Descriptive Name
  `serverA`

  Certificate Authority
  S2SCA

  Key Type
  RSA, *2048* (or higher)

  Digest Algorithm
  sha256 (or higher)

Lifetime (days)
`398`

Note: Some operating systems and software limit server certificates to a maximum lifetime of `398` days for security reasons. Clients on these platforms may reject a server certificate with a longer lifetime.

Common Name
`serverA`

Subject Component Fields
The fields contain data copied from the CA and are optional, but can be set to reflect the location of the server.

Certificate Type
Server Certificate

Warning: This setting is critical, **do not forget to set this value**!

Alternative Names
Optional extra entries, if needed, which specify alternate ways to identify the server. This can be left blank if the certificate will only be used by OpenVPN. Otherwise, add fields with additional information such as alternate hostnames, static IP addresses, and so on which are relevant to this server.

- Click **Save**

Create User Certificate

Create a user certificate for the remote site signed by the VPN CA.

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add/Sign** to create a new certificate
- Enter the settings as follows:

Method
Create an internal Certificate

Descriptive Name
`clientB`

Certificate Authority
S2SCA

Key Type
RSA, *2048* (or higher)

Digest Algorithm
sha256 (or higher)

Lifetime (days)
`3650`

Common Name
`clientB`

Subject Component Fields
The fields contain data copied from the CA and are optional, but can be set to reflect the location of the client.

Certificate Type
User Certificate

> **Warning:** This setting is critical, **do not forget to set this value**!

Alternative Names
Optional extra entries which specify alternate ways to identify the client. These can be left blank if the certificate will only be used by OpenVPN. Otherwise, add fields with additional information such as alternate hostnames, static IP addresses, and so on which are relevant to this client.

- Click **Save**

Export Certificates

The next task is to export the certificates and keys which the client requires when connecting to the OpenVPN server.

- Navigate to **System > Certificates**, **Authorities** tab
- Click  on the row for the CA to export its certificate
- Navigate to **System > Certificates**, **Certificates** tab
- Click  on the row for the client certificate to export the certificate
- Click  on the row for the client certificate to export the private key for the client certificate

> **Warning:** **Do not** export the CA key, server certificate, or server key! The client does not need these files and copying them unnecessarily significantly weakens the security of the VPN.

Configure the OpenVPN Server Instance

- Navigate to **VPN > OpenVPN**, **Servers** tab
- Click  **Add** to create a new server
- Fill in the fields as described below, with everything else left at defaults.

  Use values appropriate for this network, or the defaults if unsure.

  See also:

  - See *Server Configuration Options* for details on each of these options.
  - See *OpenVPN Data Channel Offload (DCO)* for details on options which are incompatible with DCO.

Description
Enter text to describe the connection, e.g. `Site-to-Site VPN`.

Server Mode
Peer to Peer (SSL/TLS)

> **Warning:** DCO requires SSL/TLS, shared key is not compatible.

DCO (Plus Only)
Checked

Protocol
UDP on IPv4 only

Note: DCO is not compatible with TCP.

Interface
WAN

Local Port
`1194`

TLS Configuration
Check the **Use a TLS Key** box to enable TLS authentication which provides protection for the tunnel control channel.

Leave **Automatically generate a TLS Key** checked, so the firewall will generate a new key automatically the first time this entry is saved.

Peer Certificate Authority
Select the CA created at the beginning of this process (`S2SCA`)

Peer Certificate Revocation List
Select a CRL for the CA, if one exists.

Server Certificate
Select the server certificate created at the beginning of this process (`serverA`)

Data Encryption Algorithms
DCO is limited to the AES-256-GCM, AES-128-GCM, and ChaCha20-Poly1305 encryption algorithms. Any of these options will work, provided the same algorithms are selected on the server and client side.

The best algorithm for an environment depends on the hardware and what it may be capable of accelerating. See *Cryptographic & Thermal Hardware* for details.

IPv4 Tunnel Network
Enter the chosen tunnel network, `10.18.104.0/29`

> **Warning:** This cannot be set to a `/30` or `/31` subnet, it must be a larger subnet, such as a `/29` or `/24`.

IPv4 Local Network(s)
If this setup uses static routing, enter the LAN subnet(s) for the **server** site: `10.18.0.0/24`.

This will push a route for these networks to the client.

Tip: If there are more networks on the server side that clients need to reach, such as networks reachable via static routes, other VPNs, and so on, add them as additional entries in the **IPv4 Local Network** box separated by a comma. For example, `10.18.0.0/24, 10.18.5.0/24`.

See also:

See *Configuring Routing for SSL/TLS with DCO* for details on alternative routing configurations.

IPv4 Remote Network(s)
If this setup uses static routing, enter the LAN subnet(s) for the **client** site: `10.19.0.0/24`.

This will direct the operating system to deliver traffic for these networks to this OpenVPN instance.

Tip: If there are more networks on the client side that the server side needs to reach, such as networks reachable via static routes, other VPNs, and so on, add them as additional entries in the **IPv4 Remote Network** box separated by a comma. For example, `10.18.0.0/24, 10.18.5.0/24`.

See also:

See *Configuring Routing for SSL/TLS with DCO* for details on alternative routing configurations.

Other options can remain at their default values.

- Click Save.
- Click  to edit the new server instance
- Find the **TLS Key** box
- Select all the text inside
- Copy the text to the clipboard
- Save this to a file or paste it into a text editor such as Notepad temporarily

Firewall Rules

External Traffic (WAN)

Next, add a firewall rule for the WAN interface which allows access to the OpenVPN server.

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

Protocol
UDP

Source
Any

If the client has a static IP address, set it as the source on this rule instead, either directly or by using an alias.

Destination
WAN Address

Destination port
`1194`

Description
`OpenVPN Site-to-Site DCO VPN`

- Click **Save**
- Click **Apply Changes**

Tunneled Traffic

Now add a rule to the **OpenVPN** tab to pass traffic over the VPN from the Client-side LAN to the Server-side LAN. This can be an "allow all" style rule or a set of stricter rules. This example allows all traffic using this rule:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

Protocol
Any

Source
Any

Tip: For extra security, create an alias containing only the remote hosts or subnets which must initiate contact with hosts on the sever LAN, then use that alias as the source on this rule.

Destination
Any

Tip: For extra security, create an alias containing only the local hosts or subnets on the server LAN which must accept connections from remote hosts across the VPN, then use that alias as the destination on this rule.

Description
`Allow all tunneled OpenVPN traffic from clientB`

- Click **Save**
- Click **Apply Changes**

That completes the server setup, now move on to configure the client.

35.43.6 Configuring SSL/TLS with DCO Client Instance

The client configuration consists of an imported certificate structure, OpenVPN client instance, firewall rules, and a routing configuration.

Note: Parts of this configuration will vary depending on the routing style the tunnel will use. See *Configuring Routing for SSL/TLS with DCO* for details.

Import CA and Certificate

On the client, import the CA certificate along with the client certificate and key. This is the same CA and client certificate created earlier in this document.

See also:

Certificate Management

Import these items at **System > Certificates**.

First import the CA:

- Navigate to **System > Certificates**, **Authorities** tab
- Click  **Add** to create a new certificate authority
- Enter the settings as follows:

  Descriptive Name
  `S2SCA`

  Method
  Import an existing Certificate Authority

  Certificate Data
  Open the CA certificate file in a text editor on the client PC, select all the text, and copy it to the clipboard. Then paste it into this field.

- Click **Save**

Next, import the client certificate:

- Navigate to **System > Certificates**, **Certificates** tab
- Click  **Add** to create a new certificate
- Enter the settings as follows:

  Method
  Import an existing Certificate

  Descriptive Name
  `clientB VPN Certificate`

  Certificate Type
  X.509 (PEM)

Certificate Data
Open the client certificate file in a text editor on the client PC, select all the text, and copy it to the clipboard. Then paste it into this field.

Private Key Data
Open the client certificate private key in a text editor on the client PC, select all the text, and copy it to the clipboard. Then paste it into this field.

- Click **Save**

Configure the OpenVPN Client Instance

After importing the certificates, create the OpenVPN client:

- Navigate to **VPN > OpenVPN**, **Clients** tab
- Click  **Add** to create a new client
- Fill in the fields as follows, with everything else left at defaults:

  See also:

  See *Client Configuration Options* for details on each of these options.

  Description
  Text to describe the connection (e.g. `Site A VPN`)

  Server Mode
  Peer to Peer (SSL/TLS)

  DCO (Plus Only)
  Checked

  Protocol
  UDP on IPv4 only

  Interface
  WAN

  Server host or address
  The public IP address or hostname of the OpenVPN server (`198.51.100.18` in this example)

  Server Port
  `1194`

  Use a TLS Key
  Checked

  Automatically generate a TLS key
  Unchecked

  TLS Key
  Paste in the TLS key copied from the server instance

  Peer Certificate Authority
  The CA imported at the beginning of this process

  Client Certificate
  The client certificate imported at the beginning of this process

  Data Encryption Algorithms
  Match the algorithm(s) selected in the server configuration

IPv4 Tunnel Network
Leave blank

- Click Save

Note: With SSL/TLS server/client configurations such as this example, the server automatically pushes address configuration, routes, and certain other options to the client. These options do not need to be present in the client configuration.

If the client side must reach additional networks, configure them in the **server** settings.

Firewall Rules

This next step is optional depending on whether hosts on the server network need to initiate contact with hosts on the client network. If the other site does not need to initiate contact with hosts behind this client, then no action is necessary.

Incoming connections from the server side to the client side require a firewall rule on the **OpenVPN** tab on the *client* firewall. An "allow all" style rule is OK in some cases, but a set of stricter rules is the best practice.

This example allows all traffic:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  Any

  Source
  Any

  Tip: For extra security, create an alias containing only the remote hosts or subnets which must initiate contact with hosts on the client side, then use that alias as the source on this rule.

  Destination
  Any

  Tip: For extra security, create an alias containing only the local hosts or subnets on the client side which must accept connections from remote hosts across the VPN, then use that alias as the destination on this rule.

  Description
  `Allow all tunneled OpenVPN traffic from serverA`

- Click **Save**
- Click **Apply Changes**

35.43.7 Configuring Routing for SSL/TLS with DCO

Since OpenVPN with DCO enabled cannot handle routing internally in OpenVPN, it requires extra steps to manage routes in certain cases.

Static Routing

The example in this recipe already configures static routing by using the **IPv4 Local Network(s)** and **IPv4 Remote Network(s)** configuration options.

This is the most basic method as it only requires configuring these two values on the server side, and the client side automatically picks up routes the server pushes.

Dynamic Routing with BGP

With DCO enabled the operating system can route over the OpenVPN interface using BGP. However, BGP requires a static neighbor address, so there is one prerequisite step before configuring BGP itself.

Create Client-Specific Override

On the server side, add a **Client Specific Override** for the client certificate to set a static IP address in the tunnel network. With this override entry present, the client instance will obtain the same static IP address each time the VPN connects. This is necessary so the BGP configuration on the server side can locate the client as a neighbor.

- Navigate to **VPN > OpenVPN**, **Client Specific Overrides** tab
- Click  to add a new override
- Fill in the fields on this screen as follows:

  Description
  A name for the override, such as the common name of the client (e.g. `clientB`).

  Common Name
  Enter the CN of the client node. In this example, that is `clientB`.

  Server List
  Select the server instance description.

  IPv4 Tunnel Network
  A static address inside the server **IPv4 Tunnel Network**, e.g. `10.18.104.2/29`.

  > **Warning:** The CIDR mask on this entry **must match** the CIDR mask on the server **IPv4 Tunnel Network**.

- Click **Save**

See also:

- *Client Specific Overrides*

Install and configure FRR and BGP

The FRR add-on package available for pfSense Plus software provides BGP routing functionality. BGP must be configured on **both** the server and client nodes with appropriate settings for each.

BGP configurations can be quite complex and vary significantly for each environment. As such, examples are unlikely to match the needs of most environments. There is a BGP example in *BGP Example Configuration* which works for a basic configuration between two peers, as in this recipe.

When following the *BGP Example Configuration* note the following items of interest:

- BGP peering happens on the **OpenVPN tunnel interface and addresses**. The neighbor address for each peer is the VPN interface IP address on the peer. The server side is always the first address in the subnet (e.g. `10.18.104.1`), the client side will be the IP address set in the client-specific override (e.g. `10.18.104.2`).
- Ensure firewall rules on the OpenVPN tab pass BGP traffic on TCP port 179.

Dynamic Routing with OSPF

With DCO enabled, the operating system can route over the OpenVPN interface using OSPF.

The FRR add-on package available for pfSense Plus software provides OSPF routing functionality. OSPF must be configured on **both** the server and client nodes with appropriate settings for each.

OSPF configurations can be quite complex and vary significantly for each environment. As such, examples are unlikely to match the needs of most environments. There is an OSPF example in *OSPF Example Configuration* which works for a basic configuration between two peers, as in this recipe.

When following the *OSPF Example Configuration* note the following items of interest:

- OSPF peering happens on the **OpenVPN tunnel interface and addresses**.
- Use the OpenVPN server/client interfaces as the backbone area (`0.0.0.0`) interfaces in the example.
- Firewall rules on the OpenVPN tab must pass the OSPF **Protocol**. It is not TCP or UDP. An "allow all" style rule which passes *Any* **Protocol** is also a viable choice, though less secure.
- Firewall rules on the OpenVPN tab must pass multicast traffic destinations for OSPF protocol traffic, it cannot be restricted to specific sources and destinations.

35.43.8 Testing the Connection

The configuration is now complete. OpenVPN client instances automatically start when created, so it will attempt to connect at this point and if the configuration is correct, it will successfully connect to the server.

Try initiating a ping from one LAN to the other LAN in each direction to verify connectivity.

See also:

Troubleshooting OpenVPN.

35.44 OpenVPN Site-to-Site Configuration Example with Shared Key

This section describes the configuration process for a site-to-site connection using a **shared key** style point-to-point mode OpenVPN tunnel. Other names for this style of configuration are static key or pre-shared key (PSK).

> **Danger:** Shared key mode has been deprecated by OpenVPN as it is no longer considered sufficiently secure for modern requirements.
>
> **Shared key mode will be removed from future versions of OpenVPN**. Users **should not** create any new shared key tunnels and should **immediately** convert any existing shared key tunnels to SSL/TLS mode.
>
> When an SSL/TLS instance is configured with a `/30` tunnel network it behaves similarly to shared key mode. The primary difference is the need to create and distribute the certificate structure to peers. See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for information on configuring OpenVPN in SSL/TLS mode.

In this mode each server instance can only accommodate a single client. Additionally, the server cannot push settings to the client, so routes must be added on both peers and other settings must match identically.

See also:

OpenVPN Site-to-Site Configuration Example with SSL/TLS

35.44.1 Example Configuration Overview

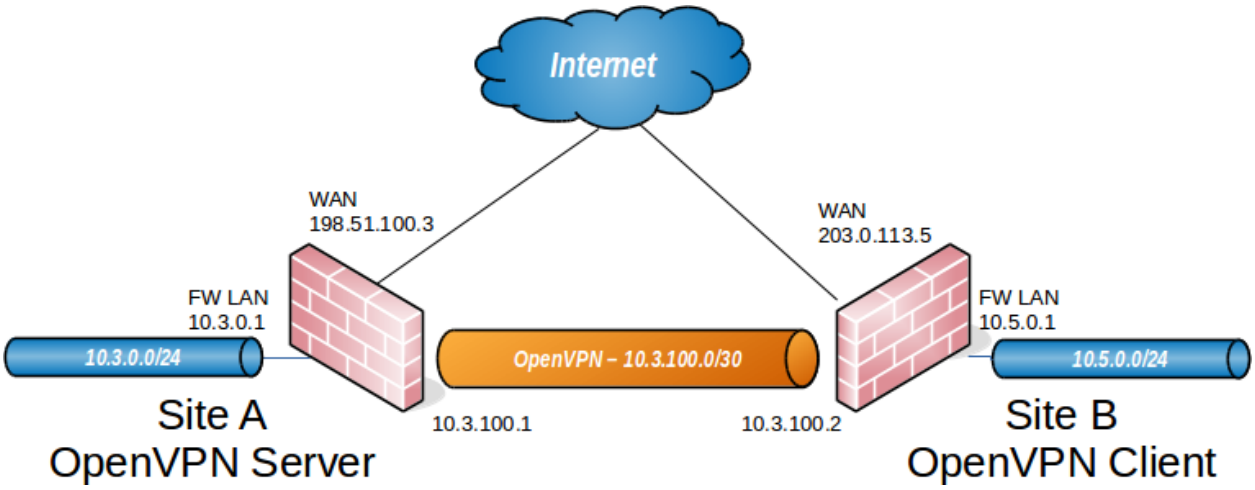


Fig. 70: OpenVPN Example Site-to-Site Network

One firewall will act as the server and the other will act as a client.

> **Tip:** Typically, the main location will be the server side and the remote office will act as a client, though the opposite is functionally equivalent.

In addition to the subnets on both ends this setup requires a dedicated subnet for the OpenVPN interconnection between networks. Figure *OpenVPN Example Site-to-Site Network* shows a depiction of this layout, using `10.3.100.0/24` as the IPv4 VPN Tunnel Network. This can be any subnet so long as it does not overlap another subnet currently in use on the network.

35.44.2 Example Configuration Settings

Table 17: OpenVPN Endpoint Settings - Site A - Server

Site A - Server	
Name	Austin Office
WAN Address	198.51.100.3
LAN Subnet	10.3.0.0/24
LAN Address	10.3.0.1
Tunnel Net	10.3.100.0/30

Table 18: OpenVPN Endpoint Settings - Site B - Client

Site B - Client	
Name	London Office
WAN Address	203.0.113.5
LAN Subnet	10.5.0.0/24
LAN Address	10.5.0.1
Tunnel Net | 10.3.100.0/30	

35.44.3 Configuring PSK Server Side

Configure the OpenVPN Server Instance

- Navigate to **VPN > OpenVPN**, **Server** tab
- Click  **Add** to create a new server entry
- Fill in the fields as follows, with everything else left at defaults:

  See also:

  See *Server Configuration Options* for details on each of these options.

  Description
  Text to describe the connection (e.g. `ExampleCo Site B VPN`)

  Server Mode
  Peer to Peer (Shared Key)

  Device Mode
  tun

  Protocol
  UDP on IPv4 only

  Interface
  WAN

  Local Port
  `1194`

  Shared key
  Check **Automatically generate a shared key**

Tunnel Network
`10.3.100.0/30`

Remote network
The LAN on the Site B side, `10.5.0.0/24`

Note: If there are more networks at Site B, such as networks reachable via static routes, other VPNs, and so on, add them as additional entries separated by a comma (`,`).

- Click **Save**
- Click  to edit this server instance again
- Find the **Shared Key** box
- Select all text inside the **Shared Key** box
- Copy the text to the clipboard
- Save the contents to a file or paste into a text editor such as Notepad temporarily

Firewall Rules

External Traffic (WAN)

Next, add a firewall rule on WAN allowing access to the OpenVPN server.

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  UDP

  Source
  Address or Alias, `203.0.113.5`

  Set the source address to match the client WAN IP address. If the client has a dynamic IP address, set the source to *Any*.

  Destination
  WAN Address

  Destination port
  `1194`

  Description
  `OpenVPN from Site B`

- Click **Save**
- Click **Apply Changes**

When finished, the rule will look like Figure *OpenVPN Example Site-to-Site WAN Firewall Rule*.

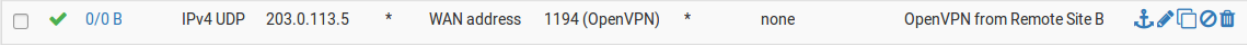


Fig. 71: OpenVPN Example Site-to-Site WAN Firewall Rule

Tunneled Traffic

Now add a rule to the **OpenVPN** tab to pass traffic over the VPN from the Client-side LAN to the Server-side LAN. This can be an "Allow all" style rule or a set of stricter rules. This example allows all traffic using this rule:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  **Add** to create a new rule at the top of the list
- Set the options as follows:

  Protocol
  any

  Source
  any

  Tip: For extra security, create an alias containing only the remote hosts or subnets which must initiate contact with hosts on the sever LAN, then use that alias as the source on this rule.

  Destination
  any

  Tip: For extra security, create an alias containing only the local hosts or subnets on the server LAN which must accept connections from remote hosts across the VPN, then use that alias as the destination on this rule.

  Description
  `Allow all on OpenVPN`

- Click **Save**
- Click **Apply Changes**

That completes the server setup, next, now move on to configure the client.

35.44.4 Configuring PSK Client Side

Configure the OpenVPN Client Instance

- Navigate to **VPN > OpenVPN**, **Client** tab on the client system
- Click  **Add** to create a new OpenVPN client instance
- Fill in the fields as follows, with everything else left at defaults:

  See also:

  See *Client Configuration Options* for details on each of these options.

Description
Text to describe the connection (e.g. `ExampleCo Site A VPN`)

Server Mode
Peer to Peer (Shared Key)

Device Mode
tun

Protocol
UDP on IPv4 only

Interface
WAN

Server host or address
The public IP address or hostname of the OpenVPN server (`198.51.100.3` in this example)

Server Port
`1194`

Shared key
Uncheck **Automatically generate a shared key**, then paste in the shared key for the connection using the key copied from the server instance created previously.

Tunnel Network
`10.3.100.0/30`

> **Warning:** This must match the server side exactly. The firewall will use the correct address for each end of the tunnel.

Remote network
The LAN on the Site A side, `10.3.0.0/24`

> **Note:** If there are more networks at Site A, such as networks reachable via static routes, other VPNs, and so on, add them as additional entries separated by a comma (`,`).

- Click **Save**

Firewall Rules

This setup does not require firewall rules on the client side WAN interface because the client only initiates outbound connections. The server never initiates connections to the client.

This next step is optional depending on whether hosts on the server network need to initiate contact with hosts on the client network. If server network hosts do not need to initiate contact with this client, then no action is necessary.

If hosts on the server side need to initiate contact, then this traffic requires a firewall rule on the **OpenVPN** tab on the *client* firewall to allow traffic from the Server-side LAN to reach the Client-side LAN. An "Allow all" style rule is OK in some cases, but a set of stricter rules is the best practice.

This example allows all traffic:

- Navigate to **Firewall > Rules**, **OpenVPN** tab
- Click  **Add** to create a new rule at the top of the list

- Set the options as follows:

  Protocol
  any

  Source
  any

  > **Tip:** For extra security, create an alias containing only the remote hosts or subnets which must initiate contact with hosts on the client LAN, then use that alias as the source on this rule.

  Destination
  any

  > **Tip:** For extra security, create an alias containing only the local hosts or subnets on the client LAN which must accept connections from remote hosts across the VPN, then use that alias as the destination on this rule.

  Description
  `Allow all on OpenVPN`

- Click **Save**
- Click **Apply Changes**

The configuration of the client is complete.

35.44.5 Testing the connection

The configuration is now complete. The OpenVPN client instance automatically starts when created, so it should already be attempting to connect at this point and if the configuration is correct, it will be connected.

Try to ping across to the remote end LAN to verify connectivity.

See also:

Troubleshooting OpenVPN.

35.45 OpenVPN Remote Access Configuration Example

The OpenVPN wizard on pfSense® software is a convenient way to set up a remote access VPN for mobile clients. The wizard configures all the necessary prerequisites for an OpenVPN remote access server:

- An authentication source (Local, RADIUS server, or LDAP server)
- A certificate authority (CA)
- A server certificate
- An OpenVPN server instance

At the end of the wizard the firewall will have a fully functioning sever, ready to accept connections from users. This server configuration can then be altered as needed.

This document uses an example setup to aid in explaining the options available in the wizard.

See also:

Server Configuration Options

35.45.1 Before Starting The Wizard

Before starting the wizard, plan the design of the VPN.

Determine an IP addressing scheme

The OpenVPN server requires a dedicated subnet for communication between the server and the OpenVPN clients. This is the **Tunnel Network** in the server configuration. The server uses the first address in this subnet for itself to act as a gateway, and it allocates IP addresses within this subnet to clients.

When selecting internal subnets for a single location, ideally choose subnets which can be CIDR summarized with other internal subnets. This example uses `10.3.0.0/24` for LAN and `10.3.201.0/24` for the remote access OpenVPN server. These two networks can be summarized with `10.3.0.0/16`, which makes routing easier to manage.

See also:

CIDR Summarization

Example Network

Figure *OpenVPN Example Remote Access Network* shows a depiction of this example deployment.

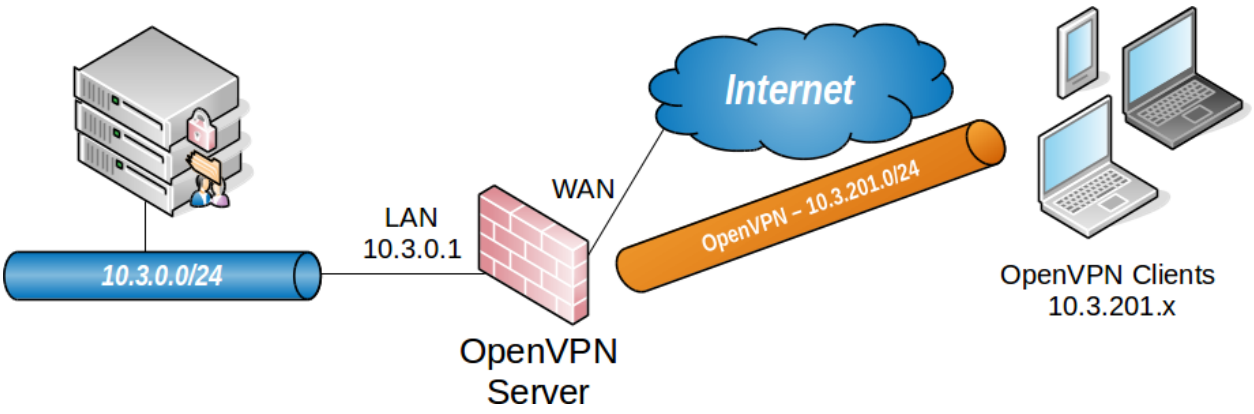


Fig. 72: OpenVPN Example Remote Access Network

Table 19: OpenVPN Remote Access Server Settings

Remote Access Server	
WAN Address	198.51.100.3
LAN Subnet	10.3.0.0/24
LAN Address	10.3.0.1
Tunnel Net	10.3.201.0/24

35.45.2 OpenVPN Wizard Walkthrough

To start the OpenVPN Remote Access Server Setup wizard:

- Navigate to **VPN > OpenVPN**
- Click the **Wizards** tab

The GUI presents the first step of the wizard automatically

Note: The option for *OpenVPN Data Channel Offload (DCO)* is not included in this wizard. To use DCO on this server, run the wizard first then after completing the wizard, edit the server instance and enable the DCO option.

Choose Authentication Type

On the first screen of the wizard, select the authentication backend server type. The choices available for **Type of Server** are *Local User Access*, *LDAP*, and *RADIUS*.

See also:

Authentication Servers

Local User Access
Manage the users, passwords, and certificates using the *User Manager* on this firewall.

Sets the server mode to *Remote Access (SSL/TLS + User Auth)* which requires user authentication as well as per-user certificates.

Local User Access easily handles per-user certificates, managed completely in the GUI. This is much more secure, but depending on the number of users which will access the service, may be less convenient than using a central authentication system.

LDAP / RADIUS
If the network has an existing authentication system already in place, such as Active Directory, pick *LDAP* or *RADIUS* depending on which method that system accepts.

LDAP and *RADIUS* both set the server mode to *Remote Access (User Auth)*, which does not require per-user certificates.

Note: The server mode can be adjusted later to require certificates, but administrators must manually create per-user certificates for *LDAP* or *RADIUS* users.

Click  **Next** to continue.

For *Local User Access*, the wizard skips the LDAP and RADIUS configuration steps.

For *LDAP* or *RADIUS* the wizard will present appropriate authentication server configuration options next. This example uses *Local User Access*, but this document discusses the other options for completeness.

Choosing an LDAP Server

If the user manager configuration on this firewall contains one or more LDAP servers, the wizard offers these LDAP servers as options it can use for this VPN.

Click  **Add new LDAP server** to create a different LDAP server entry.

Click  **Next** to continue using the server selected in the **LDAP Servers** list.

If the firewall configuration does not contain any LDAP servers, the wizard skips this step.

See also:

Authentication Servers

Adding an LDAP Server

If the user manager configuration on this firewall does not contain an LDAP server, or if the user chose to create a new LDAP server, the wizard presents a screen to define a new server.

The values for the options on this screen depend on the specific LDAP directory configuration and structure. For guidance, consult the LDAP server administrator, software vendor, or documentation.

Note: The details of LDAP servers are covered in *LDAP Authentication Servers*.

This document omits some detail since the options are discussed in-depth by that other section.

The wizard offers the following LDAP authentication server parameters:

Name
Descriptive name for this LDAP server, for reference.

Hostname or IP address
The hostname or IP address of the LDAP server.

If the firewall contacts this server using an encrypted method, this value must match the contents of the LDAP server certificate.

Port
The port on which the LDAP server is listening for requests.

The default port is `389` for standard TCP connections and `636` for SSL.

Transport
Sets the method the firewall will use when performing LDAP queries to the server.

Standard TCP
Unencrypted connections using plain TCP.

STARTTLS Encrypted
Connects to the standard TCP port and then attempts to negotiate TLS encryption.

SSL/TLS Encrypted
Secure connections using TLS encryption.

A standard TCP connection is typically sufficient for initial testing, and potentially for local servers or those only accessible over secure connections. If the server is remote or crosses any untrusted network

links, an encrypted method is essential. Using an encrypted method is always the best practice, but may not always be viable.

> **Warning:** When the firewall uses an encrypted method to contact the LDAP server, the **Hostname or IP address** above must match a value in the LDAP server certificate.

Peer Certificate Authority
To use SSL/TLS or STARTTLS transports, the firewall must trust the CA of the LDAP server. This can be accomplished by any of the following methods:

- Import the CA into the certificate manager and select it from the list in this option.
- Import the CA into the certificate manager with the **Trust Store** option set, which adds the imported CA into the list of CAs which the firewall trusts globally. Then select *global* from this list.
- If the LDAP server certificate is signed by a globally trusted CA, such as Let's Encrypt, then select *global*.

Search Scope Level
Selects how deep the firewall will search in the LDAP directory, *One Level* or *Entire Subtree*.

In almost all cases, *Entire Subtree* is the correct choice.

Search Scope Base DN
The distinguished name (DN) upon which the firewall bases its search. For example `DC=example, DC=com`.

Authentication Containers
These values specify where the directory stores user data. For example, `CN=Users;DC=example`.

LDAP Bind User DN
If the LDAP server requires authenticated binds when performing queries, this field sets the distinguished name the firewall uses for this bind action.

If this is blank the firewall performs an anonymous bind without credentials.

LDAP Bind Password
The password for authenticated binds. The firewall only uses this value if **LDAP Bind User DN** has a value.

User Naming Attribute
Varies depending on the LDAP directory software and structure.

Typically, `cn` for OpenLDAP and Novell eDirectory, and `samAccountName` for Microsoft Active Directory.

Group Naming Attribute
Varies depending on the LDAP directory software and structure, but is most typically `cn`.

Member Naming Attribute
Varies depending on the LDAP directory software and structure.

Typically, `member` on OpenLDAP, `memberOf` on Microsoft Active Directory, and `uniqueMember` on Novell eDirectory.

See also:

LDAP Authentication Servers explains the remaining options in detail, and when a server may require them.

Click  **Add new server** to continue.

Choosing a RADIUS Server

If the user manager configuration on this firewall contains one or more RADIUS servers, the wizard offers these RADIUS servers as options it can use for this VPN.

Click  **Add new RADIUS server** to create a different RADIUS server entry.

Click  **Next** to continue using the server selected in the **RADIUS Servers** list.

If the firewall configuration does not contain any RADIUS servers, the wizard skips this step.

See also:

Authentication Servers

Adding a RADIUS Server

If the user manager configuration on this firewall does not contain a RADIUS server, or if the user chose to create a new RADIUS server, the wizard presents a screen to define a new server.

The values for the options on this screen depend on the specific RADIUS configuration and structure. For guidance, consult the RADIUS server administrator, software vendor, or documentation.

Note: The details of RADIUS servers are covered in *RADIUS Authentication Servers*.

This document omits some detail since the options are discussed in-depth by that other section.

The wizard offers the following RADIUS authentication server parameters:

Name
: Descriptive name for this RADIUS server, for reference.

Hostname or IP address
: The hostname or IP address of the RADIUS server.

Authentication Port
: Port used by the RADIUS server for accepting authentication requests, typically `1812`.

Shared Secret
: The password the RADIUS server expects from this firewall when it submits authentication requests (e.g. password on the NAS entry.)

Click  **Add new server** to continue.

Choosing a Certificate Authority

If the certificate manager configuration on this firewall contains one or more certificate authorities, the wizard offers these CA entries as options it can use for this VPN.

Click  **Add new CA** to create a different certificate authority.

Click  **Next** to continue using the certificate authority selected in the **Certificate authority** list.

If the firewall configuration does not contain any CA entries, the wizard skips this step.

See also:

Certificate Management

Creating a Certificate Authority

If the certificate manager configuration on this firewall does not contain a CA, or if the user chose to create a new CA, the wizard presents a screen to define a new CA.

See also:

For more information on creating and managing CAs, see *Certificate Authority Management*.

This document omits some detail since the options are discussed in-depth by that other section.

The firewall uses this entry as a root CA which can sign server and user certificates. Clients can use this CA to validate the server, and the server can use this CA to validate clients. Because this CA is self-signed, only clients which are supplied with a copy of this CA certificate will trust other certificates signed by this CA.

The wizard offers the following CA parameters:

Descriptive Name
`ExampleCoCA`

A name for reference to identify this certificate. This is the same as **Common Name** field for other certificates.

Note: Although this field can technically contain spaces, the best practice is to conform the contents of this field to the format allowed for fully qualified domain names.

Some clients have issues handling entries with spaces properly.

Key Length
`2048`

Size of the CA private key which the wizard will generate.

Larger keys offer increased security but larger keys are generally slower to use.

Lifetime
`3650`

The time, in days, for which this CA will remain valid.

For a self-signed CA such as this, the default of `3650` is acceptable, which is approximately 10 years.

The remaining fields are optional but define additional identifying data for the CA "subject"/distinguished name. For small deployments this may not matter much, but for larger organizations with CA entries at multiple sites, this can help ensure each CA is easily identifiable.

Country Code
`US`

(Optional) Two-letter ISO country code (e.g. US, AU, CA).

ExampleCo is located in the United States which has an ISO country code of `US`.

To locate an appropriate ISO code for other countries, use the ISO Online Browsing Platform site.

State or Province
`Texas`

(Optional) Full unabbreviated State or Province name (e.g. Texas, Indiana, California).

ExampleCo is located in `Texas`.

City
`Austin`

(Optional) City or other Locality name (e.g. Austin, Indianapolis, Toronto).

ExampleCo headquarters is in `Austin`.

Organization
`ExampleCo`

(Optional) Organization name, often the Company or Group name.

> **Warning:** Do not use any special characters in this field, not even punctuation such as a period or comma.

Click  **Add new CA** finish the CA creation process.

Choosing a Server Certificate

If the certificate manager configuration on this firewall contains one or more certificates, the wizard offers these certificate entries as options it can use for this VPN.

Click  **Add new Certificate** to create a different certificate.

Click  **Next** to continue using the certificate selected in the **Certificate** list.

If the firewall configuration does not contain any certificate entries, the wizard skips this step.

See also:

Certificate Management

Adding a Server Certificate

If the certificate manager configuration on this firewall does not contain a certificate, or if the user chose to create a new certificate, the wizard presents a screen to define a new server certificate.

See also:

For more information on creating and managing certificates, see *Certificate Management*.

This document omits some detail since the options are discussed in-depth by that other section.

This server certificate verifies the identity of the server to the clients. The CA set in the previous wizard steps will sign this certificate. In most cases, as with this example, the server certificate uses the same information from the previous step and the wizard pre-fills the form automatically.

Descriptive Name
`vpn.example.com`

This is the common name (CN) field of the server certificate and the firewall also uses this name to reference the certificate.

The best practice is to set this to the fully qualified hostname of the firewall.

Note: Although this field can technically contain spaces, the best practice is to conform the contents of this field to the format allowed for fully qualified domain names.

Some clients have issues handling entries with spaces properly.

Key Length
`2048`

Size of the CA private key which the wizard will generate.

Larger keys offer increased security but larger keys are generally slower to use.

Lifetime
`398`

The time in days that this certificate will be valid. The best practice is to set this to `398` days or less.

Note: Some current operating systems and software limit server certificates to a maximum lifetime of `398` days for security reasons. Clients on these platforms may reject a server certificate with a longer lifetime.

The remaining fields are optional but define additional identifying data for the server certificate "subject"/distinguished name. For small deployments this may not matter much, but for larger organizations with many server certificates, this can help ensure each certificate is easily identifiable.

Country Code
`US`

(Optional) Two-letter ISO country code (e.g. US, AU, CA).

ExampleCo is located in the United States which has an ISO country code of `US`.

To locate an appropriate ISO code for other countries, use the ISO Online Browsing Platform site.

State or Province
`Texas`

(Optional) Full unabbreviated State or Province name (e.g. Texas, Indiana, California).

ExampleCo is located in `Texas`.

City
`Austin`

(Optional) City or other Locality name (e.g. Austin, Indianapolis, Toronto).

ExampleCo headquarters is in `Austin`.

Organization
`ExampleCo`

(Optional) Organization name, often the Company or Group name.

> **Warning:** Do not use any special characters in this field, not even punctuation such as a period or comma.

Click  **Create New Certificate** to continue.

Configuring OpenVPN Server Settings

The options on this step of the wizard configure each aspect of how the OpenVPN server itself behave as well as options the server will pass on to clients.

See also:

The options presented here are the same as those in *Server Configuration Options*. Refer to that section for details.

Because the options are covered in detail in that section, this document only mentions the settings used by this example.

General OpenVPN Server Information

These options control how the OpenVPN instance operates.

Interface
WAN

Protocol
UDP on IPv4 Only

Local Port
`1194`

The wizard suggests the first unused port number starting with port `1194`. If there is an existing OpenVPN server on that port, use a different port number.

Description
`ExampleCo Mobile VPN Clients`

Cryptographic Settings

These options control how the server encrypts and authenticates traffic in the tunnel.

TLS Authentication
Check **Enable authentication of TLS packets**

Using TLS authentication is the best practice.

Generate TLS Key
Check **Automatically generate a shared TLS authentication key**

TLS Shared Key
Blank

The wizard disables this field when **Automatically generate a shared TLS authentication key** is checked.

DH Parameters Length
2048

This value is a good balance of speed and strength.

Data Encryption Negotiation
Checked

This allows the server to automatically negotiate encryption settings with clients.

Note: Disabling this option is deprecated, but still present on this version for compatibility.

Data Encryption Algorithms
AES-256-GCM, *AES-128-GCM*, and *CHACHA20-POLY1305*

The best practice is to use the default suggested values as noted above.

Fallback Data Encryption Algorithm
AES-256-CBC

This algorithm is used when negotiation fails, for example with an old client that does not implement negotiation.

Auth Digest Algorithm
SHA256 (256-bit)

Tunnel Settings

These options control how the server routes traffic from remote clients.

Tunnel Network
`10.3.201.0/24`

This is the tunnel network from the table at the start of this example (*OpenVPN Remote Access Server Settings*).

Redirect Gateway
Unchecked

For this example, The VPN will only carry traffic destined for subnets at the main office.

Local Network
`10.3.0.0/24`

This is the server-side LAN subnet from the table at the start of this example (*OpenVPN Remote Access Server Settings*).

Concurrent Connections
Blank

This example does not limit the number of clients which can connect at the same time.

Allow Compression
Refuse any non-stub compression (Most secure)

The best practice is to disable compression for security reasons.

Compression
Disable Compression [Omit Preference]

The best practice is to disable compression for security reasons.

Type-of-Service
Unchecked

There is no traffic on this example VPN which requires prioritization/QoS.

Inter-Client Communication
Unchecked

The clients on this VPN have no need to connect to other VPN client hosts.

Duplicate Connections
Unchecked

This example uses unique certificates for every client and does not allow multiple connections per client.

Client Settings

These options control specific settings the server pushes to clients when they establish a connection.

Dynamic IP
Checked

The clients connect from all over the country and unknown mobile networks and their IP addresses are likely to change without notice.

Topology
Subnet

The method the server uses to assign IP addresses to clients.

DNS Default Domain
`example.com`

The domain name used by ExampleCo.

DNS Servers
`10.3.0.5`

A list of internal DNS servers. ExampleCo has a Windows Active Directory Domain Controller which is configured to act as a DNS server at `10.3.0.5`.

NTP Servers
`10.3.0.6`

A dedicated local NTP server exists at `10.3.0.6`.

Advanced
Blank

At this time no additional tweaks are necessary.

Click  **Next** to continue.

Firewall Rule Configuration

By default, the firewall blocks all traffic from connecting to VPNs or passing over VPN tunnels. This step of the wizard adds firewall rules automatically to allow traffic to connect to the VPN and also so connected clients can pass traffic over the VPN.

Traffic from clients to server
Checked

When checked, the wizard adds a firewall rule on the chosen interface outside the tunnel where the server is listening (e.g. *WAN*) which allows VPN clients to connect. The rule created by this option allows all clients from any source IP address to connect by default.

Since clients in this example are connecting from all over the country, the rule created by the wizard for this option is ideal.

Tip: To allow connections from a limited set of IP addresses or subnets, either make a custom rule or check this box and alter the rule it creates.

Traffic from clients through VPN tunnel
Checked

This setting allows all traffic to cross inside the OpenVPN tunnel. This is desirable for this example.

Click  **Next** to continue.

Finishing the Wizard

Click  **Finish** to complete the wizard.

At this point, the firewall now contains a full OpenVPN remote access server configuration which is ready for client connections.

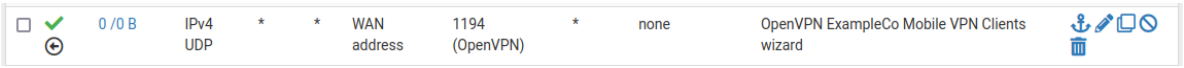
From here, the next steps are to add users and configure client devices.

If this setup requires adjustments to the automatically generated firewall rules, make them now.

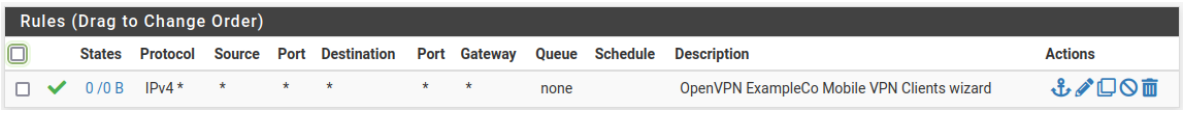
35.45.3 Verifying the Setup

Look at firewall rules (**WAN** and **OpenVPN** tabs)

- **WAN** tab rule should pass from any to the *OpenVPN* port on the *WAN address*



- **OpenVPN** tab rule should allow all traffic from any/to any



35.45.4 Adjustments

Numerous settings are not present in the wizard but might be a better fit for certain deployments than the defaults chosen by the wizard.

Server Mode

The OpenVPN **Server Mode** allows selecting a choice between requiring Certificates, User Authentication, or both. The wizard defaults to *Remote Access (SSL/TLS + User Auth)* when using local users and *Remote Access (User Auth)* for RADIUS and LDAP. The possible values for this choice and their advantages are:

Remote Access (SSL/TLS + User Auth)

- Requires both certificates **and** username/password
- Each user has a unique client configuration which includes their personal certificate and key
- Most secure as there are multiple factors of authentication (TLS Key and Certificate that the user has, and the username/password they know)

Remote Access (SSL/TLS)

- Certificates only, no authentication
- Each user has a unique client configuration which includes their personal certificate and key
- Useful if clients should not be prompted to enter a username and password
- Less secure as it relies only on something the user has (TLS key and certificate)

Remote Access (User Auth)

- Authentication only, no certificates
- Useful if the clients cannot have individual certificates
- Commonly used for external authentication (RADIUS, LDAP)
- All clients can use the same exported client configuration and/or software package
- Less secure as it relies on a shared TLS key plus only something the user knows (User-name/password)

OpenVPN Data Channel Offload (DCO)

OpenVPN Data Channel Offload (DCO), a pfSense Plus exclusive feature, can potentially increase performance of OpenVPN well beyond the capabilities of traditional OpenVPN connections.

Note: Some OpenVPN features and use cases are not compatible with DCO. See *Limitations* for a list of known DCO limitations.

Certificate Revocation

Compromised certificates can be revoked by a Certificate Revocation List (CRL). CRL entries are managed at **System > Certificates**, on the **Certificate Revocation** tab. Create a new CRL, add the certificate to it, and then select that CRL on the OpenVPN server settings.

See also:

Certificate Revocation List Management

35.45.5 Adding a User with a Certificate

If the server mode includes local user authentication, a user must exist in the user manager for each client which will connect to the VPN.

See also:

This is a simplified version of the process. For more detail, see:

- *Adding OpenVPN Remote Access Users*
- *Manage Local Users*
- *User Certificates*
- Navigate to **System > User Manager**
- Click  To add a user
- Fill in the settings as follows:

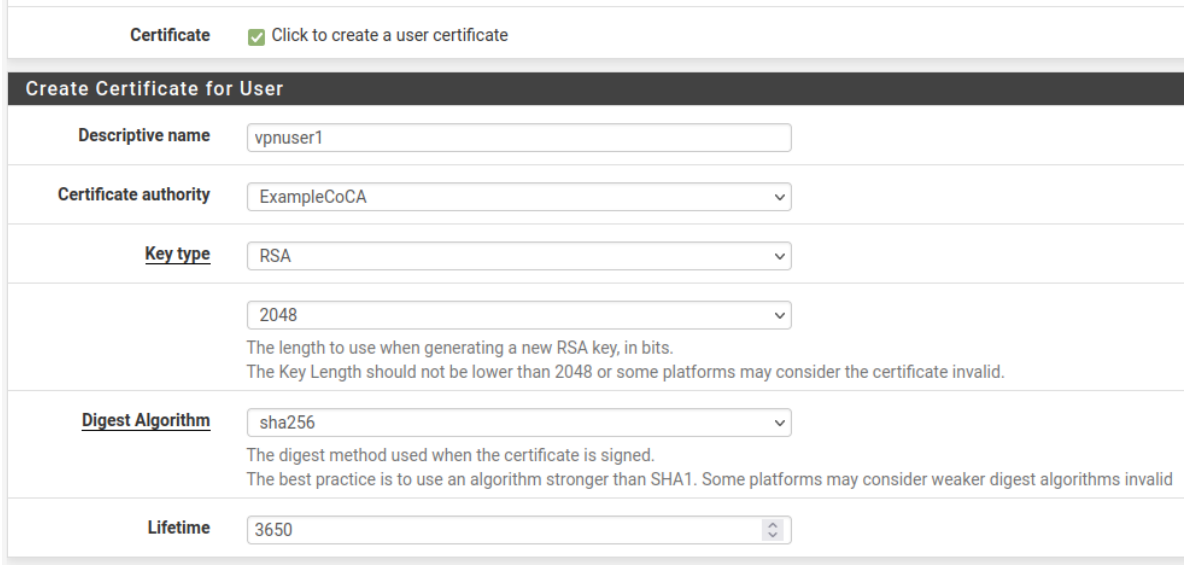
  Username
  The username for this client.

  Password/Confirm password
  The password for this client.

  Click to create a user certificate
  Checked

  Descriptive Name
  Same value as the **Username**

  Certificate Authority
  The CA used by the OpenVPN server.



- Click **Save**

35.45.6 OpenVPN Client Export Package

The OpenVPN Client Export Package can export client configurations formatted for a wide variety of platforms. It can also export a pre-packaged Windows installer executable which includes the configuration bundled inside for a painless client installation.

See also:

OpenVPN Client Export Package

35.46 Adding OpenVPN Remote Access Users

The method for adding users to the VPN depends upon the OpenVPN server authentication method and backend (e.g. Local Database, RADIUS, LDAP).

See also:

- *OpenVPN Remote Access Configuration Example*
- *Server Configuration Options*
- *User Management and Authentication*
- *User Certificates*

35.46.1 Local Database

OpenVPN authenticates local database users based on their entries in the user manager.

To create a new user with a certificate, follow these steps:

- Navigate to **System > User Manager**
- Click  To add a user
- Fill in the settings as follows:

  Username
  The username for this client.

  Password/Confirm password
  The password for this client.

  Full Name
  An optional longer name for this user.

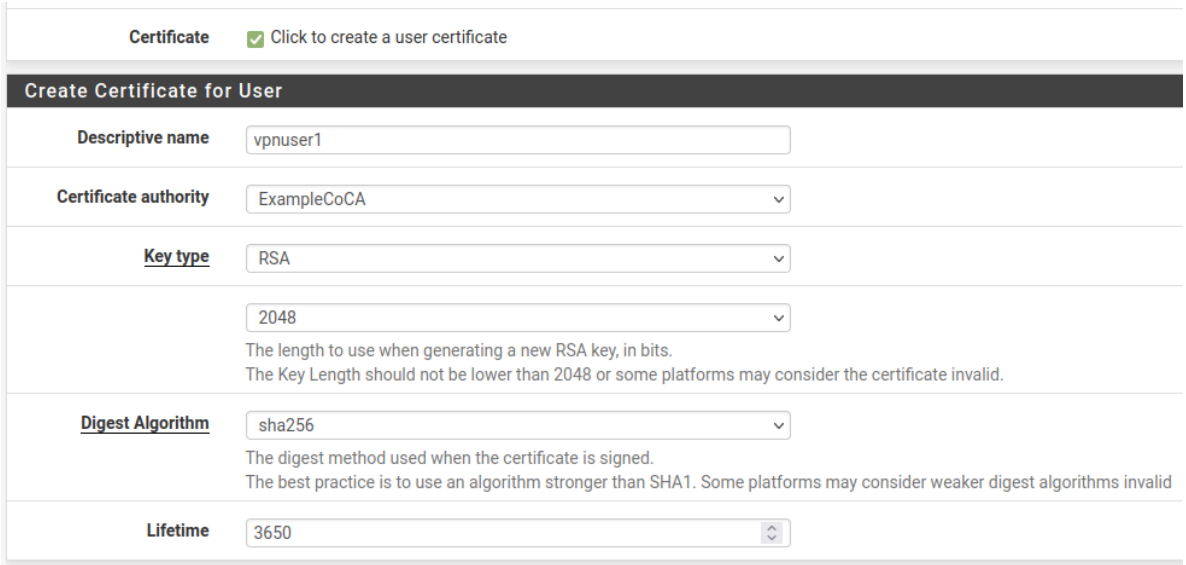
  Click to create a user certificate
  Checked

  Descriptive Name
  Same value as the **Username**

  Certificate Authority
  The CA used by the OpenVPN server.

  Key Type
  The type of private key to use for this certificate, either *RSA* or *ECDSA* and its accompanying **Key Length** (RSA) or **Curve** (ECDSA). The default is an acceptable choice.

  Lifetime
  The number of days for which the certificate is valid. The default of `3650` (approximately 10 years) is acceptable for a user certificate.



- Click **Save**

To view or change the user:

- Navigate to **System > User Manager**
- Click  next to the row containing the user to see/edit

To export a user certificate and key:

Note: The OpenVPN Client Export Package (*OpenVPN Client Export Package*) can package up the certificates and other data automatically. The client export package is a much easier way to download client configurations and installation files than exporting these items manually.

- Navigate to **System > Certificates**, **Certificates** tab
- Locate the user certificate in the list
- Click  to download the user certificates
- Click  to download the key for the certificate

  ---

  Tip: To password protect the private key, edit the certificate, enter an **Export Password**, and click **Export Private Key** on that page.

- Click  to download a PKCS#12 bundle which includes the user certificate and key, and the CA Certificate (optional).

  Tip: To password protect the PKCS#12 bundle, edit the certificate, enter an **Export Password**, and click **Export PKCS#12** on that page.

Most use cases which utilize a user certificate will also require the CA Certificate which signed the user certificate. Export the CA certificate from its entry on **System > Certificates**, **CAs** tab or use one of the PKCS#12 bundle options.

35.46.2 LDAP or RADIUS Users

Adding LDAP and RADIUS users fully depends on the server implementation and management tools, which are beyond the scope of this documentation. Contact the server administrator or software vendor for assistance.

Certificates for LDAP or RADIUS users cannot be created from within the firewall GUI in a way that reflects a user-certificate relationship, but they can be created independently as user certificates in the certificate manager. Create user certificates with the certificate manager as described in *User Certificates*

35.47 Installing OpenVPN Remote Access Clients

Most end-user devices require installation of an OpenVPN client as the client functionality is not yet built into most operating systems. This section provides an overview of OpenVPN client installation on several common operating systems.

The *OpenVPN Client Export Package* automatically generates client configuration files and installation bundles for use by client devices.

See also:

Visit the Hangouts Archive to watch the September and October 2015 Hangouts on Remote Access VPNs which covers client installations for most operating systems.

35.47.1 Installing the OpenVPN Client on Windows

The OpenVPN project provides 64-bit and 32-bit installers for Windows 7 through Windows 11 on The OpenVPN Community Downloads Page. Alternately, use *OpenVPN Client Export Package* to create a self-executable client installer bundled with an appropriate configuration file.

The client installation is straightforward, the user can accept all the default values and actions.

Current stable versions of the OpenVPN client for Windows utilize a system service, so the GUI can run without elevated privileges. Once installed, a user does not need administrative access to run the client.

While running, the OpenVPN client appears as an icon in the system tray. This icon can connect/disconnect VPNs or display additional information, such as connection logs.

The installation creates a new **Local Area Connection** adapter on the client system for OpenVPN. This interface indicates it is connected when the client has established a VPN connection and will otherwise show as disconnected. OpenVPN manages the configuration of this adapter, and it does not require any manual adjustments.

35.47.2 Installing the OpenVPN Client on macOS

There are three client options for macOS:

- The OpenVPN command line client. Most users prefer a graphical client, so this document does not cover that option.
- Tunnelblick, a free option available for download at the Tunnelblick Website.
- The commercial Viscosity client. At the time of this writing, it costs $14 USD for a single seat. Viscosity is a much nicer client and well worth the cost for frequent OpenVPN users.

Both Tunnelblick and Viscosity are easy to install, with no configuration options during installation. Both clients can accept configurations generated by the *OpenVPN Client Export Package*.

Configuring Viscosity

The Viscosity client can be configured manually, or it can import configurations from the OpenVPN Client Export package.

Viscosity provides a GUI configuration tool that can generate the underlying OpenVPN client configuration based on a manual configuration. This section covers the must simpler process of importing a Viscosity bundle generated by the OpenVPN Client Export package.

- Download a copy of the **Viscosity bundle** for the client from the OpenVPN Client Export package
- Locate the bundle file

  The bundle filename ends in `.visc.zip` indicating that it is a compressed archive.
- Copy this bundle file to a folder on the client Mac
- Double-click this file and macOS expands it to `Viscosity.visc`
- Double-click `Viscosity.visc` and Viscosity will open and import the connection as shown in Figure *Viscosity Import*

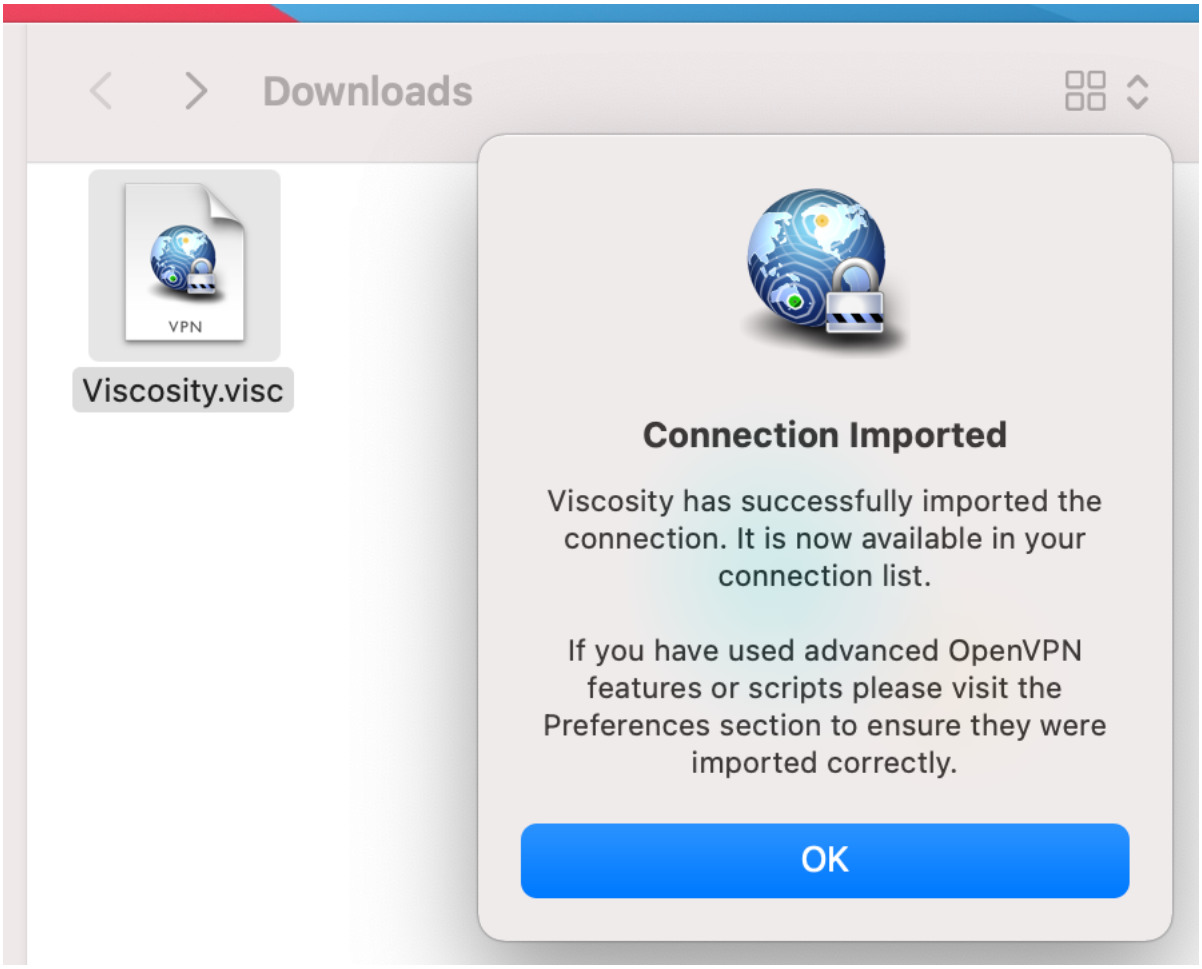


Fig. 73: Viscosity Import

- Delete the `Viscosity.visc` directory and the `.zip` archive
- Viscosity will be running after import and has an icon in the menu bar which looks like a circle with a lock

- Click the Viscosity icon in the menu bar at the top of the screen
- Click **Preferences** to check if Viscosity imported the configuration as shown in Figure *Viscosity Preferences*

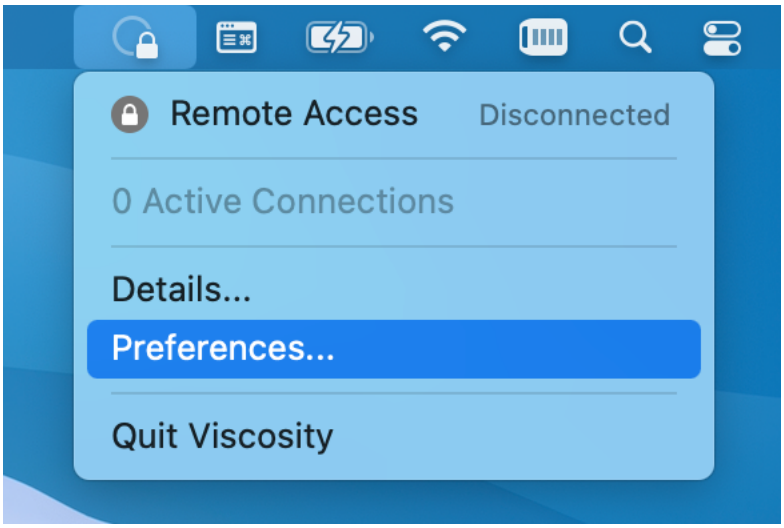


Fig. 74: Viscosity Preferences

- Check the **Connections** area to see if Viscosity imported the connection successfully as shown in Figure *Viscosity View Connections*.
- Close the Preferences screen

Connecting a VPN with Viscosity

- Click the Viscosity icon in the menu bar
- Click the name of the VPN connection to connect as shown in Figure *Viscosity Connect*

  After a few seconds the lock icon for this connection in the Viscosity menu will turn green if the connection attempt succeeds, and Viscosity displays some basic connection information.

Note: When at least one VPN is connected the appearance of the Viscosity icon also changes from a faint circle with a lock to a dark circle with a lock, but this can be easy to miss.

View Connection Status in Viscosity

To view status information about a VPN connection:

- Click the Viscosity icon
- Click **Details** as shown in Figure *Viscosity Menu*

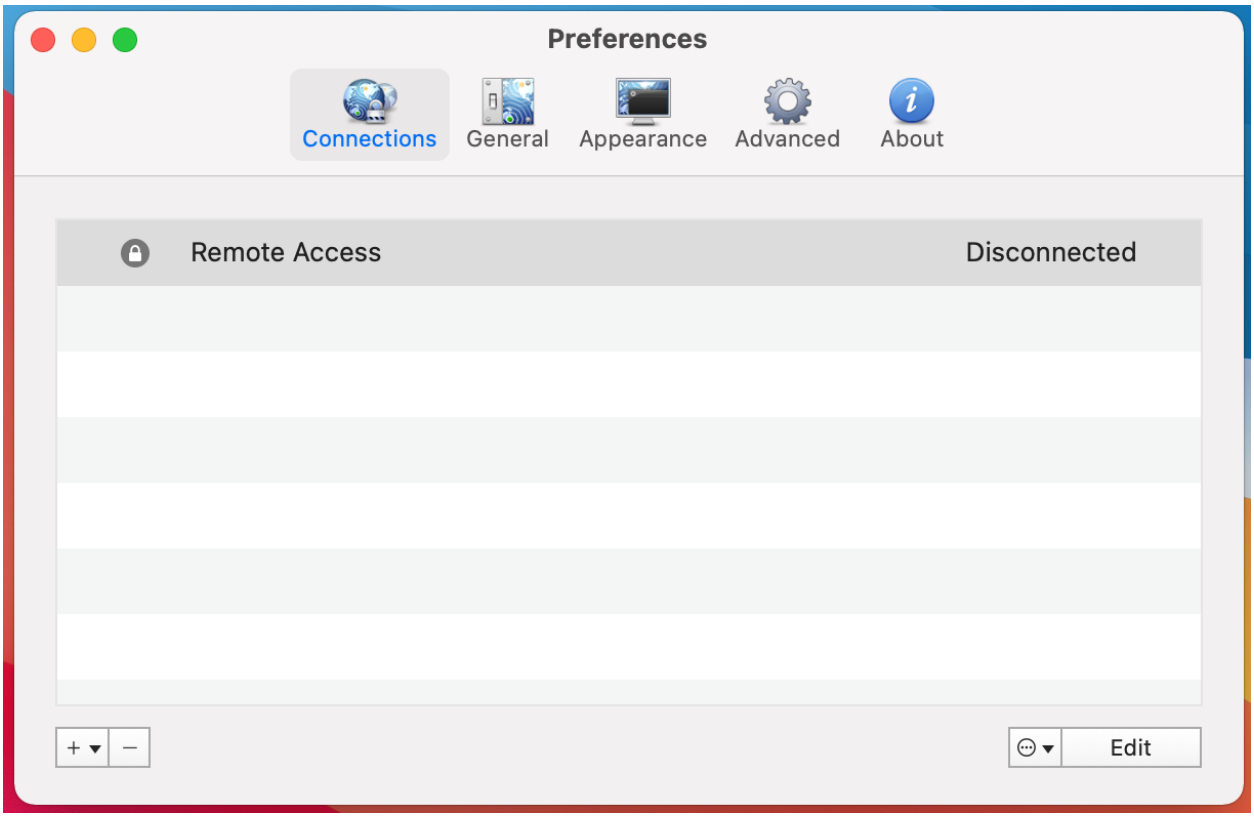


Fig. 75: Viscosity View Connections

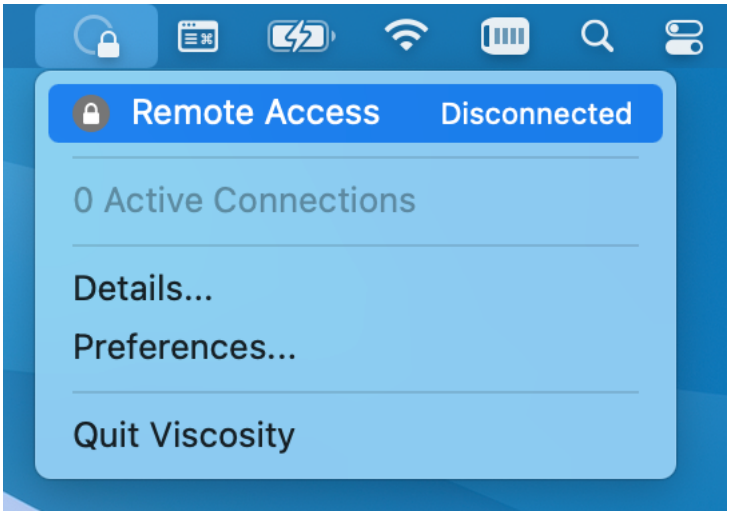


Fig. 76: Viscosity Connect

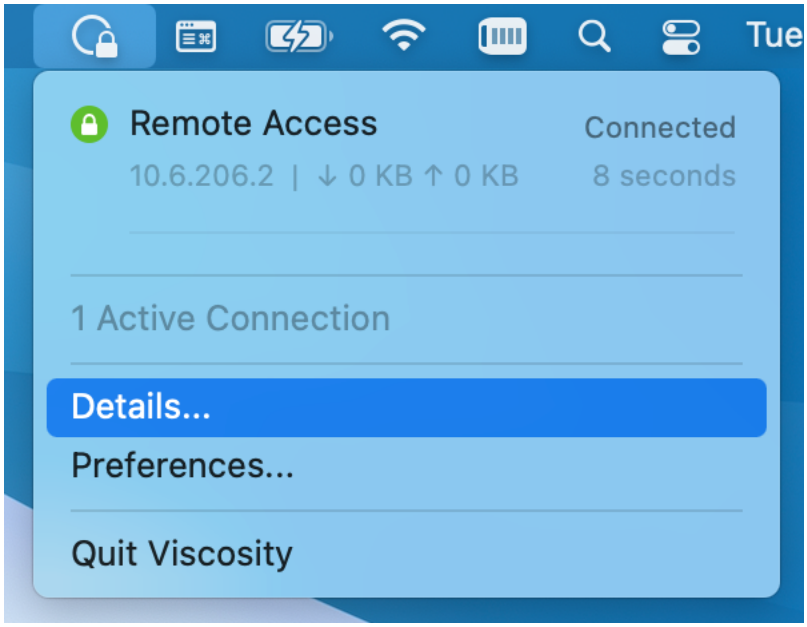


Fig. 77: Viscosity Menu

Basic Status, Bandwidth Graph

The top of the details screen (Figure *Viscosity Details: Bandwidth Graph*) shows the connection status, connected time, the client IP address, and the IP address of the server.

The bottom of the details screen contains a real-time bandwidth graph which shows the current throughput in and out of this OpenVPN connection.

Traffic Counters, Connection Details

The up/down arrow button in the middle of the details screen displays additional network traffic statistics. This function shows the total amount of traffic sent and received by the VPN client.

This screen also contains additional connection information such as DNS Servers assigned to this VPN by the server or local configuration, and the encryption algorithms used by the client to secure communications with the server.

Viscosity Client Logs

The third icon in the middle of the **Details** screen shows the OpenVPN log file (Figure *Viscosity Details: Logs*). If there is any trouble connecting, review the logs here to help determine the problem.

See also:

Troubleshooting OpenVPN.

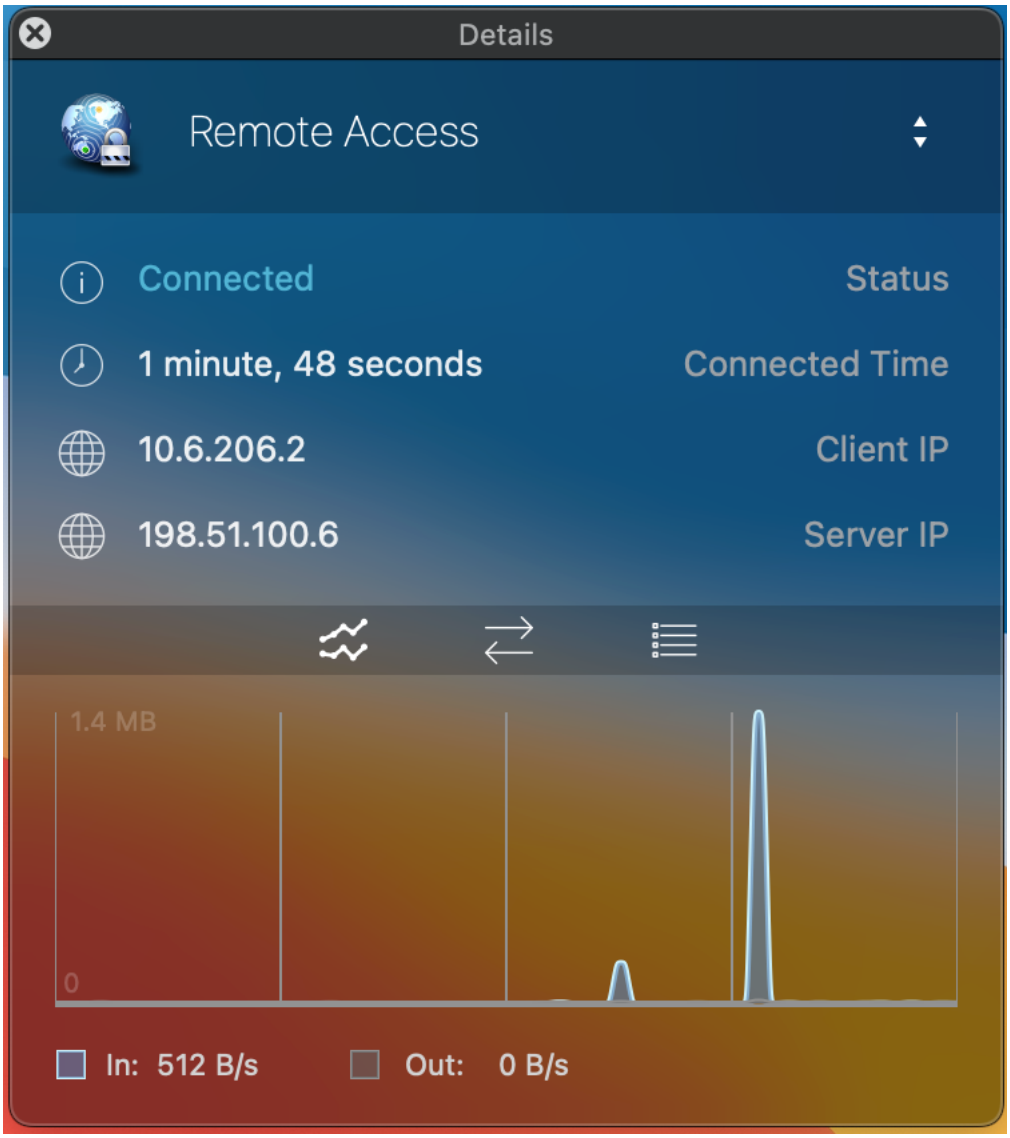


Fig. 78: Viscosity Details: Bandwidth Graph

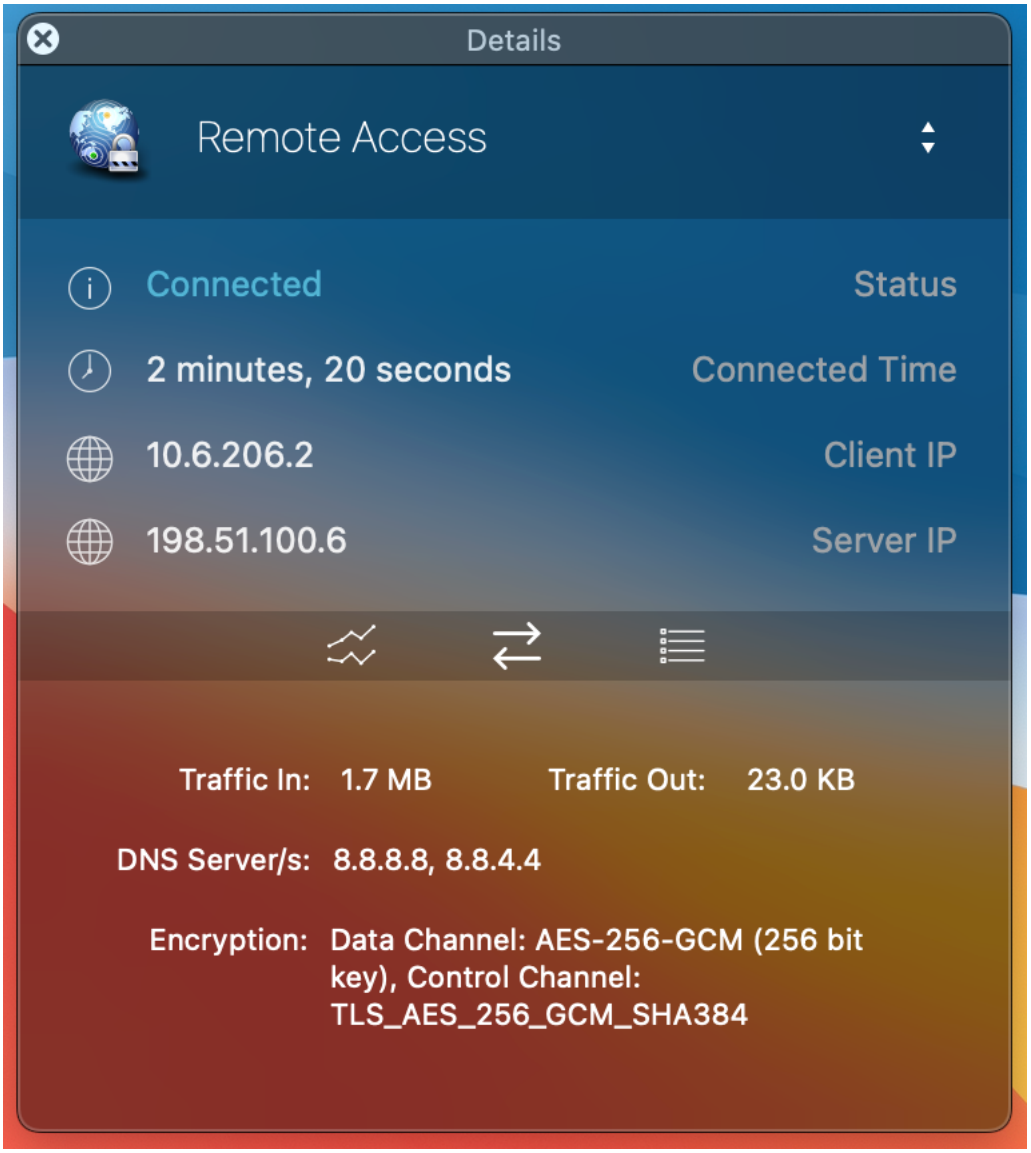


Fig. 79: Viscosity Details: Traffic Statistics

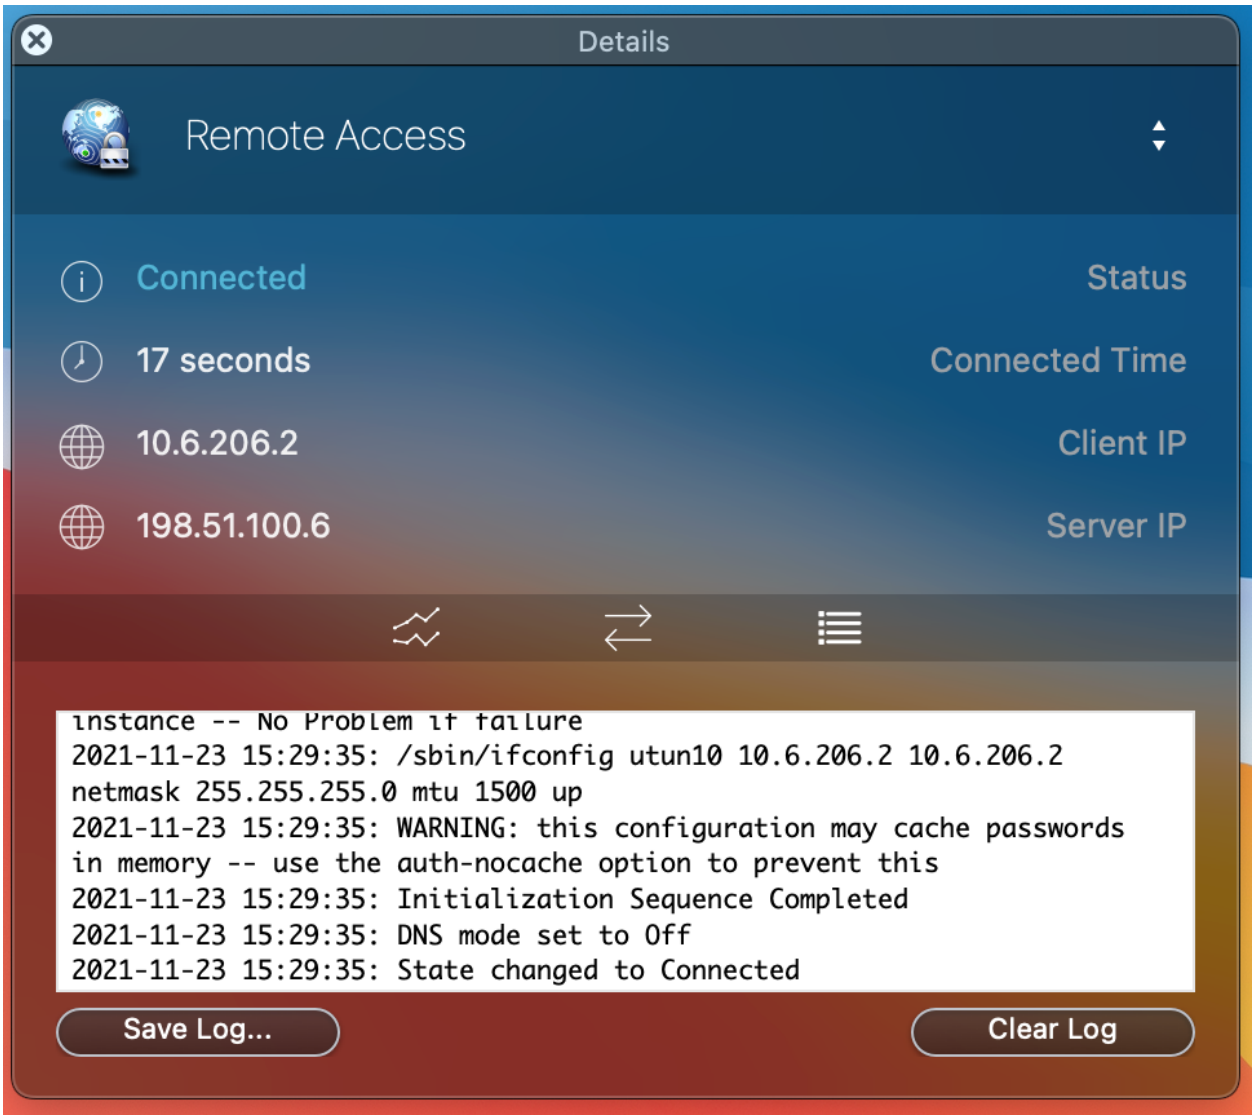


Fig. 80: Viscosity Details: Logs

35.47.3 Installing the OpenVPN Client on iOS

iOS is also capable of running OpenVPN natively using the iOS OpenVPN Connect client available in the App Store. The OpenVPN configuration and certificates must be generated outside the iOS device and then imported to the app.

The *OpenVPN Client Export Package* can export an **OpenVPN Connect** type **Inline Configuration** compatible with this app. Export a configuration file then transfer the resulting `.ovpn` file to the target device using iTunes to transfer the files into the app or e-mail it to the device.

Using other methods to get files onto the device remotely, such as Dropbox, Google Drive, or Box will work similarly to the e-mail method are generally more secure as the contents will remain private and possibly encrypted depending on the method and storage.

If using the e-mail method, use the following procedure:

- Export the **OpenVPN Connect** type **Inline Configuration** file for the VPN.
- Send the exported file in an e-mail to an account configured on the iOS device
- Install the OpenVPN Connect app on the device
- Open the Mail app on the device
- Open the e-mail message containing the attachment
- Tap the attachment. When it is tapped one of the choices will be to open it with the OpenVPN Connect app
- Tap to select the OpenVPN connect app, and it will offer to import the configuration
- Tap the + button and the profile will be imported

Using iTunes to transfer the configuration to the iOS device is simple and more secure than e-mail.

- Export the **OpenVPN Connect** type **Inline Configuration** file for the VPN.
- Connect the iOS device to the computer and open iTunes
- Find and install the OpenVPN Connect app
- Click the device icon inside of iTunes in the toolbar
- Select **Apps** on the left side of the window
- Locate the **File Sharing** section At the bottom of this screen (scroll down)
- Click the icon for OpenVPN under **File Sharing** and a list of files will show on the right under the heading **OpenVPN Documents**
- Copy the file to the device by using ONE of the following methods. The file will be immediately available on the iOS device.
  - Use Finder to drag and drop the `.ovpn` file into this area **OR**
  - Click **Add** and locate the file to import
- Open the OpenVPN Connect app, and it will offer to import the profile
- Tap the + button, and the profile will be imported

If the OpenVPN server requires user authentication the app will prompt for the credentials, which may optionally be saved. Underneath the credential prompt is a connection status which will change between **Disconnected** and **Connected** and also indicates when a connection is being attempted. Clicking this will open the OpenVPN client log which is very useful if connection problems are encountered.

To connect the VPN, move the slider at the bottom of the profile from **Off** to **On** and the app will attempt to connect. To manually disconnect, move the slider back to **Off**.

35.47.4 Installing the OpenVPN Client on Android

For devices running Android there is a free OpenVPN app in the Google Play Store that works excellently without needing root access. It is called OpenVPN for Android by Arne Schwabe.

The *OpenVPN Client Export Package* can export an **Android** type **Inline Configuration** compatible with the app.

- Export an **Android** type inline configuration from the OpenVPN Client Export package
- Transfer this `.ovpn` file to the target device

  It can be copied directly, e-mailed to the device, or by other similar methods of copying files to the device.
- Open the **OpenVPN for Android** app
- Tap Import (File folder icon at upper right)
- Find the `.ovpn` file saved above and tap it
- Tap Import (Disk icon at upper right)

The app shows the imported VPN in the connection list. Edit the entry to change the name and other details.

Tap the VPN to connect. If the OpenVPN server requires user authentication, the app will prompt for credentials when connecting.

Note: The Android OpenVPN Connect client also works on Android and does not require root. It works identically to the iOS client by the same name. It lacks the ability to fully configure the VPN in the GUI, so it is not as convenient to use. Use the **OpenVPN Connect** type **Inline Configuration** export for use with that client on both Android and iOS.

35.47.5 Installing the OpenVPN Client on FreeBSD

If the client has a stock FreeBSD installation, OpenVPN is in the pre-compiled FreeBSD packages repository as well as in the FreeBSD Ports collection.

To install OpenVPN, run the following as the `root` user or via `sudo`:

```
# pkg install openvpn
```

Alternately, compile the client from ports:

```
# cd /usr/ports/security/openvpn && make install clean
```

35.47.6 Installing the OpenVPN Client on Linux

Installing OpenVPN on Linux varies depending on the distribution, method of managing software installations, and network management software on the client device.

OpenVPN is included in the package repositories of most major Linux distributions. With all the various possibilities between countless distributions, and adequate information already available in other sources online, this documentation will not cover specifics. Search the Internet for the distribution of choice and "installing OpenVPN" to find information.

Ubuntu-based distributions have OpenVPN management integrated with Network Manager, but it requires installing an extra module.

35.47.7 Installing the OpenVPN Client Configuration Manually

Performing a manual client installation instead of using the *OpenVPN Client Export Package* requires additional steps to install the software and settings onto the client devices. Installing the client on other operating systems is left up to the reader.

After installing OpenVPN, copy the certificates to the client and create the client configuration file.

Copy certificates

Three files from the firewall are needed for each client: the CA certificate, the client certificate, and the client key. The configuration may require a fourth file, the TLS key, if the server is configured for TLS authentication.

- Export the CA certificate from **System Cert > Manager** on the **CAs** tab, save this as `ca.crt`
- Export the client certificate and key as described in *Local Database*, save these as `username.crt` and `username.key`
- Copy these files to the OpenVPN `config` directory on the client
- Copy the TLS key from the server configuration screen If TLS authentication is used on this OpenVPN server. Save this into a new text file called `tls.key` and include it in the `config` folder as well.

Create Configuration

After copying the certificates to the client, the OpenVPN client configuration file must be created. This can be done with any plain text file editor such as Notepad on Windows. The following shows the options most frequently used:

```
client
dev tun
proto udp
remote vpn.example.com 1194
ping 10
resolv-retry infinite
nobind
persist-key
persist-tun
ca ca.crt
cert username.crt
key username.key
verb 3
comp-lzo
tls-auth tls.key 1
auth-user-pass
```

remote
: The host and port of the remote OpenVPN server. Can be an IP address or FQDN.

proto
: The protocol used by the OpenVPN connection. Change this line to `proto tcp` if the OpenVPN server uses TCP.

ca, cert, key
: Must be modified accordingly for each client to reflect the filenames saved previously.

tls-auth
: If TLS authentication is not used, the `tls-auth` line may be omitted.

auth-user-pass
If the remote access VPN does not include username and password authentication, omit this line.

See also:

For a more complete reference on the OpenVPN directives, refer to the OpenVPN manual.

Distributing configuration and keys to clients

The easiest way to distribute the keys and OpenVPN configuration to clients is via the *OpenVPN Client Export Package*. If that package is not a viable choice, place the needed files in a ZIP archive or self-extracting archive automatically extracting to `C:\Program Files\OpenVPN\config` or the appropriate path for the client in question.

Warning: Transmit this configuration securely to the end user, never allow it to pass over untrusted networks unencrypted.

35.48 Authenticating OpenVPN Users with FreeRADIUS

Using OpenVPN with the *FreeRADIUS package*.

35.48.1 Purpose

This document demonstrates how to set up OpenVPN with RADIUS user authentication provided by the FreeRADIUS package.

The firewall can centrally manage usernames and passwords and this method can also use additional RADIUS-specific options. This is a plus because login times, access limits, and other options are possible.

See also:

FreeRADIUS package Controlling Client Parameters via RADIUS

35.48.2 Requirements

- A working OpenVPN remote access server (*OpenVPN Remote Access Configuration Example*)
- The FreeRADIUS Package (*FreeRADIUS package*)

35.48.3 Add an interface to FreeRADIUS

- Navigate to **Services > FreeRADIUS**, **Interfaces** tab
- Click  **Add** to create a new entry
- Enter the following settings, which may already be the default values:

  Interface IP Address
  * or `127.0.0.1` to bind only to Localhost

  Port
  `1812`

Interface Type
Authentication

IP Version
IPv4

- Click **Save**

35.48.4 Add a NAS client to FreeRADIUS

- Navigate to **Services > FreeRADIUS**, **NAS / Clients** tab
- Click  **Add** to create a new entry
- Enter the following settings:

Client IP Address
`127.0.0.1`

Client Shortname
Enter the firewall hostname (without the domain)

Client Shared Secret
Enter a secure password

Description
`Local firewall authentication` or similar text to identify this entry

- Click **Save**

35.48.5 Add Users

- Navigate to **Services > FreeRADIUS**, **Users** tab.

Note: Manage *every* user which will authenticate with FreeRADIUS/OpenVPN on this tab.

- Click  **Add** to create a new entry
- Enter the following settings:

Username / Password
The credentials for this user.

Number of simultaneous connections
(Optional) The number of active connections this user may have at the same time. Leave empty for no limit.

Session Timeout
(Optional) The amount of time, in seconds, before the user is disconnected and must log in again.

Set any other options as needed to configure or restrict the user in various ways.

See also:

Controlling Client Parameters via RADIUS

- Click **Save**

- Repeat as needed for additional users

35.48.6 Add an Authentication Server

- Navigate to **System > User Manager**, **Authentication Servers** tab
- Click  **Add** to create a new entry
- Enter the following settings:

  Descriptive name
  `Local FreeRADIUS`

  Type
  RADIUS

  Hostname or IP address
  `127.0.0.1`

  Shared Secret
  The password added to the NAS entry in a previous step

  Services offered
  Authentication

  Authentication port
  `1812`

- Click **Save**

35.48.7 Test RADIUS Authentication

- Navigate to **Diagnostics > Authentication**
- Select the newly created authentication server (e.g. *Local FreeRADIUS*)
- Fill in a **Username** and **Password** for a user entry in FreeRADIUS
- Click **Test**

If the test succeeded, continue. Otherwise, see the *Troubleshooting section*.

35.48.8 Configure OpenVPN to use RADIUS

- Navigate to **VPN > OpenVPN**, **Servers** tab
- Edit the existing remote access OpenVPN server
- Set the **Mode** to either **Remote Access (User Auth)** or **Remote Access (SSL/TLS + User Auth)** if it is not already set to one or the other.
- Set **Backend for authentication** to the FreeRADIUS authentication server (e.g. *Local FreeRADIUS*)
- Click **Save**
- Attempt to connect and authenticate with an OpenVPN client

  If the test succeeded, the setup is complete. Otherwise, see the *Troubleshooting section*.

35.48.9 Troubleshooting

The following options can be helpful in troubleshooting FreeRADIUS and OpenVPN. Commands must be run at a shell prompt either via the console or via SSH unless otherwise specified.

See also:

FreeRADIUS package contains additional troubleshooting information.

Increase the verbosity of OpenVPN Logs

- Navigate to **VPN > OpenVPN** and select the server
- Change **Verbosity level** to *7*

  This will log everything from OpenVPN
- Attempt to connect and authenticate with an OpenVPN client
- Navigate to **Status > System Logs**, **OpenVPN** tab to *check the OpenVPN log* for relevant messages

  Alternately, watch the log from an SSH or console shell prompt:

```
# tail -F /var/log/openvpn.log
```

Watch FreeRADIUS Logs

FreeRADIUS can also log attempted connections/authorizations to find potential problems.

- Navigate to **Services > FreeRADIUS**, **Settings** tab
- Set the options under **Logging Configuration** to help locate problems.

  At a minimum, set the following:

  RADIUS Logging Destination
  System Logs

  RADIUS Logging
  Enable

  The remaining options can be left at their defaults but can aid further in debugging if necessary.
- Click **Save**
- Attempt to connect and authenticate with an OpenVPN client
- Navigate to **Status > System Logs** to *check the system log* for relevant messages

  Alternately, watch the log from an SSH or console shell prompt:

```
# tail -F /var/log/system.log
```

Seek Additional Help

If the cause of the problem is not obvious from the logs, use the information gathered in the previous steps to search the web and/or post on the Netgate Forum for assistance.

35.49 Authenticating OpenVPN Users with RADIUS via Active Directory

This recipe describes the procedure to set up OpenVPN on pfSense® software with user authentication handled via RADIUS on an Active Directory server.

35.49.1 Set up the Windows Server

- Set up the Windows Server for an Active Directory role
- Add users to the Windows Server (optionally in a common group for VPN users)
- Set up the NPS role as described in *Authenticating from Active Directory using RADIUS/NPS* which allows the Windows Server to handle RADIUS requests

35.49.2 Add Authentication Server

- Navigate to **System > User Manager**, **Authentication Servers** tab
- Click  **Add** to create a new entry
- Enter the following settings:

  Descriptive name
  `Active Directory NPS`

  Type
  RADIUS

  Hostname or IP address
  `198.51.100.30` – Replace this with the IP address of the Windows server

  Shared Secret
  The password added to the NAS entry in NPS

  Services offered
  Authentication

  Authentication port
  `1812`

- Click **Save**

35.49.3 Setup OpenVPN Remote Access Server

The recipe *OpenVPN Remote Access Configuration Example* covers the OpenVPN server setup, so there is no need to duplicate the instructions here.

Choose the Active Directory NPS RADIUS authentication server entry during the wizard or configure it as the backend for authentication after completing the wizard.

35.49.4 Setup Clients

Use the *OpenVPN Client Export Package* to generate configuration files and/or installation packages for clients.

Clients are available for a wide variety of operating systems, see the installation guides at *Installing OpenVPN Remote Access Clients*.

35.50 Connecting OpenVPN Sites with Conflicting IP Subnets

One common use of NAT with OpenVPN is to mask conflicting LAN subnets between two locations. If two networks are using the exact same subnet, or overlapping subnets, as their LAN or other internal network they cannot communicate across a site-to-site VPN without NAT.

Danger: Shared key mode has been deprecated by OpenVPN as it is no longer considered sufficiently secure for modern requirements.

Shared key mode will be removed from future versions of OpenVPN. Users **should not** create any new shared key tunnels and should **immediately** convert any existing shared key tunnels to SSL/TLS mode.

When an SSL/TLS instance is configured with a `/30` tunnel network it behaves similarly to shared key mode. The primary difference is the need to create and distribute the certificate structure to peers. See *OpenVPN Site-to-Site Configuration Example with SSL/TLS* for information on configuring OpenVPN in SSL/TLS mode.

35.50.1 Site-to Site Example

In this example `10.3.0.0/24` is the LAN on both sides of a VPN. Hosts on the `10.3.0.0/24` subnet will never reach the other end of the VPN to communicate with the remote `10.3.0.0/24` subnet. Clients will always treat that network as local, attempting to reach the other systems via ARP. With NAT, however, the remote side can be made to function as if it were using a different subnet.

Note: Utilizing NAT will work for many protocols but some that are commonly desirable across VPN connections, primarily SMB/CIFS file sharing between Windows hosts, may not function in combination with NAT. If hosts use a protocol which is not capable of functioning with NAT, this is not a viable solution.

Figure *Site to Site with Conflicting Subnets* shows an example where both ends are using the same subnet. After assigning the OpenVPN interface to an OPT interface on both sides, as described in *Assigning OpenVPN Interfaces*, 1:1 NAT can be applied.

The firewall at Site A translates its LAN to `172.16.1.0/24` and the firewall at Site B translates to `172.17.1.0/24`. A 1:1 NAT entry on each firewall translates its entire /24 range.

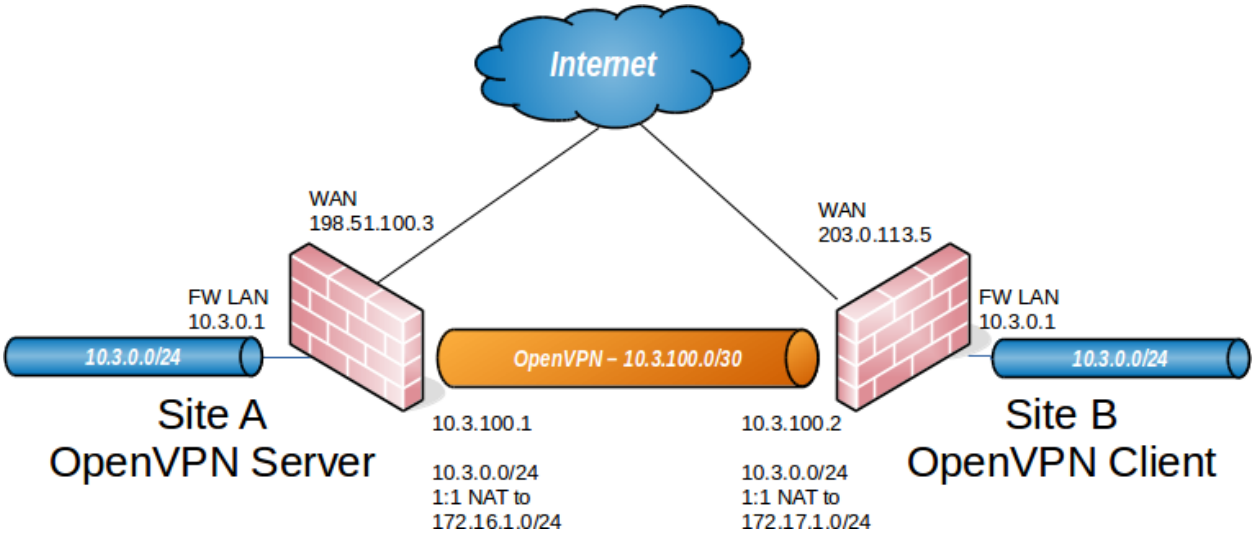


Fig. 81: Site to Site with Conflicting Subnets

To reach Site A from Site B, clients at Site B use `172.16.1.x` IP addresses. The 1:1 NAT rule translates the last octet in the `10.3.0.x` IP address to the last octet in `172.16.1.x`. A client at Site B attempting to reach `10.3.0.10` at Site A would use `172.16.1.10`. To reach `10.3.0.50` at Site B from Site A, a client would use `172.17.1.50`.

Figure *Site B 1:1 NAT Configuration* show the 1:1 NAT configuration for each side, where the *tun* interface is assigned as OPT1.

In the OpenVPN configuration on both sides, the **Remote network** must be the *translated* IP subnet, not `10.3.0.0/24`. In this example, the **Remote Network** at Site A is `172.17.1.0/24`, and Site B is `172.16.1.0/24`.

After applying the NAT configuration changes and configuring the remote networks accordingly on both sides, the networks will be able to communicate using the translated subnets.

35.50.2 Site-to-Multi-Site Example

This section describes how to map multiple subnets that have the same IP address range using OpenVPN so that they can be accessed from a central site. For example `192.168.0/24` is a very common addressing scheme and the main site may need to access all the systems on those networks.

This is the desired outcome, Site 0 is the hub:

```
Site 0: 10.1.1/24
Site 1: 192.168.0/24 -> 10.10.1/24
Site 2: 192.168.0/24 -> 10.10.2/24
Site 3: 192.168.0/24 -> 10.10.3/24
```

For example, with this type of setup in place, site 0 can access `192.168.0.33` at site 2 as `10.10.2.33`.

Note: This configuration requires 1:1 NAT which means that some things may not work, see notes in the remaining sections for details.

There are multiple alternate ways to reach this goal. This example has the remote sites (Sites 1-3) run the "server" and Site 0 runs the "clients". Either direction works.

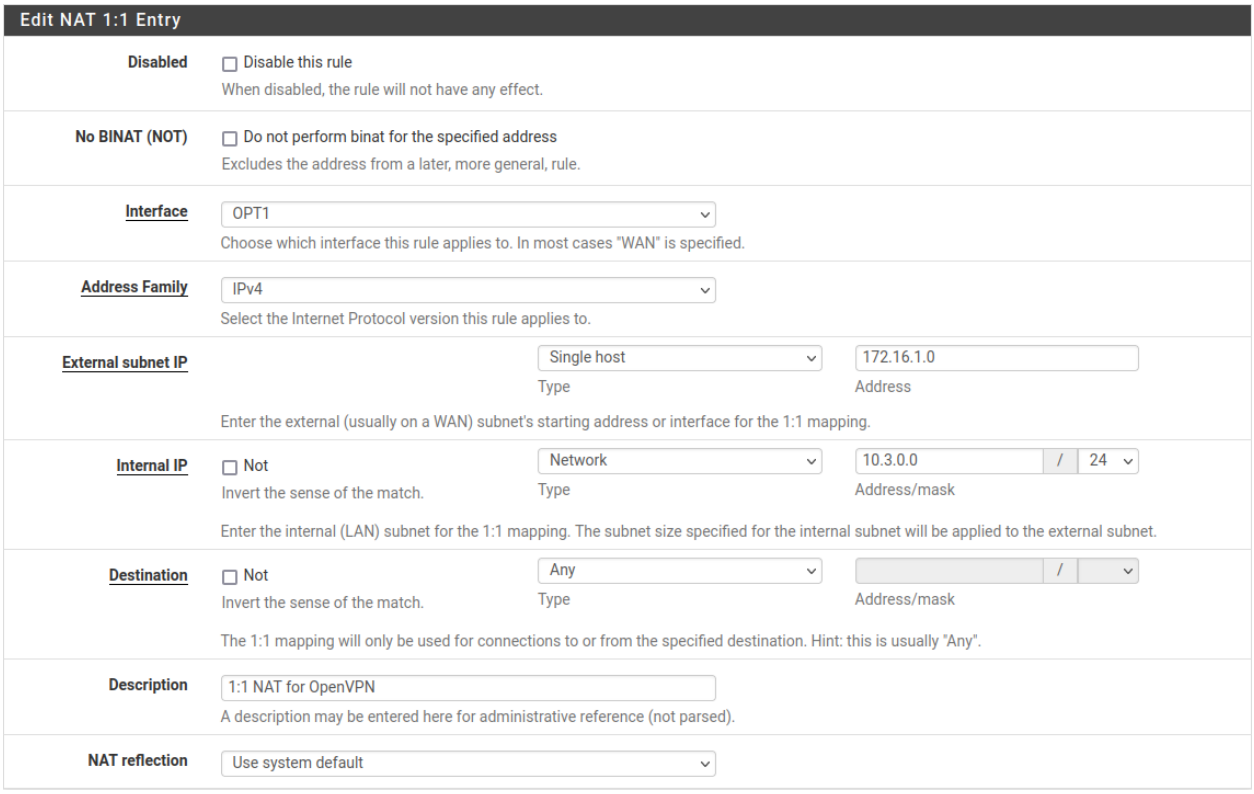


Fig. 82: Site A 1:1 NAT Configuration

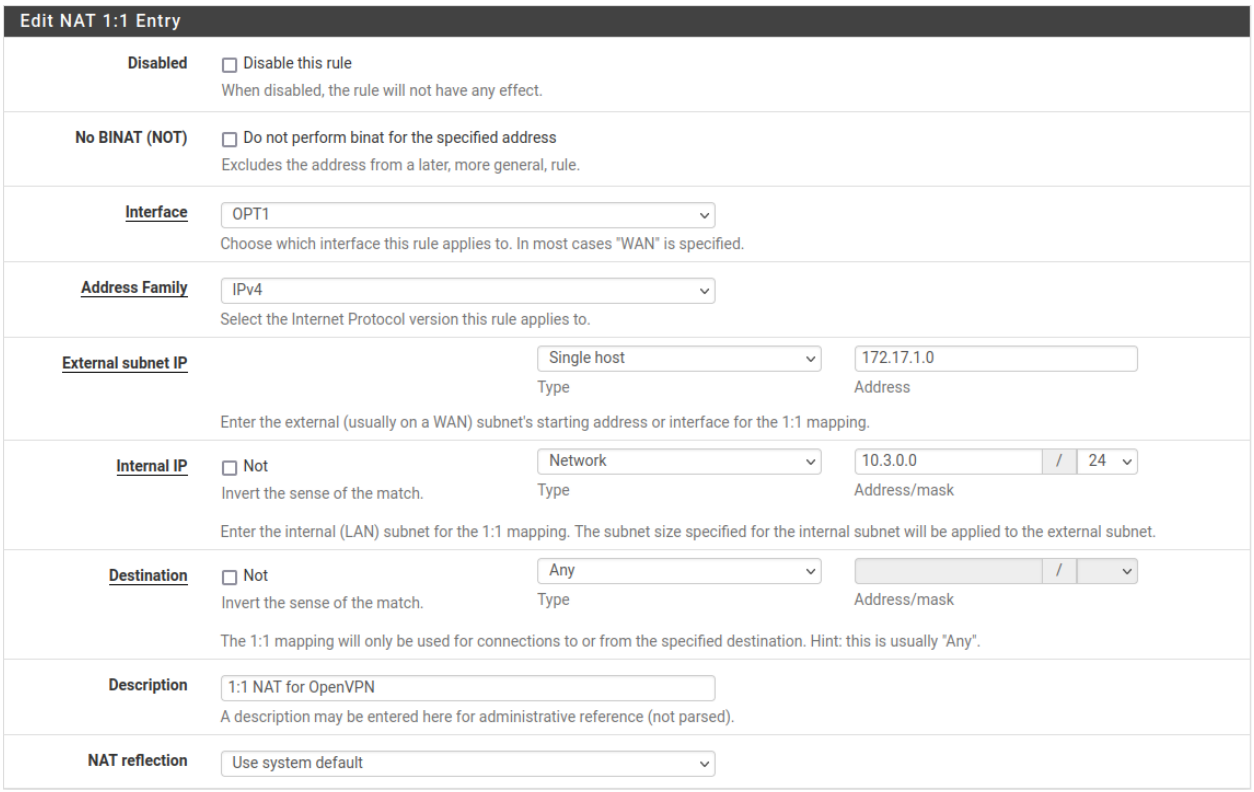


Fig. 83: Site B 1:1 NAT Configuration

Preliminary

Pick **two** ranges for this scheme. The first range is for mapping subnets and the second range is for OpenVPN client to server connections. This example uses `10.10.0.0/16` for the mappings and `10.254.100/24` for the VPN endpoints. The first can contain up to 253 mappings and the second contains 64 `/30` subnets to link it all up.

Examples before NAT:

Site 0 to Site 1:

```
           <--------             (OpenVPN)            ---------->

10.1.1/24 --- 10.254.100.2/30 --- INTERNET --- 10.254.100.1/30 --- 192.168.0/24
```

Site 0 to Site 2:

```
           <--------             (OpenVPN)            ---------->

10.1.1/24 --- 10.254.100.6/30 --- INTERNET --- 10.254.100.5/30 --- 192.168.0/24
```

Note: Notice how the addresses on the right are not unique - they are always `192.168.0.0/24`.

NAT can make each remote site unique:

```
10.1.1/24 --- (OpenVPN) --- INTERNET --- OpenVPN_Interface - NAT - 10.10.1/24 <=> 192.
↪168.0/24
10.1.1/24 --- (OpenVPN) --- INTERNET --- OpenVPN_Interface - NAT - 10.10.2/24 <=> 192.
↪168.0/24
10.1.1/24 --- (OpenVPN) --- INTERNET --- OpenVPN_Interface - NAT - 10.10.3/24 <=> 192.
↪168.0/24
```

Site 0 sends packets to a destination in the mapped subnet down the tunnel which corresponds to the desired Site 1, 2, or 3 and the NAT at the other end translates to and from that mapping and sends the result back to Site 0.

Recipe

In the following examples, only necessary changes from default are given. Pick suitable transports (UDP by default) and ports (1194 by default) Put in a suitable firewall rule on the server WAN interfaces to allow the inbound connection VPN. This should be restricted to the WAN IP address of Site 0. Both ends must have a suitable firewall rule on the OpenVPN interface for traffic to pass.

At Site 1-3

OpenVPN server

At each remote site, create a new OpenVPN server:

Server Mode
: Peer to Peer

Description
: Link to Site 0

TLS authentication
Check "Automatically generate a shared TLS authentication key."

IPv4 Tunnel Network
Link subnet, e.g. `10.254.100.0/30`

IPv4 Remote networks
IP address range(s) at Site 0, e.g. `10.1.1.0/24`

If Site 0 has multiple IP ranges then specify them all in IPv4 Remote networks, comma separated. After saving, edit the VPN instance and copy the shared key to the other end (see below)

1:1 NAT

Add one of these for *each* network range at Site 0:

Interface
OpenVPN

External subnet IP
Mapping subnet, first IP address, e.g. `10.10.1.0`

Internal IP
LAN net

Destination
Network, IP range at Site 0, e.g. `10.1.1.0/24`

Not the default gateway

If this firewall is not the default gateway for the site then use an outbound NAT rule on LAN to ensure that replies from the clients return via the OpenVPN tunnel. Without this, the systems at Sites 1-3 will reply via their default gateway because they will be unaware of the Site 0 network. Another option is to put a suitable route on the site gateway via the LAN address of the OpenVPN system, but this will introduce an asymmetric route and which will potentially break things even more than the double NAT.

At Site 0

OpenVPN client

Create a separate OpenVPN client for each remote subnet (Where examples are given they are for Site 1):

Server mode
Peer to Peer (Shared key)

Server host or address
IP address of the remote site

Description
Text to describe the connection, such as the site name, mapping subnet, and link subnet. For example, `Site 1 10.254.100.0/30 10.10.1.0/24`

Shared Key
Copy from the server for the site link

IPv4 Tunnel Network
Link subnet, e.g. `10.254.100.0/30`

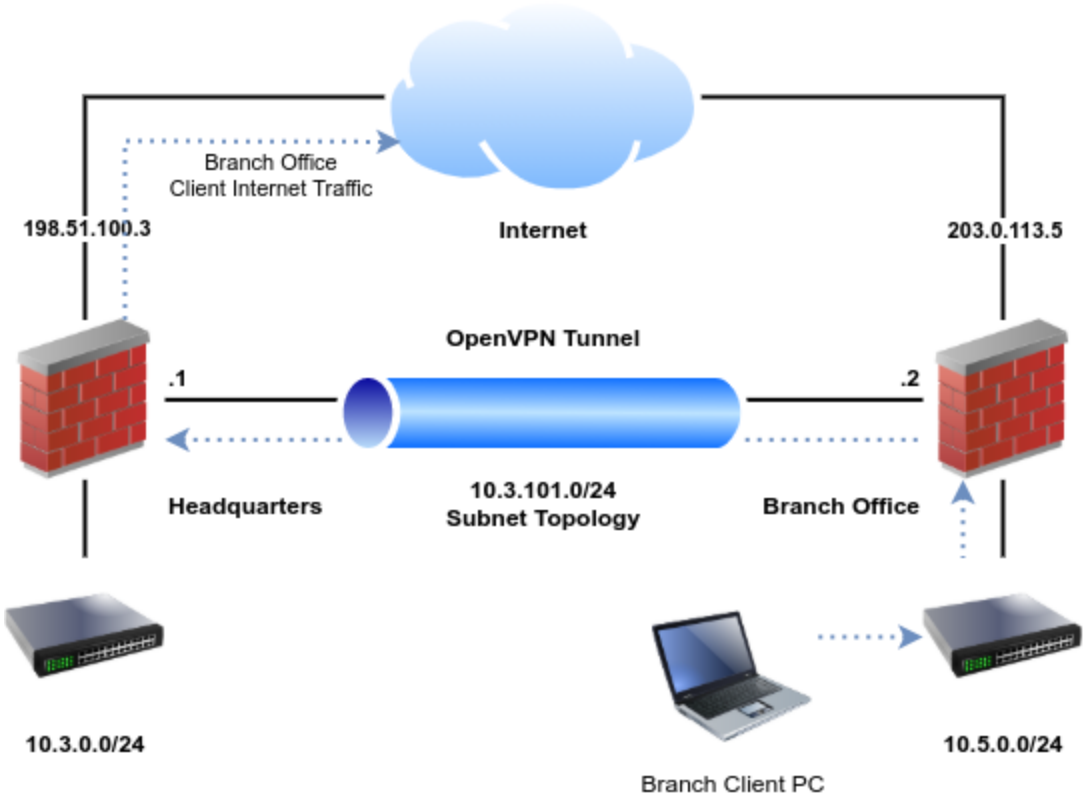
IPv4 Remote networks
Mapping network, e.g. `10.10.1.0/24`

Notes

- SIP and RTP can be tricky with NAT involved
- DNS may not work well under this scheme
- Anything relying on DNS may have issues, such as web links that do not use the passed host name but use the built-in name
- Do not forget to put suitable firewall access rules on the various OpenVPN interfaces

35.51 Routing Internet Traffic Through A Site-To-Site OpenVPN Tunnel

This article shows how route Internet traffic from one site through a second site over OpenVPN on pfSense® software.



See also:

This is similar to using IPsec to accomplish the same task, as described in *Routing Internet Traffic Through a Site-to-Site IPsec Tunnel*

The OpenVPN portion is similar to the basic site-to-site setup detailed in *OpenVPN Site-to-Site Configuration Example with SSL/TLS*.

35.51.1 Design

In this scenario, the headquarters site will be the server and branch offices will be clients. The basic settings are identical to the *Example Configuration Settings* from the SSL/TLS recipe:

Table 20: OpenVPN Endpoint Settings - Headquarters

Headquarters - Server	
WAN Address	198.51.100.3
LAN Subnet	10.3.0.0/24
LAN Address	10.3.0.1
CA Name	S2SCA
Cert CN	serverA
Tunnel Net	10.3.101.0/24

Table 21: OpenVPN Endpoint Settings - Branch Office

Branch Office - Client	
Cert CN	clientB
WAN Address	203.0.113.5
LAN Subnet	10.5.0.0/24
LAN Address	10.5.0.1

35.51.2 OpenVPN Configuration

Set up the entire site-to-site VPN as detailed in *OpenVPN Site-to-Site Configuration Example with SSL/TLS* which will result in a usable base from which the remainder of the settings can be built. That example uses two remote offices, so only the first remote office is necessary here. Configure as many as the use case requires.

35.51.3 Configure outbound NAT

For clients at the remote office to reach the Internet through the headquarters, the headquarters site must perform outbound NAT on the traffic from the remote office LAN (`10.5.0.0/24`) as it leaves the WAN.

Note: While it is also possible to perform NAT on the client side firewall OpenVPN interface, that would negatively impact non-Internet VPN traffic and also results in double NAT for Internet traffic which can cause problems with certain protocols.

To set up outbound NAT on the headquarters firewall, first change the outbound NAT mode:

- Navigate to **Firewall > NAT**, **Outbound** tab on the headquarters firewall
- Set the **Outbound NAT Mode** to **Hybrid Outbound NAT**

  Note: If the headquarters firewall is already on this mode or set to **Manual**, then do not change the mode.

- Click **Save**

Using this mode allows the default automatic NAT rules to continue working without needing a full manual ruleset. Now add a custom rule to the top of the list which matches the remote office LAN:

- Click  **Add**
- Set the following values:

  Source
  Network, `10.5.0.0/24`

  Destination
  Any

  Translation Address
  Interface Address

  Description
  `NAT for OpenVPN remote office`

- Click **Save**
- Click **Apply changes**.

The new entry is now in the outbound NAT rule list.

35.51.4 Configure Client OpenVPN Gateway Behavior

The OpenVPN server can optionally instruct the client to send all of its Internet traffic over the VPN, including traffic from the client firewall itself. This can be convenient as it is less to configure, but it does not lend itself well to selective control over the traffic which crosses the VPN.

Note: If the use case only requires redirecting LAN client traffic via policy routing then skip this step.

To configure this:

- Navigate to **VPN > OpenVPN**, **Servers** tab on the headquarters firewall
- Edit the OpenVPN server instance
- Check **Redirect IPv4 Gateway**
- Click **Save**

Next time the client connects, OpenVPN will automatically set the default gateway for the firewall to the VPN server while it is connected.

35.51.5 Assign OpenVPN Interfaces

On the server and every client, assign this OpenVPN instance as an interface, following the procedure in *Assigning OpenVPN Interfaces*.

This will set up the necessary gateways and routing behavior which take care of some parts of the setup automatically. The gateways can be used for policy routing and failover to selectively handle client traffic in different ways.

35.51.6 Setup Gateway Groups

Note: This step is optional if the use case does not call for policy routing or failover.

Assigned OpenVPN interfaces automatically create gateways which can be used for gateway groups and policy routing. These are located under **System > Routing**.

Using one of these gateways a gateway group can enable several different types of scenarios:

- Send client traffic over the VPN, allow it to exit the Internet directly if the VPN is down.
- If there are multiple VPNs, client traffic could fail between them, or even load balance connections between them.
- Failover between the VPN and other types of private circuits or higher cost connections if it fails.

Create a gateway group as described in *Gateway Groups* using whichever parameters best fit the use case.

35.51.7 OpenVPN Firewall Rules

Since this tunnel must pass traffic from the Internet, the firewall rules must be fairly lenient. The rules at the headquarters site will need to pass traffic from a source of the remote office LAN (`10.5.0.0/24`) to a destination of *any*.

These firewall rules should be placed on the assigned OpenVPN interface tab where possible, and not on the OpenVPN tab of the firewall rules. This helps ensure proper routing and return routing.

Tip: To prevent the remote office from reaching sensitive local resources at headquarters or sites connected to additional VPNs, place block rules for those sensitive destinations above the rule passing the Internet traffic.

The rules on the remote office firewall do not necessarily have to allow much traffic back through unless there are public resources at the remote office which must be reached across the tunnel (e.g. 1:1 NAT, port forwards).

35.51.8 Setup Policy Routing

Policy routing (*Policy Routing Configuration*) allows the firewall to selectively match and route client traffic over the VPN that otherwise would follow the default routing table when exiting the firewall. For example, firewall rules could match only HTTP and HTTPS traffic and send them across the VPN. The other traffic would exit the remote office Internet connection directly.

This gives firewall administrators fine-grained control over traffic flow, rather than directly all traffic over the VPN. The downside is that it doesn't control traffic from the firewall itself, only traffic coming from other interfaces on the firewall.

Note: These policy routing rules go on the local interfaces which contain the clients that initiate traffic, such as LAN. They do not go on VPN interfaces.

In this example, the firewall will route all traffic from hosts on the LAN across the VPN:

- Navigate to **Firewall Rules**, **LAN** tab on the remote office firewall
- Edit the default rule which matches LAN traffic (e.g. `Default allow LAN to any rule`)
- Click **Display Advanced**

- Set the **Gateway** to the assigned OpenVPN interface gateway, or a suitable gateway group.
- Click **Save**
- Click **Apply Changes**

Note: If this rule is set to both IPv4 and IPv6, this change may not work. Duplicate the rule and set one to only IPv4 and one to only IPv6, then set appropriate gateways on each as needed.

35.51.9 Test the Configuration

From this point, new connections made from the remote office LAN will be sent over the VPN. Open a site to check the client IP address, such as doing a Google search for `what is my ip address`. The result should be the WAN IP address of the headquarters firewall.

If it is not, then check the LAN firewall rules, the headquarters OpenVPN firewall rules, and the headquarters outbound NAT rules.

35.51.10 Bonus Topics

Adding More Clients

The example in *OpenVPN Site-to-Site Configuration Example with SSL/TLS* already shows how to accommodate multiple clients, but for this type of setup it takes a couple more steps:

- On the headquarters firewall, add more outbound NAT rules to cover the new client LAN subnets.
- On the headquarters firewall, ensure the OpenVPN interface firewall rules allow traffic from the new client LAN subnets.
- On the remote office firewall, assign the OpenVPN interface.

  Since there is only one server, there is no need to do this again on the headquarters firewall.
- On the remote office firewall, setup gateway groups and policy routing.

Note: This assumes the setup was based on SSL/TLS in client/server mode. If the setup used a /30 tunnel network, that will require an additional server and an additional interface assignment at headquarters.

Port Forwarding or 1:1 NAT to hosts at the Remote Office

It is possible to forward traffic initiated by hosts on the Internet to a server at the remote office in a couple different ways. This can be useful, for example, if a server was relocated, but there are still outdated DNS records or links pointing to the old location.

This can work in one of two ways:

- If the client gateway is redirected using the OpenVPN server option, there is no need for additional configuration. Add the port forwards or 1:1 NAT and it will work as-is.
- If the OpenVPN interfaces are assigned, this can work by allowing the `reply-to` keyword in PF to handle return routing. The remainder of this section covers this process.

For the `reply-to` method to work, a few points must be followed:

- The OpenVPN interface on the remote office must be assigned
- The firewall rules on the **OpenVPN** tab at the remote office **must not** match this traffic
- The firewall rules on the assigned OpenVPN interface tab **must** match this traffic

With the appropriate configuration in place on the remote office firewall, then port forwards or 1:1 NAT on the headquarters site can be set to a destination on the remote office LAN.

35.52 Bridging OpenVPN Connections to Local Networks

The examples in most other OpenVPN recipes are routed using *tun* interfaces which operate at layer 3 and are generally the best practice. OpenVPN also offers the option of using *tap* interfaces, which operate at layer 2 and allow bridging clients directly onto the LAN or other internal network. This can make the remote clients appear to be on the local LAN.

See also:

See *Device Mode* for information on differences between *tun* and *tap* interfaces.

35.52.1 OpenVPN Server Settings

Most of the settings for a bridged remote access VPN are the same as for a traditional remote access VPN (*OpenVPN Remote Access Configuration Example*). The differences are noted here.

Device Mode
tap

A bridged connection requires a **Device Mode** of *tap*.

Tunnel Network
Empty

Remove values from the IPv4 Tunnel Network and IPv6 Tunnel Network boxes so they are empty. The way a *tap* bridge OpenVPN functions it does not need a tunnel network as OpenVPN does not use the same address assignment techniques that it does for *tun* mode.

Bridge DHCP
When selected, OpenVPN passes DHCP through to the bridged interface configured later. In the most common scenario, this is *LAN*. Using this method connecting clients would receive IP addresses from the same DHCP pool used by directly wired LAN clients.

Bridge Interface
LAN

> **Warning:** This setting does not create the bridge, it only indicates to OpenVPN which interface will be a member of the bridge.

This controls which existing IP address and subnet mask OpenVPN will use for the bridge. Setting this to *none* will cause the **Server Bridge DHCP** settings below to be ignored.

Server Bridge DHCP Start/End
When using *tap* mode as a multipoint server, a DHCP range may optionally be configured to use on the interface to which this *tap* instance is bridged.

If these settings are left blank, OpenVPN will pass DHCP through to the bridge interface, and it will ignore the interface setting above. This allows administrators to set aside a range of IP addresses

for use only by OpenVPN clients so they may be contained within a portion of the internal network rather than consuming IP addresses from the existing DHCP pool. Enter the **Server Bridge DHCP Start** and **Server Bridge DHCP End** IP address values as needed.

35.52.2 Creating the Bridge

Once the OpenVPN *tap* server has been created, the OpenVPN interface must be assigned and bridged to the internal interface.

Assign OpenVPN interface

The VPN interface must be assigned before it can become a bridge member. The procedure for assigning an OpenVPN interface is covered in *Assigning OpenVPN Interfaces*.

Create Bridge

Once the VPN interface has been assigned, create the bridge as follows:

- Navigate to **Interfaces > Assignments**, **Bridges** tab
- Click  **Add** to create a bridge
- Ctrl-click both the VPN interface and the interface to which it will be bridged (e.g. *LAN*)
- Click **Save**

See also:

More information on bridging can be found in *Bridging*.

35.52.3 Connect with Clients

Clients connecting to the VPN must also be set to use *tap* mode. Once that has been set, connect with a client such as one exported using the OpenVPN Client Export package. The clients will receive an IP address inside the internal subnet as if they were on the LAN.

Note: Bridged OpenVPN clients also receive broadcast and multicast traffic which can greatly increase the amount of traffic passing over the VPN.

35.53 OpenVPN Site-to-Site with Multi-WAN and OSPF

One way to configure a redundant OpenVPN deployment that uses multiple WAN interfaces for failover is by using a dynamic routing protocol such as OSPF.

See also:

- *Open Shortest Path First v2 (OSPF)*

35.53.1 Design

The general layout of this scenario is in Figure *Example OpenVPN Setup Involving OSPF Across Multiple WANs*. In this example, both sites have multiple WANs.

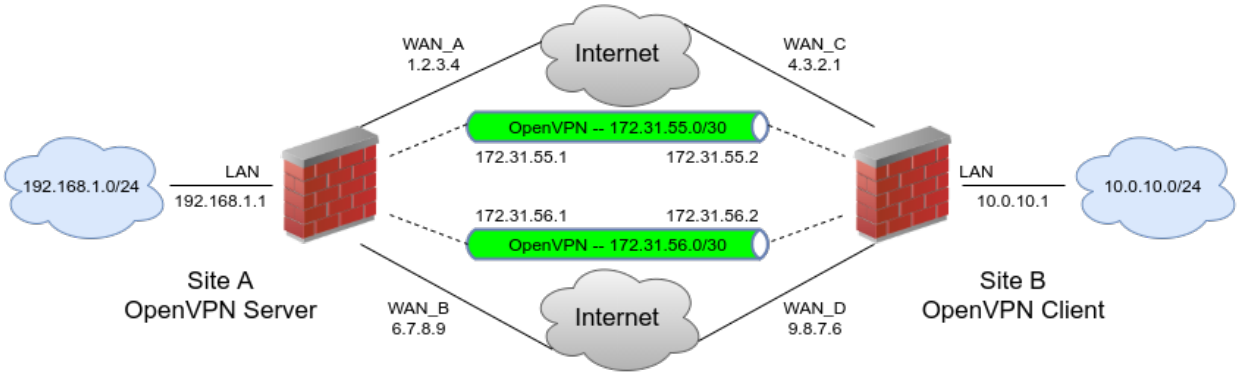


Fig. 84: Example OpenVPN Setup Involving OSPF Across Multiple WANs

35.53.2 OpenVPN Configuration

First, set up a site-to-site OpenVPN instance on each WAN for the remote sites using SSL/TLS with a /30 subnet (*OpenVPN Site-to-Site Configuration Example with SSL/TLS*).

- **Do not fill in** the **Remote Networks** fields on any of the server or client instances, only fill in **Tunnel Network** addresses
- Setup two servers on the local side, each on a different WAN and port

  Use two distinct, non-overlapping tunnel networks (e.g. `172.31.55.0/30` and `172.31.56.0/30`.)
- Setup two clients on the remote firewall, each paired up with one of the above servers, matching the IP addresses and port numbers involved
- Ensure the clients are set for their specific WAN

  Choose the interface from the drop-down menu or a CARP VIP on one of the WANs.
- Ensure these OpenVPN connections link up between client and server

  The tunnel address on both sides will respond to a ping when they are working correctly and appropriate firewall rules are in place. If the tunnels do not establish, see *Troubleshooting OpenVPN* for suggestions on troubleshooting the connection.
- Ensure the OpenVPN firewall rules allow all traffic or at least allow OSPF traffic from a source of the tunnel networks to a destination of any

  The destination on the traffic will be a multicast address, which firewall rules can use to filter specifically if needed, but there isn't much to be gained in the way of security if the source is locked down in the rules as the traffic cannot leave that segment.

35.53.3 FRR OSPF Configuration

Once both instances are connected, configure OSPF.

- Install the FRR package from **System > Packages**, **Available Packages** tab **on both firewalls**
- Navigate to **Services > FRR OSPF**, **Interfaces tab**
- Add each OpenVPN interface

  Set the cost to `10` on the primary link and `20` on the secondary, and so on.
- Add the LAN and other internal interfaces as **passive** interfaces
- Navigate to the **[Global Settings]** tab
- Enter a Master Password

  The content of the password does not matter significantly; FRR uses it internally, it is not used by administrators.
- Set the Router ID to an IP-address-like value, (e.g. `10.3.0.1`)

  The Router ID is unique on each device, which is why setting it to the LAN IP address of a router is a good practice.
- Set the Area ID which is also an IP-address-like value

  The Area ID is typically set to `0.0.0.0` or `0.0.0.1`, but any properly formatted value may be used. The Area ID is the **same** for **all** routers involved in this VPN
- Click **Save**

35.53.4 Validation and Testing

Once OSPF is active on all routers, they will attempt to form a neighbor relationship.

The **Status** tab in FRR will show a full peering with each instance on each wan if they connected properly. It will also list the routes obtained via OSPF. Once that happens, try unplugging and reconnecting WANs and refreshing the status while running test traffic across the VPN, such as an ICMP ping.

35.54 WireGuard Remote Access VPN Configuration Example

This recipe covers configuring a basic *WireGuard* remote access style VPN tunnel.

Note: Though WireGuard does not have a concept of "Client" and "Server" per se, in this style of deployment the firewall cannot initiate connections to remote peers. In this way the firewall acts like a "Server" and may be referred to as such in this documentation. Remote peers may also be referred to as "clients".

35.54.1 Required Information

The following basic information must be determined before starting the VPN configuration.

Item	Value
Design	Remote access, one tunnel+many peers
Firewall WAN	`198.51.100.6`
Listen Port	`51820`
Tunnel Subnet	`10.6.210.0/24`
Tunnel Address	`10.6.210.1/24`
Peer Addresses	`10.6.210.2` - `10.6.210.254`
Peer Endpoints	Dynamic

Generating Keys

WireGuard requires public/private key pairs for each peer, including this firewall.

Warning: Keys cannot be reused between clients, as WireGuard requires unique keys to identify clients and where to send their traffic.

Tunnel Keys

To generate keys for the firewall itself, click the **Generate** button when configuring a tunnel. The GUI will populate the private and public key fields automatically.

The peers will need the public key for their configuration.

Peer Keys

Each peer will need its own public/private key pair. The private key will be needed on the peer client software while the public key will be needed on the firewall itself for the peer definition.

These keys can be generated by the clients themselves, or via command line on a system which has the WireGuard utilities installed. This includes the firewall itself; these commands may be run from a console or SSH shell or from **Diagnostics > Command Prompt**.

From a command line, execute the following:

```
$ wg genkey | tee privatekey | wg pubkey > publickey
```

This command outputs files named `privatekey` and `publickey` which respectively contain a private key and its associated public key. This key pair can be used for a WireGuard peer.

To view the keys, inspect the contents of the files:

```
$ cat privatekey
WGpL3/ejM5L9ngLoAtXkSP1QTNp4eSD34Zh6/Jfni1Q=
$ cat publickey
b9FjbupGC7fomO5U4jL5Irt1ZV5rq4c+utGKj53HXgU=
```

Repeat the commands as needed as many times as is necessary for the number of peers required by this tunnel. Note the keys in a secure place.

Tip: Change the commands to output files named for their associated peer, then store the resulting files in a secure location.

Alternately, the keys can be output in one command without storing them persistently. This behavior is not available on all platforms, but is available on the firewall itself.

```
$ wg genkey | tee /dev/stderr | wg pubkey
4BSH81zC3/OWl25XrzqWy7WnAiARXySHd+K+KFxNrWU=
rzWOC0zH9v2zF6r92uCbjs7JOmhqy8N+cUdA+GCynSM=
```

35.54.2 Tunnel Configuration

Now it's time to create the WireGuard tunnel.

- Navigate to **VPN > WireGuard > Tunnels**
- Click  **Add Tunnel**
- Fill in the options using the information determined earlier:

  Enable
  Checked

  Description
  `Remote Access`

  Listen Port
  `51820`

  Interface Keys
  Click **Generate** to create a new set of keys.

  Interface Addresses
  `10.6.210.1/24`

- Click **Save**

35.54.3 Peer Configuration

Peers can be added when editing a tunnel. To edit a tunnel:

- Navigate to **VPN > WireGuard > Peers**
- Click  **Add Peer**
- Fill in the options using the information determined earlier:

  Enable
  Checked

  Tunnel
  tun_wg<num> (Remote Access)

Description
The name of this client (e.g. The name of a person, device, username, or other uniquely identifying information.)

Dynamic Endpoint
Checked

Keep Alive
Typically left blank, but may be filled in if clients have problems traversing certain firewalls.

Public Key
The public key for this peer. Obtained from the key generation process earlier, or from the peer itself if it was generated by client software directly.

Pre-Shared Key
Not used in this example, but for additional security this pre-shared key can be generated and copied to the peer. Must match on the client and server.

Allowed IPs
The tunnel IP address for this peer, from the list determined above, with a `/32` CIDR mask. For example, the first peer will be `10.6.210.2/32`, the second will be `10.6.210.3/32`, and so on.

- Click **Save Peer**
- Repeat the steps to add additional peers as needed.

35.54.4 Firewall Rules

First add a rule to pass external WireGuard traffic on the WAN:

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  **Add** to add a new rule to the top of the list
- Use the following settings:

Action
Pass

Interface
WAN

Protocol
UDP

Source
any

Destination
WAN Address

Destination Port Range
(other), `51820`

Description
`Pass traffic to WireGuard`

- Click **Save**
- Click **Apply Changes**

Next, add a rule to pass traffic inside the WireGuard tunnel:

- Navigate to **Firewall > Rules**, **WireGuard** tab
- Click  **Add** to add a new rule to the top of the list
- Use the following settings:

  Action
  Pass

  Interface
  WireGuard

  Protocol
  Any

  Source
  any

  Destination
  any

  Description
  `Pass VPN traffic from WireGuard peers`

- Click **Save**
- Click **Apply Changes**

35.54.5 Client Configuration

Client configuration varies by platform, see WireGuard documentation for details. This section covers a basic configuration.

This is an example configuration from a WireGuard client for a split-tunnel configuration:

```
[Interface]
PrivateKey = WGpL3/ejM5L9ngLoAtXkSP1QTNp4eSD34Zh6/Jfni1Q=
ListenPort = 51820
Address = 10.6.210.2/24

[Peer]
PublicKey = PUVBJ+zuz/0mRPEB4tIaVbet5NzVwdWMX7crGx+/wDs=
AllowedIPs = 10.6.210.1/32, 10.6.0.0/24
Endpoint = 198.51.100.6:51820
```

This is an example configuration from a WireGuard client for a full-tunnel configuration:

```
[Interface]
PrivateKey = WGpL3/ejM5L9ngLoAtXkSP1QTNp4eSD34Zh6/Jfni1Q=
ListenPort = 51820
DNS = 10.6.210.1, pfSense.home.arpa
Address = 10.6.210.2/24

[Peer]
PublicKey = PUVBJ+zuz/0mRPEB4tIaVbet5NzVwdWMX7crGx+/wDs=
AllowedIPs = 0.0.0.0/0
Endpoint = 198.51.100.6:51820
```

The fields in that file are as follows:

Interface
Settings for this client.

PrivateKey
The private key for this peer. Obtained from the key generation process earlier, or from the peer itself if it was generated by client software directly.

ListenPort
A static port to listen on, or omit the line to use a random port instead.

DNS
The DNS server(s) and search domain that should be used by the system when the tunnel is enabled.

Address
The tunnel address for this client. Not compatible with all platforms, as some require configuring the address using command-line utilities. However, clients on Windows and Android, for example, allow this directive.

This should use the same CIDR mask as the **Tunnel** address. In this example, the first peer is `10.6.210.2/24`.

Peer
Configuration for the firewall end of the tunnel.

PublicKey
The public key from the **Tunnel** configuration on the firewall.

AllowedIPs
The **Tunnel** address, and any additional networks which should be routed across the VPN in a comma-separated list. This could be a LAN subnet (e.g. `10.6.0.0/24`) or use `0.0.0.0/0` to route all traffic, including Internet traffic, across the tunnel.

Dynamic Endpoint
Unchecked

Endpoint
The firewall WAN IP address and WireGuard **Listen Port**

Note: This only covers the basics, there are numerous other fields which can be used to control client behavior plus additional client options which vary by platform. For additional details, see the WireGuard documentation and the documentation for the WireGuard software used by a peer.

Transfer the resulting client configuration file to the peer securely. Methods vary by platform and client software.

35.54.6 Finish Up

After configuring the client and activating the VPN, the client should be able to pass traffic to the networks listed in the `AllowedIPs` list in its configuration.

See also:

- *WireGuard*
- *Routing*
- *WireGuard Site-to-Site VPN Configuration Example*

- *WireGuard Site-to-Multisite VPN Configuration Example*
- *WireGuard VPN Client Configuration Example*

35.55 WireGuard Site-to-Site VPN Configuration Example

This recipe explains how to set up a VPN tunnel between two firewalls using *WireGuard*.

This example is a minimal configuration, more complicated scenarios are possible, see *WireGuard* for details.

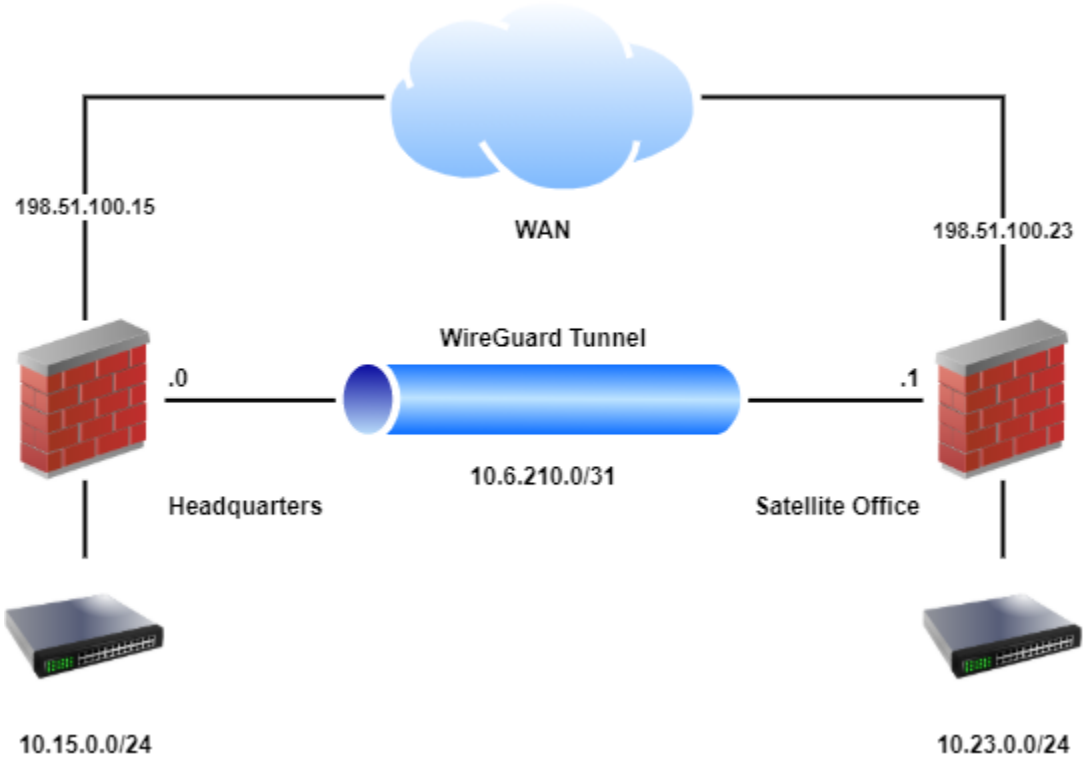


Fig. 85: WireGuard Example Site-to-Site Network

35.55.1 Required Information

General Values

Item	Value
Design	Site-to-Site, one peer per tunnel
Tunnel Subnet	`10.6.210.0/31`

HQ:

Item	Value
WAN IP Address	`198.51.100.15`
Tunnel Address	`10.6.210.0/31`
Listen Port	`51820`
LAN Subnet	`10.15.0.0/24`

Satellite Office:

Item	Value
WAN IP Address	`198.51.100.23`
Tunnel Address	`10.6.210.1/31`
Listen Port	`51820`
LAN Subnet	`10.23.0.0/24`

35.55.2 WireGuard Configuration

- Navigate to **VPN > WireGuard > Settings**
- Fill in the following options:

  Enable
  Checked

  Interface Group Membership
  Only Unassigned Tunnels

- Click **Save**

Tip: When allowing inbound connections from arbitrary remote networks, use rules only on assigned WireGuard interface tabs only to ensure proper return routing.

Note: Rules on assigned WireGuard interface tabs get `reply-to` which ensures that traffic entering a specific assigned WireGuard interface exits back out the same interface. Without that, return traffic will follow the default gateway.

35.55.3 Tunnel Configuration

First create the WireGuard tunnel on both sites:

- Navigate to **VPN > WireGuard > Tunnels**
- Click  **Add Tunnel**
- Fill in the options using the information determined earlier, with variations noted for each site:

  Enabled
  Checked

  HQ Settings

Description
`Satellite Office VPN`

Satellite Office Settings

Description
`HQ VPN`

Listen Port
`51820`

Interface Keys
Click **Generate** to create a new set of keys.

- Copy the public key from each firewall and note which is which
- Click **Save**

35.55.4 Peer Configuration

The peer entry for the server can be added when editing the tunnel. Follow these steps on both sites, with the differences in settings noted inline.

Edit the tunnel:

- Navigate to **VPN > WireGuard > Tunnels**
- Locate the WireGuard tunnel for this VPN
- Click  at the end of the row for the tunnel

From the tunnel editing page, add a peer:

- Click  **Add Peer**
- Fill in the options using the information determined earlier, with variations noted for each site:

HQ Settings

Description
`Satellite Office Peer`

Dynamic Endpoint
Unchecked

Endpoint
`198.51.100.23` (the WAN IP address of the Satellite Office)

Endpoint Port
`51820`

Public Key
The public key *from the Satellite Office firewall*

Allowed IPs
`10.6.210.0/31` and `10.23.0.0/24` (Tunnel network and Satellite Office LAN)

Satellite Office Settings

Description
`HQ VPN Peer`

Dynamic Endpoint
Unchecked

Endpoint
`198.51.100.15` (the WAN IP address of HQ)

Endpoint Port
`51820`

Public Key
The public key *from the HQ firewall*

Allowed IPs
`10.6.210.0/31` and `10.15.0.0/24` (Tunnel network and HQ LAN)

- Click **Save Peer**

35.55.5 Assign Interface

These steps should be done on both sites.

First, fix the default gateway so WireGuard isn't automatically selected before it's ready:

- Navigate to **System > Routing**
- Set **Default Gateway IPv4** to a specific gateway (e.g. *WANGW*) or group
- Set **Default Gateway IPv6** similarly if this VPN will also carry IPv6 traffic
- Click **Save**
- Click **Apply Changes**

Next, assign the interface (*Assign a WireGuard Interface*):

- Navigate to **Interfaces > Assignments**
- Select the appropriate `tun_wg<number>` interface in the **Available network ports** list
- Click  **Add** to assign the interface as a new OPT interface (e.g. OPT1)
- Navigate to the Interface configuration page, **Interfaces > OPTx**
- Check **Enable**
- Enter an appropriate **Description** which will become the interface name (e.g. `VPN_HQ` or `VPN_SATELLITE`)
- Configure an appropriate **MTU** value for the WireGuard interface (e.g. `1420` for IPv4+IPv6 or `1440` for IPv4 only).

  For details on calculating the correct MTU, see in *Assign a WireGuard Interface*.
- Fill in the options for the **HQ** endpoint using the information determined earlier:

  IPv4 Configuration Type
  `Static IPV4`

  IPv4 Address
  `10.6.210.0/31`

  IPv4 Upstream Gateway

  - Click **Add a new gateway**
  - Fill in the options:

Gateway Name
`WG_VPN_SAT_V4`

Gateway IPv4
`10.6.210.1`

– Click  **Add**

- Fill in the options for the **Satellite Office** endpoint using the information determined earlier:

IPv4 Configuration Type
`Static IPV4`

IPv4 Address
`10.6.210.1/31`

IPv4 Upstream Gateway

– Click **Add a new gateway**

– Fill in the options:

Gateway Name
`WG_VPN_HQ_V4`

Gateway IPv4
`10.6.210.0`

– Click  **Add**

- Click **Save**
- Click **Apply Changes**

35.55.6 Firewall Rules

First, add a rule to the WAN on both firewalls to allow traffic to reach WireGuard:

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  **Add** to create a new firewall rule at the top of the list so that it matches before other rules
- Configure the firewall rule as follows:

Action
Pass

Protocol
UDP

Source
This can typically be left at *Any*, but it is more secure to fill in the IP address of the opposing firewall.

Destination
WAN Address

Destination Port Range
(other), `51820`

Description
`Pass traffic to WireGuard`

- Click **Save**
- Click **Apply Changes**

Next, add a rule to pass traffic inside the WireGuard tunnel on both firewalls:

- Navigate to **Firewall > Rules**
- Click the tab for the assigned WireGuard interface (e.g. **VPN_SATELLITE** or **VPN_HQ**)
- Click  **Add** to add a new rule to the top of the list
- Use the following settings:

Action
Pass

Protocol
Any

Source
any

Destination
any

Description
`Pass VPN traffic from WireGuard peers`

Note: This rule allows all traffic between sites, which is easy but not a secure practice. Traffic between the sites can be restricted as needed with less permissive rules.

- Click **Save**
- Click **Apply Changes**

35.55.7 Routing

Specific networks can be routed across the VPN by adding a static route for the network(s) under **System > Routing** on the **Static Routes** tab.

These steps should be done on both sites.

- Navigate to **System > Routing > Static Routes**
- Click **Add**
- Fill in the options using the information determined earlier, with variations noted for each site:

HQ Settings

Destination Network
`10.23.0.0/24` (e.g. Satellite office LAN segment)

Gateway
`WG_VPN_SAT_V4`

Satellite Office Settings

Destination Network
`10.15.0.0/24` (e.g. HQ LAN segment)

Gateway
`WG_VPN_HQ_V4`

- Click **Save**
- Click **Apply Changes**

See also:

As an alternative to static routing in this way, dynamic routing protocols can also work with WireGuard. See *WireGuard Routing* for more information.

Tip: These gateways can also be used for policy routing if needed.

35.55.8 Finish Up

The configuration is now complete! The two sites should now have full LAN-to-LAN connectivity.

See also:

- *WireGuard*
- *Routing*
- *WireGuard Remote Access VPN Configuration Example*
- *WireGuard Site-to-Multisite VPN Configuration Example*
- *WireGuard VPN Client Configuration Example*

35.56 WireGuard Site-to-Multisite VPN Configuration Example

This recipe explains how to set up a VPN tunnel between three firewalls in a site-to-multisite configuration using *WireGuard*.

This example is a minimal configuration, more complicated scenarios are possible, see *WireGuard* for details.

35.56.1 Required Information

General Values

Item	Value
Design	Site-to-Multisite, one peer per tunnel (spokes), multiple peers per tunnel (hubs)
Tunnel Subnet	`172.16.0.0/29`

HQ:

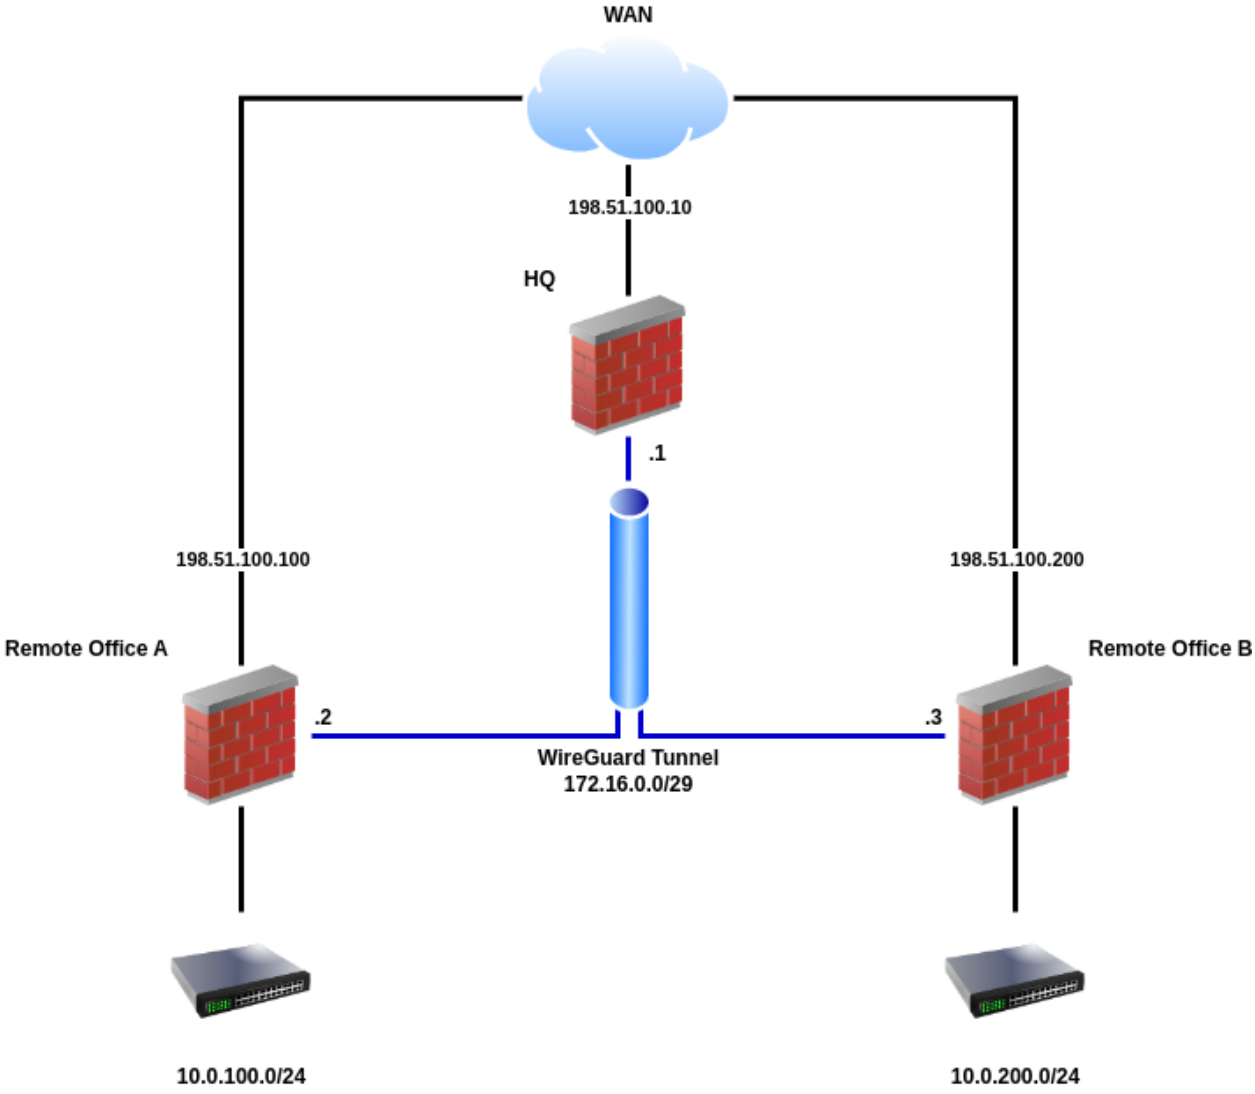


Fig. 86: WireGuard Example Site-to-Multisite Network

Item	Value
WAN IP Address	`198.51.100.10`
Tunnel Address	`172.16.0.1/29`
Listen Port	`51820`

Remote Office A:

Item	Value
WAN IP Address	`198.51.100.100`
Tunnel Address	`172.16.0.2/29`
Listen Port	`51820`
LAN Subnet	`10.0.100.0/24`

Remote Office B:

Item	Value
WAN IP Address	`198.51.100.200`
Tunnel Address	`172.16.0.3/29`
Listen Port	`51820`
LAN Subnet	`10.0.200.0/24`

35.56.2 WireGuard Configuration

- Navigate to **VPN > WireGuard > Settings**
- Fill in the following options:

  Enable
  Checked

  Interface Group Membership
  Only Unassigned Tunnels

- Click  **Save**

Tip: When allowing inbound connections from arbitrary remote networks, use rules only on assigned WireGuard interface tabs only to ensure proper return routing.

Note: Rules on assigned WireGuard interface tabs get `reply-to` which ensures that traffic entering a specific assigned WireGuard interface exits back out the same interface. Without that, return traffic will follow the default gateway.

35.56.3 Tunnel Configuration

First, create the WireGuard tunnel on all three sites:

- Navigate to **VPN > WireGuard > Tunnels**
- Click  **Add Tunnel**
- Fill in the options using the information determined earlier, with variations noted for each site:

  Enabled
  Checked

  HQ Settings

  Description
  `Remote Sites VPN`

  Remote Office A Settings

  Description
  `HQ VPN`

  Remote Office B Settings

  Description
  `HQ VPN`

  Listen Port
  `51820`

  Interface Keys
  Click **Generate** to create a new set of keys.

- Copy the public key from each firewall and note which is which
- Click  **Save Tunnel**

35.56.4 Peer Configuration

The peer entry for the server can be added when editing the tunnel. Follow these steps on both sites, with the differences in settings noted inline.

Edit the tunnel:

- Navigate to **VPN > WireGuard > Tunnels**
- Locate the WireGuard tunnel for this VPN
- Click  at the end of the row for the tunnel

From the tunnel editing page, add a peer:

- Click  **Add Peer**
- Fill in the options using the information determined earlier, with variations noted for each site:

  HQ Settings

Description
`Remote Office A Peer`

Dynamic Endpoint
Unchecked

Endpoint
`198.51.100.100` (the WAN IP address of Remote Site A)

Endpoint Port
`51820`

Public Key
The public key *from the Remote Office A firewall*

Allowed IPs
`172.16.0.2/32` and `10.0.100.0/24` (Remote Site A Tunnel Interface and LAN)

HQ Settings

Description
`Remote Office B Peer`

Dynamic Endpoint
Unchecked

Endpoint
`198.51.100.200` (the WAN IP address of Remote Site B)

Endpoint Port
`51820`

Public Key
The public key *from the Remote Office B firewall*

Allowed IPs
`172.16.0.3/32` and `10.0.200.0/24` (Remote Site B Tunnel Interface and LAN)

Remote Office A Settings

Description
`HQ Peer`

Dynamic Endpoint
Unchecked

Endpoint
`198.51.100.10` (the WAN IP address of HQ)

Endpoint Port
`51820`

Public Key
The public key *from the HQ firewall*

Allowed IPs
`172.16.0.0/29` and `10.0.200.0/24` (Tunnel Network and Remote Site B LAN)

Remote Office B Settings

Description
`HQ Peer`

Dynamic Endpoint
Unchecked

Endpoint
`198.51.100.10` (the WAN IP address of HQ)

Endpoint Port
`51820`

Public Key
The public key *from the HQ firewall*

Allowed IPs
`172.16.0.0/29` and `10.0.100.0/24` (Tunnel Network and Remote Site A LAN)

- Click  **Save Peer**

35.56.5 Assign Interface

These steps should be done on all sites.

First, fix the default gateway so WireGuard isn't automatically selected before it's ready:

- Navigate to **System > Routing**
- Set **Default Gateway IPv4** to a specific gateway (e.g. *WANGW*) or group
- Set **Default Gateway IPv6** similarly if this VPN will also carry IPv6 traffic
- Click  **Save**
- Click  **Apply Changes**

Next, assign the interface (*Assign a WireGuard Interface*):

- Navigate to **Interfaces > Assignments**
- Select the appropriate `tun_wg<number>` interface in the **Available network ports** list
- Click  **Add** to assign the interface as a new OPT interface (e.g. OPT1)
- Navigate to the Interface configuration page, **Interfaces > OPTx**
- Check **Enable**
- Enter an appropriate **Description** which will become the interface name (e.g. `VPN_HQ`, `VPN_SITEA`, or `VPN_SITEB`)
- Configure an appropriate **MTU** value for the WireGuard interface (e.g. `1420` for IPv4+IPv6 or `1440` for IPv4 only).

  For details on calculating the correct MTU, see in *Assign a WireGuard Interface*.
- Fill in the options using the information determined earlier, with variations noted for each site:

  HQ Settings (VPN_HQ)

  IPv4 Configuration Type
  `Static IPV4`

  IPv4 Address
  `172.16.0.1/29`

Remote Office A Settings (VPN_SITEA)

IPv4 Configuration Type
`Static IPV4`

IPv4 Address
`172.16.0.2/29`

Remote Office B Settings (VPN_SITEB)

IPv4 Configuration Type
`Static IPV4`

IPv4 Address
`172.16.0.3/29`

- Click  **Save**
- Click  **Apply Changes**

35.56.6 Firewall Rules

First, add a rule to the WAN on both firewalls to allow traffic to reach WireGuard:

Note: Traffic flowing between peers is subject to firewall rules on the associated WireGuard interface.

- Navigate to **Firewall > Rules**, **WAN** tab
- Click  **Add** to create a new firewall rule at the top of the list so that it matches before other rules
- Configure the firewall rule as follows:

Action
Pass

Protocol
UDP

Source
This can typically be left at *Any*, but it is more secure to fill in the IP address of the opposing firewall.

Destination
WAN Address

Destination Port Range
(other), `51820`

Description
`Pass traffic to WireGuard`

- Click  **Save**
- Click  **Apply Changes**

Next, add a rule to pass traffic inside the WireGuard tunnel on both firewalls:

- Navigate to **Firewall > Rules**
- Click the tab for the assigned WireGuard interface (e.g. **VPN_HQ**, **VPN_SITEA**, or **VPN_SITEB**)
- Click  **Add** to add a new rule to the top of the list
- Use the following settings:

  Action
  Pass

  Protocol
  Any

  Source
  any

  Destination
  any

  Description
  `Pass VPN traffic from WireGuard peers`

  Note: This rule allows all traffic between sites, which is easy but not a secure practice. Traffic between the sites can be restricted as needed with less permissive rules.

- Click  **Save**
- Click  **Apply Changes**

35.56.7 Routing

These steps should be performed on all sites:

- Navigate to **System > Routing > Gateways**
- Click  **Add**
- Fill in the options using the information determined earlier, with variations noted for each site:

  HQ Settings

  Interface
  `tun_wg<number>`

  Address Family
  `IPv4`

  Name
  `WG_SITEA_GW4`

  Gateway
  `172.16.0.2`

HQ Settings

Interface
`tun_wg<number>`

Address Family
`IPv4`

Name
`WG_SITEB_GW4`

Gateway
`172.16.0.3`

Remote Office A/B

Interface
`tun_wg<number>`

Address Family
`IPv4`

Name
`WG_HQ_GW4`

Gateway
`172.16.0.1`

- Click  **Save**
- Click  **Apply Changes**

Specific networks can be routed across the VPN by adding a static route for the network(s) under **System > Routing** on the **Static Routes** tab.

These steps should be done on all sites.

- Navigate to **System > Routing > Static Routes**
- Click  **Add**
- Fill in the options using the information determined earlier, with variations noted for each site:

HQ Settings

Destination Network
`10.0.100.0/24` (e.g. Remote Office A LAN segment)

Gateway
`WG_SITEA_GW4`

HQ Settings

Destination Network
`10.0.200.0/24` (e.g. Remote Office B LAN segment)

Gateway
`WG_SITEB_GW4`

Remote Office A Settings

Destination Network
`10.0.200.0/24` (e.g. Remote Office B LAN segment)

Gateway
`WG_HQ_GW4`

Remote Office B Settings

Destination Network
`10.0.100.0/24` (e.g. Remote Office A LAN segment)

Gateway
`WG_HQ_GW4`

- Click  **Save**
- Click  **Apply Changes**

See also:

As an alternative to static routing in this way, dynamic routing protocols can also work with WireGuard. See *WireGuard Routing* for more information.

Tip: These gateways can also be used for policy routing if needed.

35.56.8 Finish Up

The configuration is now complete! The two remote office sites should now have full LAN-to-LAN connectivity via HQ.

See also:

- *WireGuard*
- *Routing*
- *WireGuard Remote Access VPN Configuration Example*
- *WireGuard Site-to-Site VPN Configuration Example*
- *WireGuard VPN Client Configuration Example*

35.57 WireGuard VPN Client Configuration Example

This recipe explains how to set up *WireGuard* as a "client" to a remote VPN service through which Internet traffic will be routed.

Note: Though WireGuard does not have a concept of "Client" and "Server" per se, in this style of deployment the firewall initiates connections to a remote peer, but the peer never initiates back to the firewall. In this way, the firewall behaves like a "Client" and may be referred to as such in this document. The remote peer may also be referred to as "server".

35.57.1 Required Information

The following basic information must be determined before starting the VPN configuration. This example information was obtained from a popular WireGuard VPN Provider.

Item	Value
Tunnel Addresses	`10.68.140.33/32` and `fc00:bbbb:bbbb:bb01::5:8c20/128`
Tunnel Private Key	ADRM6pyoYpofcDd0TkX4sb7UkR+Zj4AYeZOE2WWg2tI=
Peer Public Key	EPLh6pVel06dND8cE4Prix9GP4hGLYNhQhn5mSN2yzM=
Peer Endpoint	`86.106.143.236`
Peer Port	`51820`
Peer WG Address	Same as tunnel addresses for `/32` and `/128` routes
Peer DNS Server	`193.138.218.74`
Allowed IPs	`0.0.0.0/0` and `::0/0`

Hint: Start with configuring IPv4 connectivity first. Once IPv4 connectivity is established and working, then circle back and configure IPv6 connectivity if desired.

Some or all of these values must be obtained from the VPN provider or server administrator. Methods vary, but some may have a web-based portal which shows settings or generates a configuration file. Others may opt to send settings in a more secure manner.

Keys

In this role, the source of the keys can vary. Ideally, a private and public key set for this firewall should be generated by this firewall and the private key should never leave. The public key should be copied and submitted to the administrator of the server side so it can be used for this client.

Some providers insist on generating the keys themselves so they can pre-allocate addresses and other settings based on keys they already know. This also allows them to easily generate configurations for clients. It's less secure this way, but more convenient.

35.57.2 Tunnel Configuration

First create the WireGuard tunnel.

- Navigate to **VPN > WireGuard > Tunnels**
- Click  **Add Tunnel**
- Fill in the options using the information determined earlier:

  Enabled
  Checked

  Description
  `VPN Provider`

  Listen Port
  This does not likely matter unless the server requires a specific source port. In most cases it can be left blank or at the default `51820`.

Interface Keys
Enter the private key supplied by the provider `ADRM6pyoYpofcDd0TkX4sb7UkR+Zj4AYeZOE2WWg2tI=`.

Note: Click **Generate** to generate a new key pair if the provider accepts user-generated keys.

- Click **Save**

35.57.3 Peer Configuration

The peer entry for the server can be added when editing the tunnel. To edit the tunnel:

- Navigate to **VPN > WireGuard > Tunnels**
- Locate the WireGuard tunnel for this VPN provider
- Click  at the end of the row for the tunnel

From the tunnel editing page, add a peer as follows:

- Click  **Add Peer**
- Fill in the options using the information determined earlier:

Enable
Checked

Tunnel
The WireGuard tunnel for this VPN provider.

Description
The name of this server or VPN provider.

Dynamic Endpoint
Unchecked

Endpoint
The server hostname or IP address, `86.106.143.236` in this example.

Endpoint Port
The server WireGuard port, `51820` in this example.

Public Key
The public key for the **VPN provider endpoint**, given by the VPN provider `EPLh6pVel06dND8cE4Prix9GP4hGLYNhQhn5mSN2yzM=`.

Pre-Shared Key
Not used in this example, but for additional security this pre-shared key can be generated and copied to the peer. Must match on the client and server.

Most VPN providers are not utilizing pre-shared keys at this time.

Allowed IPs
List of networks to route to the remote side. Since this example will be sending all traffic through the VPN provider, enter `0.0.0.0/0` and `::0/0`.

- Click **Save Peer**

35.57.4 Confirm Handshakes

At this point it is possible to confirm basic connectivity with the VPN provider.

- Navigate to **VPN > WireGuard > Status**
- Click **Show Peers**
- Confirm peer connectivity and recent handshaking with the peer

35.57.5 Assign Interface

First, fix the default gateway so WireGuard isn't automatically selected before it's ready:

- Navigate to **System > Routing**
- Set **Default Gateway IPv4** to a specific gateway (e.g. *WANGW*) or group
- Set **Default Gateway IPv6** similarly if this VPN will also carry IPv6 traffic
- Click **Save**
- Click **Apply Changes**

Next, assign the interface (*Assign a WireGuard Interface*):

- Navigate to **Interfaces > Assignments**
- Select the appropriate `tun_wg<number>` interface in the **Available network ports** list
- Click  **Add** to assign the interface as a new OPT interface (e.g. OPT1)
- Navigate to the Interface configuration page, **Interfaces > OPTx**
- Check **Enable**
- Enter an appropriate **Description** which will become the interface name (e.g. `WG_VPN`)
- Configure an appropriate **MTU** value for the WireGuard interface (e.g. `1420` for IPv4+IPv6 or `1440` for IPv4 only).

  For details on calculating the correct MTU, see in *Assign a WireGuard Interface*.
- Fill in the options using the information determined earlier:

  IPv4 Configuration Type
  `Static IPV4`

  IPv6 Configuration Type
  `Static IPv6`

  IPv4 Address
  `10.68.140.33/32`

  IPv4 Upstream Gateway

  - Click **Add a new gateway**
  - Fill in the options:

    Gateway Name
    `WG_VPN_v4`

    Gateway IPv4
    `10.68.140.33`

- Click  **Add**

IPv6 Address
`fc00:bbbb:bbbb:bb01::5:8c20/128`

IPv6 Upstream Gateway

- Click **Add a new gateway**
- Fill in the options:

Gateway Name
`WG_VPN_v6`

Gateway IPv6
`fc00:bbbb:bbbb:bb01::5:8c20`

- Click  **Add**

- Click **Save**
- Click **Apply Changes**

35.57.6 Gateways and Groups

These gateways can be added to a gateway group for failover or load balancing of outbound traffic. This example assumes there are no existing groups. If there are groups already, the new gateway can be added to them like any other.

This example sets up a Gateway Group which prefers WireGuard and fails over to WAN. Traffic directed to this group will use WireGuard when it is up, and WAN when it is down. In practice this specific behavior may or may not be desirable, but can be used as a template for other scenarios.

Note: This will only function properly if gateway monitoring is possible. For that to work, edit the WireGuard interface gateways and fill in a different **Monitor IP** address which responds to ICMP echo (ping) requests over the WireGuard tunnel.

To create a new group:

- Navigate to **System > Routing**, **Gateway Groups** tab
- Click  **Add**
- Configure the group as follows:

Group Name
`Prefer_WireGuard_V4`

Gateway Priority

WG_VPN_v4
Tier 1

WANGW
Tier 2

Description
`Prefer VPN, fail to WAN`

- Repeat for IPv6 if required
- Click **Save**
- Click **Apply Changes**

35.57.7 Outbound NAT

By default, the VPN will not have outbound NAT applied to its traffic. Most VPN providers will require this, so that all traffic appears to originate from the address of the VPN interface, and not LAN.

To set up outbound NAT for the VPN:

- Navigate to **Firewall > NAT**, **Outbound** tab
- Set **Mode** to **Hybrid Outbound NAT**

  This example assumes the firewall starts out on **Automatic Outbound NAT**. If the firewall is using **Manual Outbound NAT**, there is no need to change the mode.
- Click **Save**
- Click | fa-turn-up| **Add** to create a new outbound NAT rule at the top of the list
- Configure the NAT rule as follows:

  Interface
  The assigned WireGuard interface (e.g. *WG_VPN*)

  Source
  The LAN subnet of this firewall (e.g. `192.168.1.0/24`)

  Description
  A description of the rule, if desired: `Outbound NAT for LAN to WireGuard VPN Provider`

  Leave all remaining options at their default values
- Repeat for IPv6 if required
- Click **Save**
- Click **Apply Changes**

35.57.8 Firewall Rules

This scenario should not require any firewall rules on the WAN or VPN interface. No connections will be made inbound on the WAN, only outbound. Traffic from the Internet will not be allowed back into the VPN interface.

35.57.9 Routing Traffic

Policy Routing

Rules can be added to local interfaces, such as LAN, for policy routing which utilize the gateway for the WireGuard interface.

For example, to policy route all traffic from a host on the LAN out through WireGuard:

- Navigate to **Firewall > Rules**, **LAN** tab

- Click  **Add** to create a new firewall rule at the top of the list so that it matches before other rules
- Configure the firewall rule as follows:

  Action
  Pass

  Interface
  LAN

  Protocol
  Any

  Source
  Set this to match the client whose outbound traffic will be routed across the VPN. For example:

  Address or Alias, `192.168.1.23`

  Destination
  Any

  Gateway
  WG_VPN_WGV4

  Note: Click **Display Advanced** to show this option.

  Leave all remaining options at their default values
- Repeat for IPv6 if required
- Click **Save**
- Click **Apply Changes**

This concept can be adapted for a number of different scenarios. For example, match all LAN traffic and send it across the VPN, or match traffic and use a gateway group to prefer the VPN, etc.

Static Routing

Specific networks can be routed across the VPN by adding a static route for the network(s) under **System > Routing** on the **Static Routes** tab.

Default Gateway

This is an optional step that some users may want to perform if they want all traffic from the firewall to cross the VPN, not only LAN client traffic.

Policy routing is the most flexible way to direct traffic over this type of connection, but it does not influence traffic from the firewall itself. To send traffic from the firewall across the VPN to Internet destinations, the VPN must be set as the default gateway.

To avoid a chicken-and-egg problem, a manual static route is required for the VPN provider peer endpoint address:

- Navigate to **System > Routing**, **Static Routes** tab
- Click  **Add**

- Configure the routes as follows:

  Destination network
  The VPN provider peer endpoint IP address. For this example, `86.106.143.236`. Use a CIDR mask of `32` (or `128` if the peer endpoint is an IPv6 address.)

  Gateway
  WANGW so that traffic for this endpoint is routed over WAN

  Description
  `Route to VPN provider endpoint`

- Click **Save**

With the peer route in place, now set the default gateway:

- Navigate to **System > Routing**, **Gateways** tab
- Set **Default Gateway IPv4** to *WG_VPN_V4*, or a gateway group which includes that gateway, such as the previously created *Prefer_WireGuard*.
- Set **Default Gateway IPv6** similarly if the VPN also carries IPv6 traffic.
- Repeat for IPv6 if required
- Click **Save**
- Click **Apply Changes**

At this point, all traffic that doesn't match entries in the routing table will be sent across the VPN.

35.57.10 DNS Configuration

DNS privacy is also important, and there are a few factors to consider. For this example, DNS requests will be sent to a DNS server at the VPN peer, but without TLS.

See also:

Options such as *DNS over TLS* are covered elsewhere, but can help as well.

First, set the VPN provider DNS server:

- Navigate to **System > General**
- Remove any DNS servers present in the list under **DNS Server Settings**
- Set a DNS Server entry as follows:

  Address
  The address of the DNS server at the peer, in this example, `193.138.218.74`.

  DNS Hostname
  If this server offers DNS over TLS, enter its hostname here. Otherwise, leave it blank.

  Gateway
  Select the VPN gateway, *WG_VPN_V4*.

- *Uncheck* **DNS Server Override** to prevent this firewall from using DNS servers from dynamic WANs.
- Set **DNS Resolution Behavior** based on the requirements of this environment:

> **Warning:** This can help prevent DNS requests from leaking to other servers not using the VPN, but it can cause a chicken-end-egg scenario where DNS requests will fail unless the VPN is working. Ensure that DNS is not required to establish the VPN.

Use local DNS, fall back to remote DNS Servers
Use this option when using the DNS Resolver in forwarding mode and when the DNS server does not need DNS over TLS. This is the best fit for this example.

Use local DNS, ignore remote DNS Servers
Use this option when using DNS over TLS with the DNS Resolver in forwarding mode. This ensures that no DNS query will be sent without TLS.

Use remote DNS Servers, ignore local DNS
Use this option if the firewall itself shouldn't use the DNS Resolver, but communicate directly with the DNS server without TLS.

Next, configure the DNS Resolver for Forwarding mode:

- Navigate to **Services > DNS Resolver**
- *Uncheck* **Enable DNSSEC Support**
- *Check* **Enable Forwarding Mode**
- Repeat for IPv6 if required
- Click **Save**
- Click **Apply Changes**

> **Warning:** If there are any **Custom Options** in the DNS Resolver, it is possible that switching to forwarding mode will change the context of the options. Add `server:` to the beginning of the **Custom Options** box content, above any existing options. If the **Custom Options** box is empty, it can remain empty.

35.57.11 Finish Up

The configuration is now complete! Depending on which sections were followed, the firewall should be able to at least communicate with the remote peer, networks, and clients should be able to pass traffic through the VPN provider out to the Internet.

Make any final adjustments or additional configurations as needed.

See also:

- *WireGuard*
- *Routing*
- *Policy Routing Configuration*
- *Configuring DNS over TLS*
- *WireGuard Remote Access VPN Configuration Example*
- *WireGuard Site-to-Site VPN Configuration Example*
- *WireGuard Site-to-Multisite VPN Configuration Example*

35.58 Accessing Port Forwards from Local Networks

By default, pfSense® software does not redirect internally connected devices to forwarded ports and 1:1 NAT on WAN interfaces. For example, if a client on LAN attempts to reach a service forwarded from WAN port `80` or `443`, the connection will hit the firewall web interface and not the service they intended to access. The client will be presented with a certificate error if the GUI is running HTTPS, and a DNS rebinding error since the GUI rejects access for unrecognized hostnames.

Split DNS and NAT reflection are two techniques which allow clients to reach these types of resources from local networks. Split DNS allows clients to avoid passing this traffic through the firewall, while NAT Reflection employs techniques to redirect these connections through the firewall. Split DNS is the best practice because it allows for retaining of the original source IP address and avoids unnecessarily looping internal traffic through the firewall. Both techniques are explained in this document.

35.58.1 Method 1: Split DNS

Split DNS is a configuration where internal and external clients resolve hostnames differently. Split DNS is the best practice to solve this problem, and it is a much more elegant solution than NAT reflection.

In this scenario, internal clients access resources by hostname, not IP address. Clients on the local network resolve that hostname to the actual local IP address of the server, and not the external (e.g. WAN) IP address as others outside the network would see.

For this to work using the DNS Resolver or Forwarder in pfSense software, clients must use the IP address of the firewall as their primary DNS server.

Note: If the clients all use some other internal DNS server not on the firewall, such as Active Directory, split DNS can still work. Configure the internal DNS server similarly to what is described in this section.

Example:

- `www.example.com` resolves to public IP address `1.2.3.4`, which is the WAN IP address of the firewall.
- The firewall is configured to forward port `80` on `1.2.3.4` to port `80` on `192.168.1.5`, the internal web server.
- Override `www.example.com` using **Services > DNS Resolver** (or **DNS Forwarder**, if that is active instead) and point `www.example.com` to `192.168.1.5`

Screenshots that show the above in practice:

35.58.2 Method 2: NAT Reflection

To access ports forwarded on the WAN interface from internal networks using NAT reflection, the feature must be explicitly enabled:

- Navigate to **System > Advanced**, **Firewall & NAT** tab
- Configure the following options in the **Network Address Translation** section of the page:

  NAT Reflection mode for port forwards
  Pure NAT

  Pure NAT mode is the best type of NAT reflection, but it may not work for all scenarios. See *NAT Reflection mode for Port Forwards* for details on each of the NAT reflection modes.

  Enable NAT Reflection for 1:1 NAT
  Checked

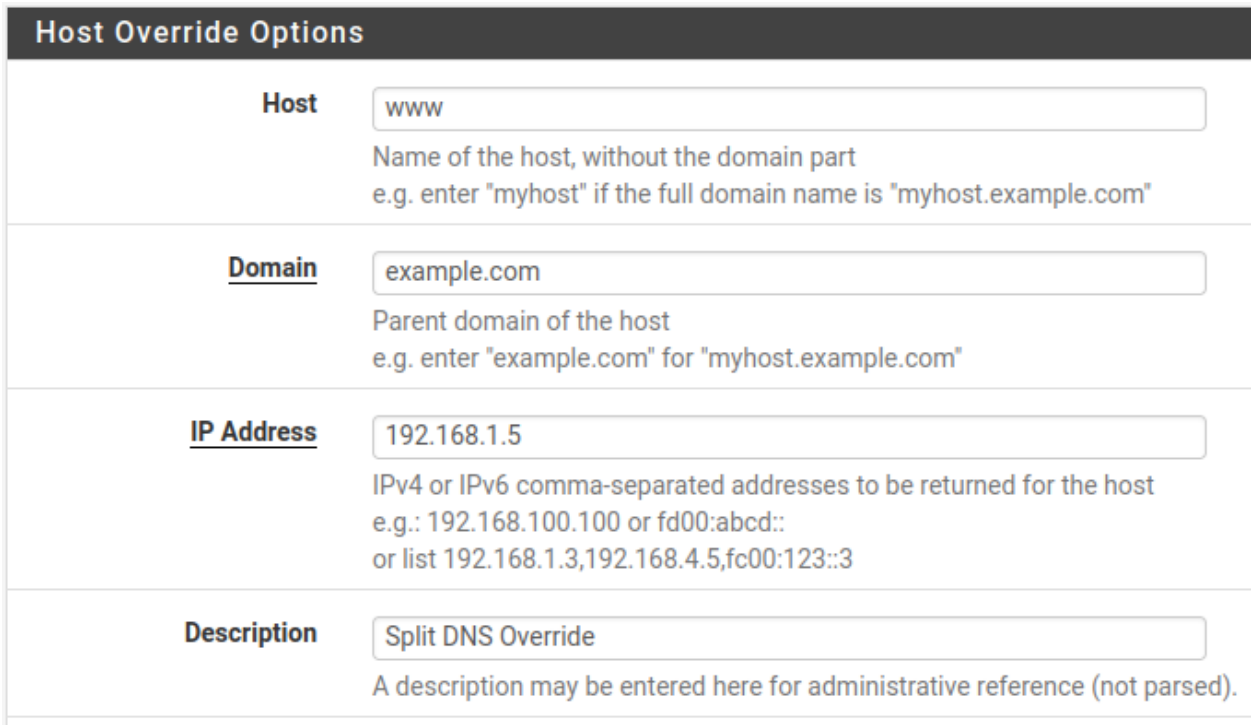


Fig. 87: Adding a DNS Resolver host override for split DNS

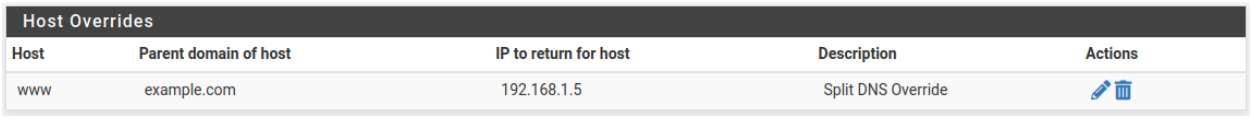


Fig. 88: Split DNS entry in the list of host overrides

Enable automatic outbound NAT for Reflection
Checked

- Click *Save*

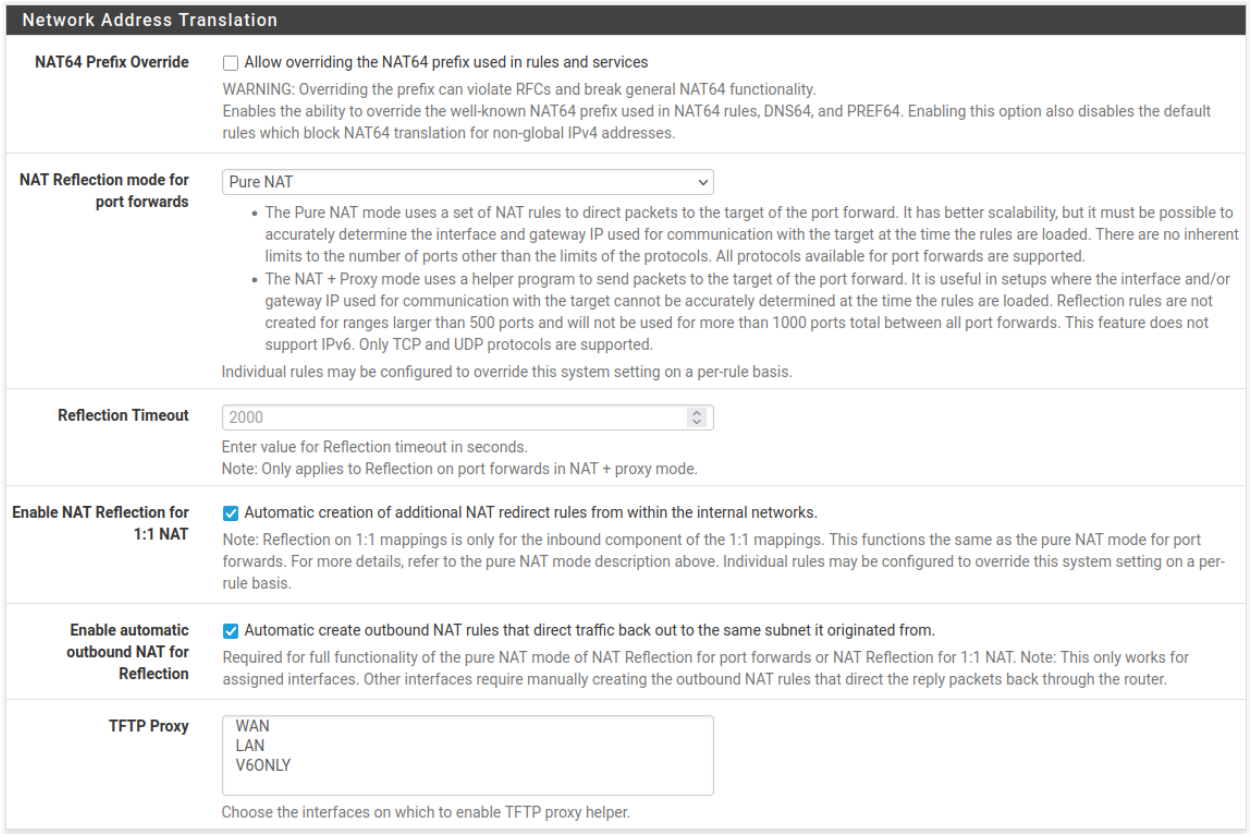


Fig. 89: NAT Reflection Settings

35.59 Authenticating from Active Directory using RADIUS/NPS

Windows Servers can be configured as a RADIUS server using the Microsoft Network Policy Server (NPS). This allows a Windows Server to handle authentication for OpenVPN, Captive Portal, the PPPoE server, or even the firewall GUI itself. NPS can authenticate based on Windows Server local user accounts or Active Directory.

Note: While NPS capability has been present since Windows Server 2008, this document focuses on current versions of Windows Server software.

The options may vary slightly depending on the version of Windows Server software.

35.59.1 Choosing a server for NPS

NPS requires a minimal amount of resources and is suitable for addition to an existing Windows Server in most environments. Microsoft recommends installing it on an Active Directory domain controller to improve performance in environments where NPS is authenticating against Active Directory.

Tip: NPS can also be installed on a member server, which may be desirable in some environments to reduce the attack footprint of domain controllers. Each network-accessible service provides another potential avenue for compromising a server. NPS has a solid security record, especially compared to other services that must be running on domain controllers for Active Directory to function, so this isn't much of a concern in most network environments.

Most environments install NPS on one of their domain controllers. Microsoft recommends running it on each domain controller in the forest and using NPS proxies to share the load for a busy environment.

35.59.2 Installing NPS

- Open the Server Manager Dashboard
- Click **Add Roles and Features**

  This may be on the main screen or under the **Manage** menu.
- Click **Next** until the wizard displays the server selection screen
- Select this server from the list
- Click **Next** again
- Check **Network Policy and Access Services** on the list of roles
- Click **Add Features** if it appears
- Click **Next** on each screen until the end of the wizard
- Click **Finish** or **Install**, depending on the Windows Server version
- Click **Close** once the installation completes

35.59.3 Configuring NPS

To configure NPS, bring up the Server Manager and select the new role. The name varies on different versions of Windows Server but may be NPAS (2022), NAP (2012), Network Policy and Access Services, or a similar name.

First configure a RADIUS client for the firewall, then setup remote access policies.

Adding a RADIUS Client

- Open the **Server Manager** dashboard
- Click **NPAS** or its equivalent name (**NAP**, etc.)
- Right-click on this server in the server list
- Click **Network Policy Server**
- Expand **RADIUS Clients and Server**
- Click **RADIUS Clients**

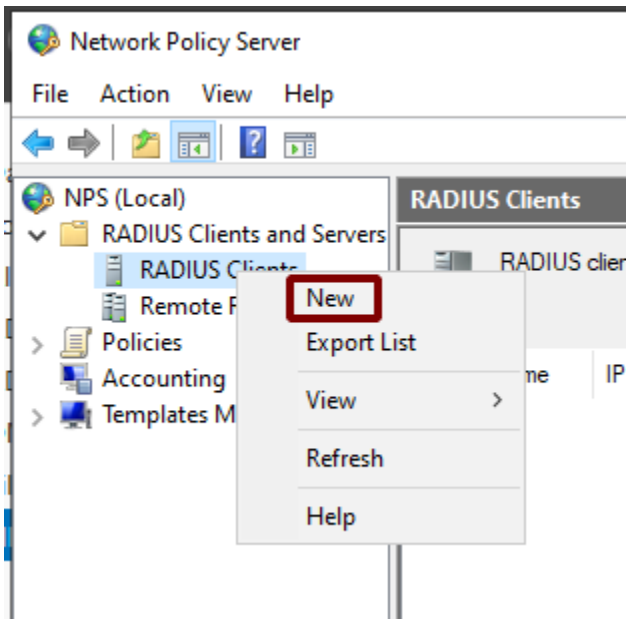


Fig. 90: Add New RADIUS Client

Add the new RADIUS client:

- Right-click on **RADIUS Clients**
- Click **New**, as shown in Figure *Add New RADIUS Client*
- Enter a **Friendly name** for the firewall, as shown in Figure *Add New RADIUS Client Address*.

  This can be the hostname or an FQDN.
- Enter the **Address (IP or DNS)** for the firewall.

  This must be the IP address from which the firewall will initiate RADIUS requests or an FQDN which resolves to that IP address.

  Note: This is the IP address of the firewall interface closest to the RADIUS server. If the RADIUS server is reachable via the firewall LAN interface, this will be the LAN IP address of the firewall. In deployments where the firewall is not the perimeter firewall, and the WAN interface resides on the internal network where the RADIUS server resides, the WAN IP address would be the correct address.

- Enter a **Shared secret**, as shown in Figure *Add New RADIUS Client Shared Secret*.

  This shared secret is used by the firewall to authenticate itself when making RADIUS access requests.

  Windows can automatically create a shared secret using the **Generate** option.
- Click OK.

The NPS configuration for the RADIUS client is now complete. The RADIUS Client is visible as in Figure *Listing of the RADIUS Client*.

Refer to other sections in this documentation describing the service to be used with RADIUS for more guidance on how to utilize the service. The *User Manager* can use NPS as an authentication server which also enables RADIUS for IPsec, OpenVPN, and *Captive Portal*. Other services such as the *PPPoE server* can use it directly as well.

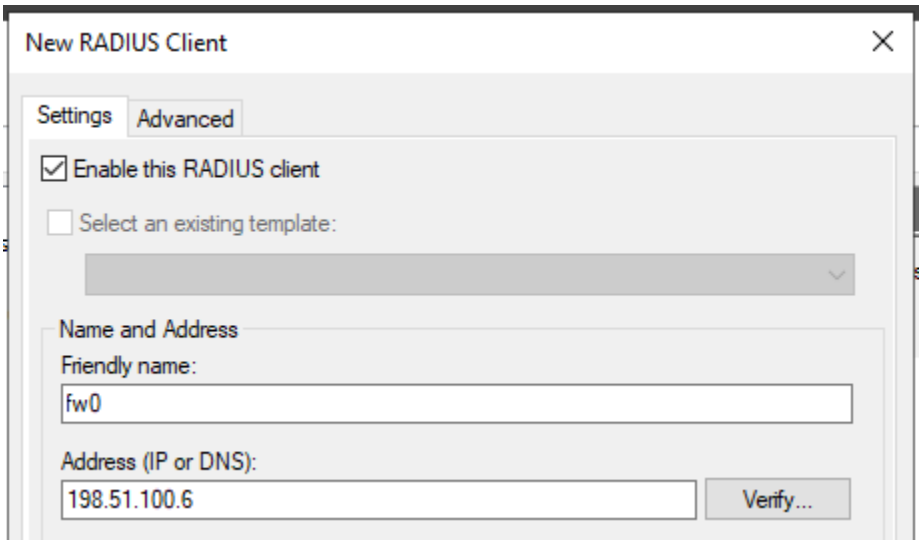


Fig. 91: Add New RADIUS Client Address

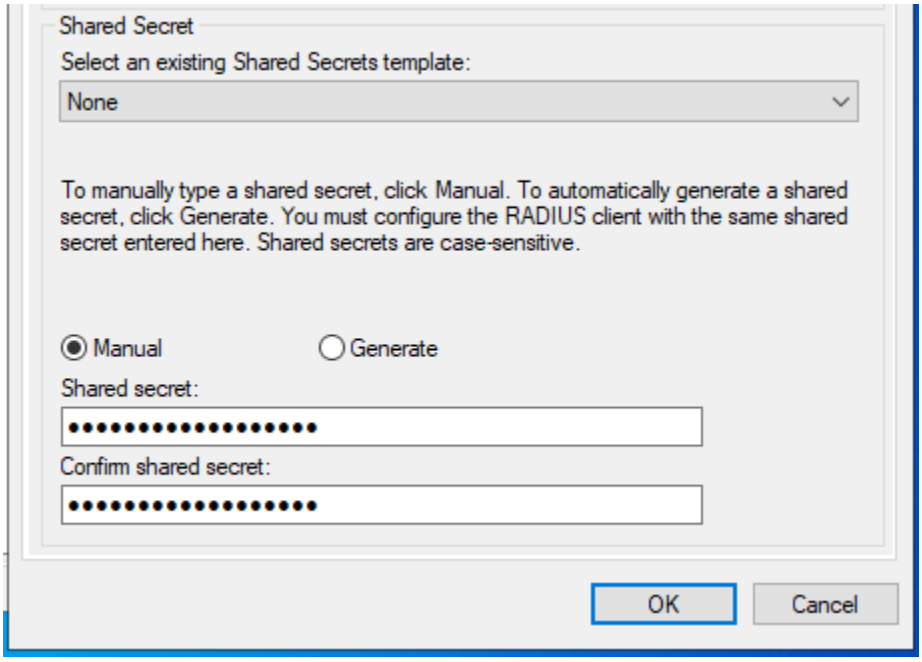


Fig. 92: Add New RADIUS Client Shared Secret

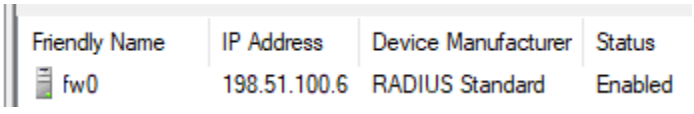


Fig. 93: Listing of the RADIUS Client

Configuring Users and Network Policies

Network Policies control whether a user can authenticate via RADIUS. Using Network Policies, an administrator can place a user in a specific Active Directory group to allow VPN access and also offer more advanced capabilities such as time of day restrictions.

More information on remote access policies can be found in Microsoft's documentation at http://technet.microsoft.com/en-us/library/cc785236%28WS.10%29.aspx.

Adding a Network Policy

- Open the **Server Manager** dashboard
- Click **NPAS** or its equivalent name (**NAP**, etc.)
- Right-click on this server in the server list
- Click **Network Policy Server**
- Expand **NPS (Local)**, **Policies**, then **Network Policies**
- Right-click on **Network Policies**
- Click **New**
- Enter `Allow from Firewall` in the **Policy name**
- Leave the **Type of network access server** set to *Unspecified*
- Click **Next**
- Click **Add** in the Specify Conditions window
- Select **Windows Groups**
- Click **Add**
- Enter or select the name of the user group which contains VPN users, e.g. `VPNUsers`
- Click **OK**
- Click **Next**
- Choose **Access granted**
- Click **Next**
- Add **EAP Types / Authentication Methods** as needed:
  - Leave existing authentication methods selected
  - Add or Select **Microsoft: Secured Password (EAP-MSCHAP v2)** if the firewall will use this policy for IPsec IKEv2 EAP-RADIUS authentication
  - Select **Encrypted Authentication (CHAP)**
  - Select **Unencrypted Authentication (PAP, SPAP)**
- Click **Next**
- Click **No** or **Decline** if the wizard prompts to view a help topic about security
- Configure any additional access constraints, if necessary
- Click **Next** on the remaining screens until the final screen is reached
- Click **Finish**

Editing an Existing Network Policy

Existing policies can be altered to change their constraints or other properties. For example, to edit an older policy to enable it for use by IPsec for IKEv2 EAP-RADIUS:

- Open the **Server Manager** dashboard
- Click **NPAS** or its equivalent name (**NAP**, etc.)
- Right-click on this server in the server list
- Click **Network Policy Server**
- Expand **NPS (Local)**, **Policies**, then **Network Policies**
- Edit the policy currently in use (e.g. right-click, click **Properties**)
- Click the **Constraints** tab
- Click **Authentication Methods**
- Click **Add**
- Select **Microsoft: Secured Password (EAP-MSCHAP v2)**
- Click **OK**
- Click **Apply** to restart NPS
- Click **OK**

35.59.4 Check Users and Groups

These steps are only necessary if the use case for this setup requires group authentication on the firewall.

Before proceeding, ensure any users who must authenticate using NPS are members of the correct groups (e.g. `VPNUsers`).

Create a matching group with a remote scope on the firewall (*Manage Local Groups*).

Edit the NPS policy on the Windows server so it returns the group name:

- Open the **Server Manager** dashboard
- Click **NPAS** or its equivalent name (**NAP**, etc.)
- Right-click on this server in the server list
- Click **Network Policy Server**
- Expand **NPS (Local)**, **Policies**, then **Network Policies**
- Edit the policy currently in use (e.g. right-click, click **Properties**)
- Click the **Settings** tab
- Click **Standard** under **RADIUS Attributes**
- Select **Class** from the list
- Click **Add**
- Select **String** for the attribute value type
- Enter a group name which matches a group on the firewall (e.g. `VPNUsers`)
- Click **OK**

- Click **Close**
- Click **Apply** to restart NPS
- Click **OK**

35.59.5 Add Authentication Server

Now that NPS is ready to accept authentication requests, the next step is to add an authentication server entry on the firewall.

See also:

RADIUS Authentication Servers

- Open the firewall GUI
- Navigate to **System > User Manager**, **Authentication Servers** tab
- Click  **Add** to create a new entry
- Enter the following settings:

  Descriptive name
  `Active Directory NPS`

  Type
  RADIUS

  Hostname or IP address
  `198.51.100.30` – Replace this with the IP address of the Windows server

  Shared Secret
  The password added to the NAS entry in NPS

  Services offered
  Authentication

  Authentication port
  `1812`

- Click **Save**

35.59.6 Test Authentication

On the firewall GUI, test the authentication:

- Navigate to **Diagnostics > Authentication**
- Set **Authentication Server** to the entry for NPS
- Enter a username and password for a user which should have access
- Click **Test**

If that test succeeded, then configure other services such as IPsec or OpenVPN to use the new RADIUS server and attempt authentication there.

35.59.7 Troubleshooting NPS

This section describes the most common problems users encounter with NPS.

Verify port

First ensure NPS is using the default port `1812`. If the NPS server was already installed, it may have be using a non-standard port.

- Open the **Server Manager** dashboard
- Click **NPAS** or its equivalent name (**NAP**, etc.)
- Right-click on this server in the server list
- Click **Network Policy Server**
- Right-click on **NPS (Local)** at the top left of the console
- Click **Properties**
- Click the **Ports** tab
- Verify that the **Authentication** port set includes port `1812`

  NPS can use multiple ports separated with commas, as shown in figure *NPS Ports*.
- Verify the **Accounting** port set includes port `1813` (optional)

  This is only necessary if the use case requires RADIUS accounting.

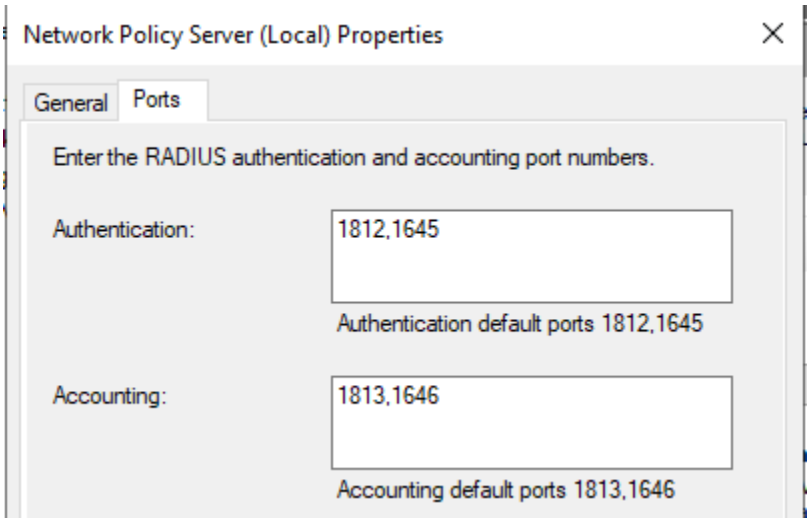


Fig. 94: NPS Ports

Check Event Viewer

When NPS handles a RADIUS authentication request it creates a log entry in the Security log in Event Viewer with the result of the authentication request. If it denies access, it logs the reason in the event log.

These log entries can be viewed in one of two ways:

View the **Security** log. This method is easier to identify success vs failure but on a busy server it may be difficult to isolate entries specific to NPS.

- Open Event Viewer on the Windows Server
- Expand **Windows Logs**
- Click **Security**
- Look for entries in the log which reference NPS

Use the custom view which only displays NPS log entries:

- Open Event Viewer on the Windows Server
- Expand **Custom Views**
- Expand **Server Roles**
- Click **Network Policy and Access Services**

Similar messages are available in both locations thought their format may vary slightly.

The contents of the log message contain a **Reason:** line which explains why authentication failed. The common two failures are:

- "Authentication failed due to a user credentials mismatch"

  This indicates that the user supplied an invalid username or password.

- "The Network Access Permission setting in the dial-in properties of the user account in Active Directory is set to Deny access to the user."

  Indicates that the user account is set to deny access or the network policies in NPS do not allow access for that user. For example, they may not be a member of the correct group.

If NPS is logging that authentication was successful, but the client is receiving a bad username or password message, ensure that the RADIUS secret configured in NPS and on the firewall match.

35.60 Allowing Remote Access to the GUI

Several ways exist to remotely administer a firewall running pfSense® software that come with varying levels of recommendation. They all work, but their use may vary for any number of reasons (Client restrictions, corporate policies, etc.)

35.60.1 Use a VPN

The safest way to accomplish the task is to set up a VPN that will allow access to the firewall and the network it protects. There are several VPN options available in pfSense software, such as

- *IPsec*
- *OpenVPN*
- *SSH tunneling*

Once a VPN is in place, reach the GUI safely using a local address on the firewall, such as the LAN IP address. The exact details vary depending on the VPN configuration.

35.60.2 Restricted Firewall Access

If the GUI port must be accessible to the Internet, restrict it by IP address or range as much as possible. Ideally, if there is a static IP address at the location to manage from, allow traffic from that IP address or subnet and nowhere else. Aliases also help, and they can include fully qualified domain names as well. If the remote management clients have a dynamic DNS address, add it to a management alias.

35.60.3 Use HTTPS

The best practice is to always use HTTPS to encrypt access to the GUI port. Modern browsers may complain about the certificate, but an exception can usually be stored so it will only complain the first time.

To disable (or re-enable) HTTPS for the GUI, navigate to **System > Advanced**, under the **Admin Access** tab, using the **Protocol** option in the webConfigurator section. See *Admin Access* for details.

35.60.4 Move the GUI to an Alternate Port

Moving the GUI to a non-standard, random port is also beneficial. This does not improve the actual security of the GUI itself, but can potentially reduce the number of brute force attempts. The GUI can still be found by scanners unless the port is properly filtered.

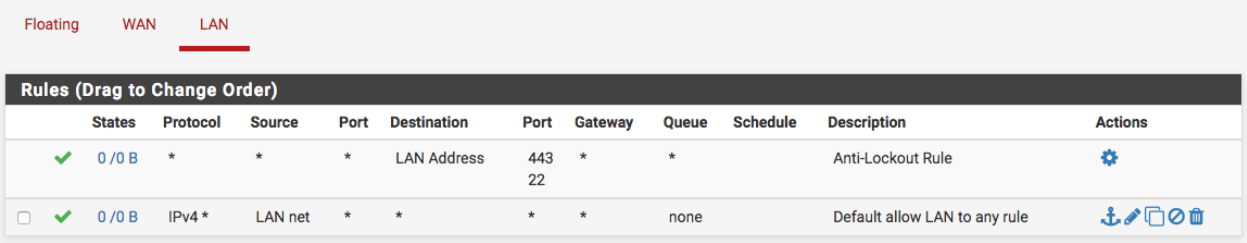
The port for the GUI can be changed under **System > Advanced**, **Admin Access** tab, using the *TCP Port* option in the **webConfigurator** section. Avoid common ports like `443`, `31337`, `8080`, `8888`, etc.

35.60.5 Strict Management

To enhance the security of a network, in many environments access to the firewall GUI is limited by firewall rules. Restricting access to the management interface is the best practice, for reasons as to why, see the blog post Securely Managing Web-administered Devices.

The default configuration of pfSense software allows management access from any machine on the LAN and denies it to anything outside the local network. There is also an anti-lockout rule enabled by default that prevents firewall rules from being configured in a way that will lock the user out of the web interface.

To restrict management access first ensure the LAN rules allow access to the port used for the GUI. This depicts the default LAN rule, which allows access to the web interface.

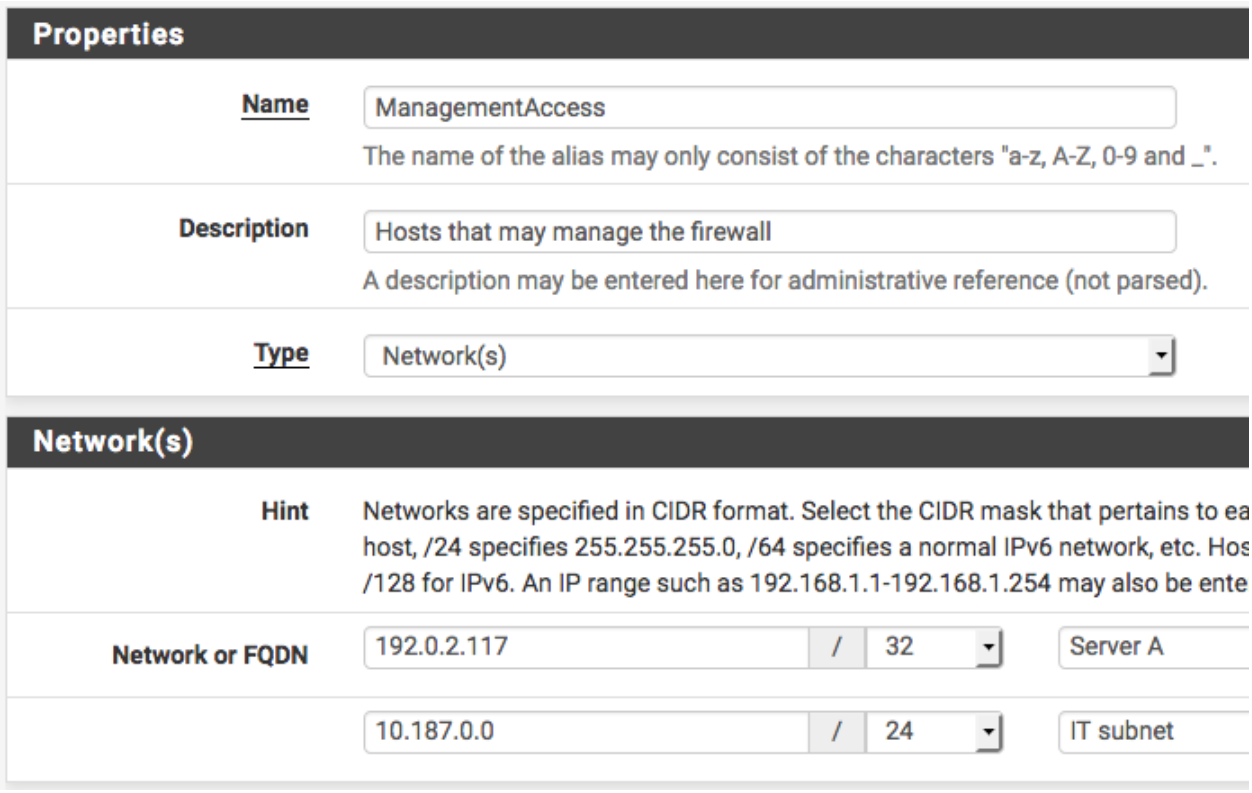


If a restrictive ruleset is in place on the LAN, make sure it permits access to the web interface before continuing.

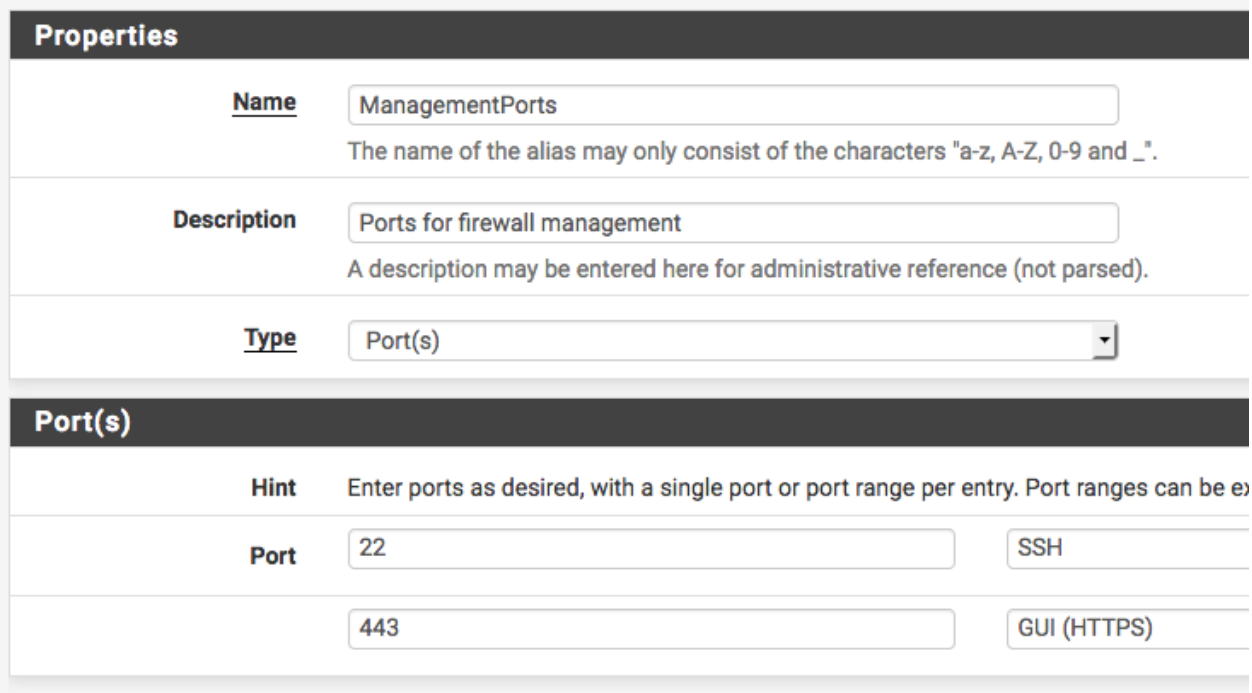
Now disable the anti-lockout rule. Navigate to **System > Advanced**, **Admin Access** tab and check **Disable webConfigurator anti-lockout rule**. Click **Save** and the rule will be removed.

Using a network alias for management access is another useful best practice. If both web and SSH administration are used, add an alias for those ports. The following are examples:

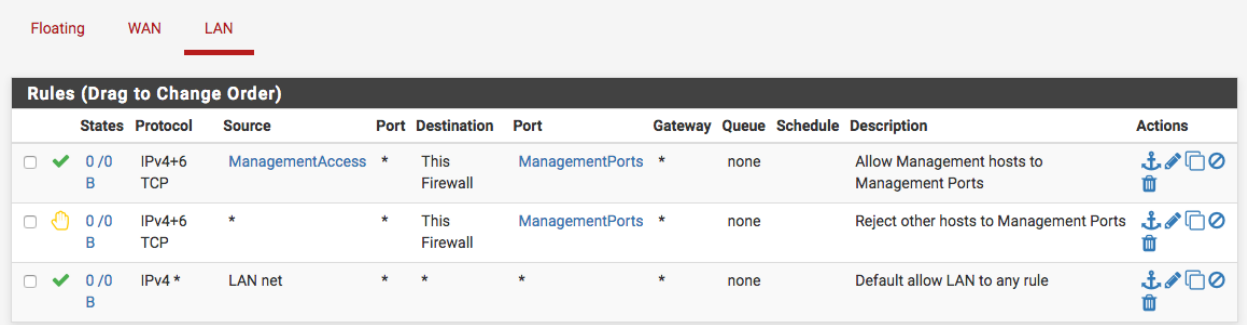
1. Example alias for networks allowed to access management interface



2. Example alias for ports allowed to access management interface



Now add a firewall rule allowing the sources defined in the management alias to the destination of the firewall, with the port used or alias created for those using multiple ports. **Make sure this rule is first in the list**. Then add a rule based on that rule (click + next to the rule), changing action to *block* or *reject* (reject is preferred on internal networks), source to *any*, and destination the same. When finished the ruleset should look like the following.



Click **Apply Changes** and the management interface is now restricted to only the defined hosts.

35.60.6 I Don't Care About Security, How Do I Open Access To The GUI?

To open the firewall GUI, create a firewall rule to allow remote firewall administration.

Note: Do not create a port forward or other NAT configuration.

Firewall > Rules, WAN Tab

Action
pass

Interface
WAN

Protocol
TCP

Source
The IP address or subnet of the client, an alias containing management hosts/networks, or (as a last resort only) *Any*

Destination
WAN Address

Destination port range
HTTPS (Or the custom port)

Description
Allow remote management from anywhere (Dangerous!)

35.61 Preventing RFC 1918 Traffic from Exiting a WAN Interface

RFC 1918 addresses are blocks of network IP addresses reserved for private use. These addresses are commonly used behind firewalls to allow a single public IP address to be shared with multiple devices using NAT. The default pfSense® software installation assigns the `192.168.1.0/24` address space to the LAN interface, but RFC 1918 also defines other CIDR ranges for private use:

- `10.0.0.0/8`
- `172.16.0.0/12`
- `192.168.0.0/16`

As a general rule, it is good practice to prevent network traffic intended for RFC 1918 subnets from leaving the firewall via the WAN interface. This avoids unnecessary traffic on the WAN link and also provides a small security benefit by keeping information about the LAN network behind the firewall.

An example where this rule can be helpful is if a machine on the local LAN (e.g. `192.168.1.5`) is configured to access private LAN addresses that are routed across a VPN tunnel (e.g. `192.168.100.0/24`). If the VPN link were to go down, the firewall would no longer have an active route for `192.168.100.0/24` and a packet intended for `192.168.100.0/24` will be routed out the WAN interface using the default route. This could potentially provide information about the private LAN to someone with access to the ISP network. A malicious user could even set up an impostor machine on the WAN with a `192.168.100.0/24` address and pretend to be a machine on the inactive VPN link.

While the chance of this being a problem is small, the probability of unintentional RFC 1918 traffic routing through the WAN interface will increase for installations with more complex LAN topologies, numerous users (typos, etc.), or routes that may frequently change (VPN, etc.). In these scenarios, it may be beneficial to add a firewall rule preventing RFC 1918 traffic from being routed out of the WAN interface.

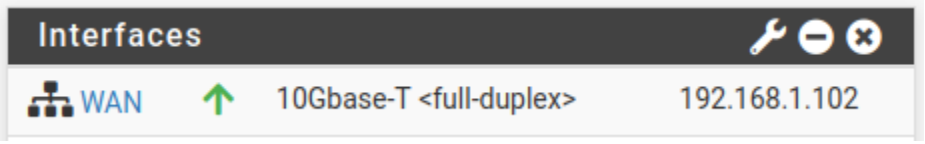
35.61.1 Scenarios where RFC 1918 addresses should NOT be blocked on the WAN interface

The default configuration of pfSense software will not block RFC 1918 addresses routed from the LAN subnet to the outside WAN because there are two common scenarios where blocking this traffic is not desirable:

ISP assigns an RFC 1918 address to end users
Some ISPs assign private network addresses to their customers and perform their own NAT for customer traffic to the public internet. Verify this by looking at the WAN interface IP address on the dashboard. If the assigned address is from one of the private IP ranges listed above, RFC 1918 traffic should NOT be blocked.

The firewall is behind another firewall or router
In this case, pfSense software performs NAT for devices connected to the LAN. The WAN interface forwards traffic to the upstream device where it undergoes a second NAT operation before entering the public internet. This is verified using the same steps as above – if the WAN IP address is from the RFC 1918 range, do NOT block this traffic from exiting the WAN

This is an example of an RFC 1918 address assigned to the WAN:



Warning: If either of these scenarios apply to this installation of pfSense software, do NOT add additional RFC 1918 traffic blocking to the WAN interface as this may prevent LAN users from accessing the WAN.

35.61.2 Steps to block RFC 1918 traffic from leaving the WAN interface

For installations where the above scenarios do not apply an additional firewall rule can be put in place to prevent RFC 1918 traffic from leaking out of the WAN interface. This provides a small increase in security and privacy by preventing information about the local LAN from being routed further upstream to the ISP.

To add a block rule for RFC 1918 traffic:

- Navigate to **Firewall > Aliases**
- Click **Add** to create a new alias
- Configure the alias with these settings:

  Name
  `private_networks` (Or another relevant name such as `RFC1918`)

  Description
  `RFC 1918 Private Networks`

  Type
  Network(s)

  Network or FQDN
  Add entries for each of the private RFC 1918 subnets:

  - `10.0.0.0/8`

  - `172.16.0.0/12`
  - `192.168.0.0/16`

- Click **Save**
- Navigate to **Firewall > Rules**, **Floating** tab
- Click  to add a new rule to the top of the list
- Configure the rule with these settings:

  Action
  Reject

  Quick
  Checked

  Interface
  WAN

  Optionally select multiple WAN interfaces or interface groups here, do NOT select the local LAN

  Direction
  out

  Address Family
  IPv4

  Protocol
  any

  Source
  any

  Destination
  Address or Alias, `private_networks`

- Click **Save**
- Click **Apply Changes**

Lastly, verify that local LAN and internet connectivity are still functional.

35.61.3 Notes

Adding this rule to the firewall will block access to bridge devices like cable modems or upstream routers outside the WAN interface. For example, many cable modems use an IP address of `192.168.100.1` by default. This may or may not be desirable behavior for users. The RFC 1918 firewall rule needs to be bypassed with a specific pass rule above it, or disabled, if clients inside the LAN require access to this type of device.

On the interface options (**Interfaces > WAN**, for example) there is an option to *Block private networks*. This is a rule blocking **inbound** traffic, not outbound like the rule described here. As long as the firewall is not behind a WAN that uses private addressing, both types of rules are desirable and should be enabled.

35.62 Routing Public IP Addresses

This section covers the routing of public IP addresses where a public IP subnet is assigned to an internal interface on a single firewall deployment.

See also:

If a High Availability cluster is in use, see *High Availability Configuration Example without NAT*.

35.62.1 IP Assignments

At least two public IP subnets must be assigned by the ISP. One is for the WAN of the firewall, and one for the inside interface. This is commonly a /30 subnet for the WAN, with a second subnet assigned for the internal interface. This example will use a /30 on WAN as shown in Table *WAN IP Block* and a /29 public subnet on an internal OPT interface as shown in Table *Inside IP Block*.

Table 22: WAN IP Block

198.51.100.64/30	
IP Address	Assigned To
198.51.100.65	ISP router (pfSense® default gateway)
198.51.100.66	pfSense WAN interface IP address

Table 23: Inside IP Block

192.0.2.128/29	
IP Address	Assigned To
192.0.2.129	pfSense OPT interface
192.0.2.130	Internal hosts
192.0.2.131	
192.0.2.132	
192.0.2.133	
192.0.2.134	

35.62.2 Interface Configuration

First configure the WAN and OPT interfaces. The LAN interface can also be used for public IP addresses if desired. In this example, LAN is a private IP subnet and OPT1 is the public IP subnet.

Configure WAN

Add the IP address and gateway accordingly. Figure *WAN IP and Gateway Configuration* shows the WAN configured as shown in Table *WAN IP Block*.



Fig. 95: WAN IP and Gateway Configuration

Configure OPT1

Now enable OPT1, optionally change its name, and configure the IP address and subnet mask. Figure *Routing OPT1 Interface Configuration* shows OPT1 configured as shown in Table *Inside IP Block*.

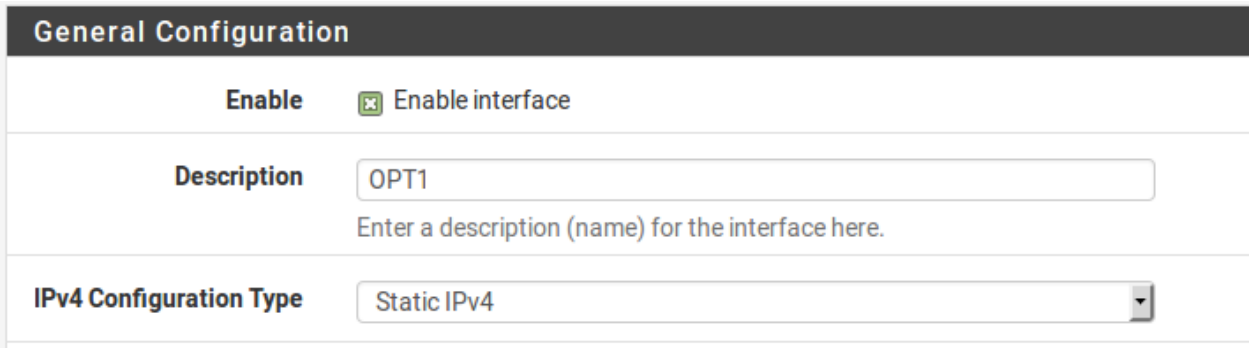


Fig. 96: Routing OPT1 Interface Configuration



Fig. 97: Routing OPT1 IP Address Configuration

35.62.3 NAT Configuration

The default of translating internal traffic to the WAN IP must be overridden when using public IP addresses on an internal interface.

- Browse to **Firewall > NAT**
- Click the Outbound tab
- Select Hybrid Outbound NAT rule generation
- Click **Save**
- Click  to add a new rule to the top of the list with the following settings:

Do not NAT
Checked, so that NAT will be disabled

Interface
WAN

Protocol
Any

Source
Network, enter the local public IP subnet, `192.0.2.128/29`

Destination
Any

- Click **Save**

This will override the default automatic rules which translate all traffic from local interfaces leaving the WAN interface to the WAN IP address. Traffic sourced from the OPT1 network `192.0.2.128/29` is not translated because of the manually added rule excluding it from NAT. This configuration maintains the automatic behavior for other internal interfaces, so that the advantages of automatic outbound NAT rules are not lost. This configuration is shown in Figure *Outbound NAT Configuration*.

If public IP addresses are used on **all** local interfaces, then set **Disable Outbound NAT** rather than using Hybrid mode.

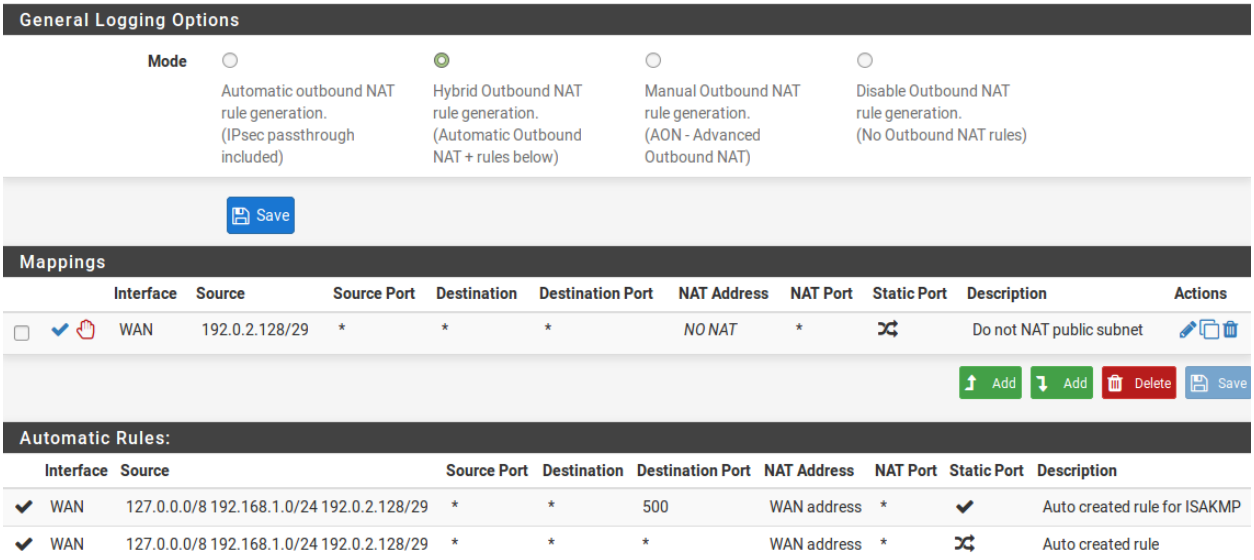


Fig. 98: Outbound NAT Configuration

35.62.4 Firewall Rule Configuration

The NAT and IP address configuration is now complete. Firewall rules will need to be added to permit outbound and inbound traffic. Figure *OPT1 Firewall Rules* shows a DMZ-like configuration, where all traffic destined for the LAN subnet is rejected, DNS and pings to the OPT1 interface IP address are permitted, and HTTP is allowed outbound.

To allow traffic from the Internet to the public IP addresses on an internal interface, add rules on the WAN using the public IP addresses as the **Destination**. Figure *WAN Firewall Rules* shows a rule that allows HTTP to 192.0.2.130, one of the public IP addresses on the internal interface as shown in Table *Inside IP Block*.

After configuring the firewall rules as desired, the setup is complete.

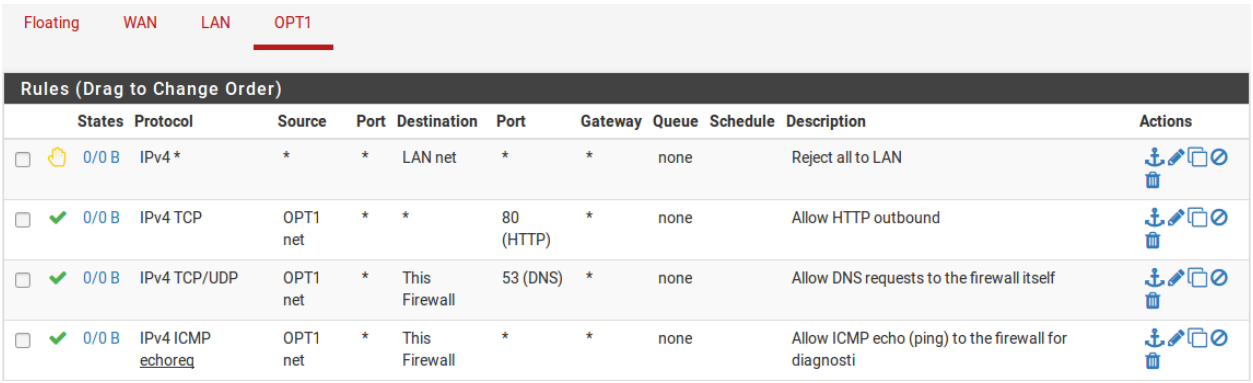


Fig. 99: OPT1 Firewall Rules

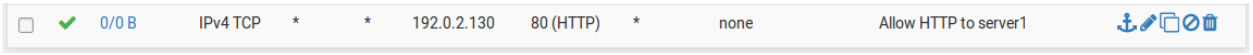


Fig. 100: WAN Firewall Rules

Note: Traffic will flow from LAN to this public subnet by default without NAT. If this behavior is not desired, adjust the LAN firewall and NAT rules accordingly. Additionally, policy routing may need to be bypassed to allow from LAN to this interface.

35.63 Accessing Files on the Firewall with SCP

Files may be transferred to and from the firewall with `scp`, which is part of the functionality that comes with having ssh access enabled.

See also:

- *Secure Shell (SSH)*
- *Granting Users Access to SSH*

To connect to the firewall with SCP for file transfers, use the `root` account with the same credentials as `admin`, or a user account with sufficient privileges.

Note: In most recent versions, the `admin` account can also be used for `scp`, but using `root` is still the best practice.

Users with shell access may transfer files, as well as users with the *User - System - Copy files* privilege.

Note: Users other than `root` can only transfer or write files for which their account has permission to read or modify.

Any SCP/SFTP-compatible program may be used to transfer files. Popular choices include `scp`, FileZilla, and WinSCP.

An example: Logged into a FreeBSD machine, copy a file from `/tmp` on a remote firewall to the current local working directory:

```
myuser@somebox:~/$ scp root@192.168.1.1:/tmp/lan-traffic.rrd-4h.png .
```

35.63.1 Troubleshooting

If `scp` is failing with the error `Illegal option -t` or printing a usage message, check that the connection is using the `root` account, as mentioned above.

35.64 Granting Users Access to SSH

This recipe explains how to enable *Secure Shell (SSH)* access to the firewall.

SSH is typically used for debugging and troubleshooting, but has many other useful purposes.

Note: The SSH daemon is not required by the firewall for operation, so it is disabled by default.

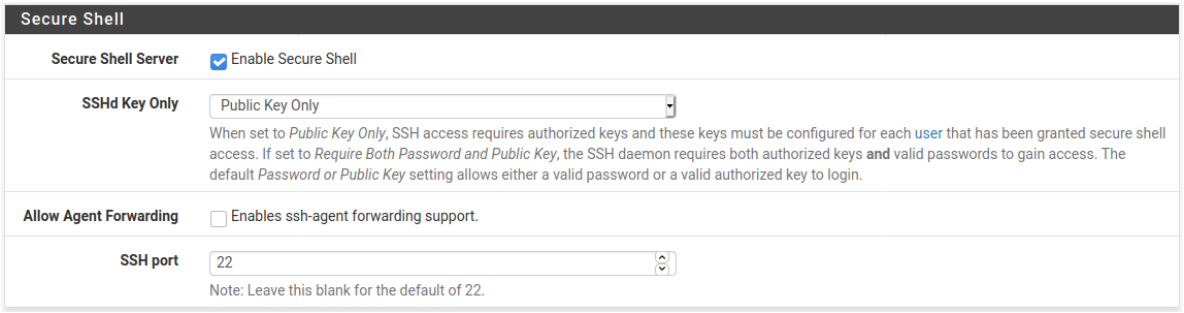
See also:

Secure Shell (SSH)

35.64.1 Enable SSH via GUI

This example enables SSH access using only public key authentication, which is more secure than allowing access by password alone.

- Navigate to **System > Advanced**, **Admin Access** tab
- Check **Enable Secure Shell**
- Set **SSHd Key Only** to *Public Key Only* to allow only key-based SSH authentication



- Enter a port number in **SSH Port** if the SSH daemon should listen on a non-default port

  Leave the field blank for the daemon to use port `22`

- Click **Save**

35.64.2 SSH Keys

When the SSH daemon is set for key-based authentication, it uses the keys defined on user accounts. Add keys to individual user accounts under **System > User Manager**. The `admin` user and `root` user share keys.

> **Warning:** Do not attempt to manage keys from the shell directly.

See also:

Manage Local Users

35.64.3 Enable SSH via Console

Connect to the console (VGA or Serial) and use option `14` to enable or disable SSH.

To change the port number or key authentication options, use the GUI as directed above.

35.64.4 SSH Daemon Security

With a default ruleset, SSH may only be accessed by clients on the LAN. If SSH access must be allowed for clients the WAN, the best practice is to restrict access to Key-based authentication to avoid issues with brute force attacks. Moving the daemon to an alternate port is also a good practice, but moving the port alone is not sufficient protection.

The firewall will automatically block users who attempt to authenticate unsuccessfully. This behavior, and settings to control it, are described in *Login Protection*.

If password authentication is active, ensure that all user accounts with shell access have strong passwords that cannot be easily guessed.

See also:

See *Best Practices for SSH* for more on SSH security.

35.64.5 User Access

By default, only `admin` and `root` have SSH access. Additional users with limited access may be granted the *User - System - Shell account access* privilege to login via SSH.

Note: Additional users do not have full root privileges in the shell, so the system does not display the console menu for those users. Many commands and other files are inaccessible as well. For a normal user to get much use from the shell, the *Sudo Package* can delegate additional privileges to run commands as *root* or other users.

35.64.6 SCP File Transfers

For information on using SCP file transfers via SSH, see *Accessing Files on the Firewall with SCP*.

35.65 Configuring Switches with VLANs

This section provides guidance on configuring a few varieties of switches for use with VLANs. This offers generic guidance that will apply to most if not all 802.1Q capable switches, then goes on to cover configuration on specific switches from Cisco, HP, Netgear, and Dell.

Note: This is the bare minimum configuration needed for VLANs to function, and it does not necessarily show the ideal secure switch configuration for any specific environment. An in depth discussion of switch security is outside the scope of this documentation.

35.65.1 Switch configuration overview

Generally three or four things must be configured on VLAN capable switches:

1. **Add/define the VLANs**

   Most switches have a means of defining a list of configured VLANs, and they must be added before they can be configured on any ports.

2. **Configure the trunk port**

   The port to which the firewall running pfSense® software will be connected must be configured as a trunk port, tagging all possible VLANs on the interface.

3. **Configure the access ports**

   Configure ports for internal hosts as access ports on the desired VLANs, with **untagged** VLANs.

4. **Configure the Port VLAN ID (PVID)**

   Some switches require configuring the PVID for access ports. This specifies which VLAN to use for the traffic entering that switch port. For some switches this is a one-step process, by configuring the port as an access port on a particular VLAN, it automatically tags traffic coming in on that port. Other switches require this to be configured in one or two places. Check the switch documentation for details if it is not one detailed in this chapter.

35.65.2 Cisco IOS based switches

Configuring and using VLANs on Cisco switches with IOS is a fairly simple process, taking only a few commands to create and use VLANs, trunk ports, and assigning ports to VLANs. Many switches from other vendors behave similarly to IOS, and will use nearly the same if not identical syntax for configuration.

Create VLANs

VLANs can be created in a standalone fashion, or using VLAN Trunk Protocol (VTP). Using VTP may be more convenient, as it will automatically propagate the VLAN configuration to all switches on a VTP domain, though it also can create its own security problems and open up possibilities for inadvertently wiping out the VLAN configuration.

With VTP, to add another VLAN it only needs to be configured on a single switch, and then all other trunked switches in the group can assign ports to that VLAN. If VLANs are configured independently, they must be added to each switch by hand. Refer to Cisco's documentation on VTP to ensure a secure configuration use used, and that it is not prone to accidental destruction.

In a network with only a few switches where VLANs do not change frequently, VTP may be overkill and avoiding it will also avoid its potential downfalls.

Standalone VLANs

To create standalone VLANs:

```
sw# vlan database
sw(vlan)# vlan 10 name "DMZ Servers"
sw(vlan)# vlan 20 name "Phones"
sw(vlan)# exit
```

VTP VLANs

To set up a switch for VTP and VLANs, create a VTP database on the master switch and then create two VLANs:

```
sw# vlan database
sw(vlan)# vtp server
sw(vlan)# vtp domain example.com
sw(vlan)# vtp password SuperSecret
sw(vlan)# vlan 10 name "DMZ Servers"
sw(vlan)# vlan 20 name "Phones"
sw(vlan)# exit
```

Configure Trunk Port

For handing off VLANs to pfSense software a switch port not only has to be in trunk mode, but also must be using 802.1q tagging. This can be done like so:

```
sw# configure terminal
sw(config)# interface FastEthernet 0/24
sw(config-if)# switchport mode trunk
sw(config-if)# switchport trunk encapsulation dot1q
```

Note: On some newer Cisco IOS switches, the Cisco-proprietary ISL VLAN encapsulation method is deprecated. If a switch does not allow the `encapsulation dot1q` configuration option, it can only use 802.1Q and the encapsulation does not need to be specified.

Add Ports to the VLAN

To add ports to these VLANs, assign them as follows:

```
sw# configure terminal
sw(config)# interface FastEthernet 0/12
sw(config-if)# switchport mode access
sw(config-if)# switchport access vlan 10
```

35.65.3 Cisco CatOS based switches

Creating VLANs on CatOS is a little different, though the terminology is the same as using VLANs under IOS. Standalone VLANs and VTP are both possible to maintain the VLAN database:

```
# set vtp domain example mode server
# set vtp passwd SuperSecret
# set vlan 10 name dmz
# set vlan 20 name phones
```

Then configure a trunk port to automatically handle every VLAN:

```
# set trunk 5/24 on dot1q 1-4094
```

Then add ports to the VLAN:

```
# set vlan 10 5/1-8
# set vlan 20 5/9-15
```

35.65.4 HP ProCurve switches

HP ProCurve switches only implement 802.1q trunking, so no configuration is needed for encapsulation. First, ssh or telnet into the switch and bring up the management menu.

Enable VLAN Support

First, VLAN support needs to be enabled on the switch if it is not already:

1. Choose **Switch configuration**
2. Choose **Advanced Features**
3. Choose **VLAN Menu...**
4. Choose **VLAN Support**
5. Set **Enable VLANs** to *Yes* if it is not already, and choose a number of VLANs. Each time this value is changed the switch must be restarted, so ensure it is large enough to contain as many VLANs as necessary.
6. Restart the switch to apply the changes.

Create VLANs

Before the VLANs can be assigned to ports, The VLANs must be created. At the switch configuration menu:

1. Choose **Switch configuration**
2. Choose **Advanced Features**
3. Choose **VLAN Menu…**
4. Choose **VLAN Names**
5. Choose **Add**
6. Enter the **VLAN ID**, `10`
7. Enter the **name**, `DMZ`
8. Choose **Save**
9. Repeat the steps from **Add** to **Save** for any remaining VLANs

Assigning Trunk Ports to VLANs

Next, configure the trunk port for the firewall as well as any trunk ports going to other switches containing multiple VLANs.

1. Choose **Switch configuration**
2. Choose **VLAN Menu…**
3. Choose **VLAN Port Assignment**
4. Choose **Edit**
5. Find the port to assign
6. Press **space** on Default VLAN until it shows **No**
7. Move over to the column for each of the VLANs on this trunk port, and Press **space** until it shows **Tagged**. Every VLAN in use must be tagged on the trunk port.

Assigning Access Ports to VLANs

1. Choose **Switch configuration**
2. Choose **VLAN Menu…**
3. Choose **VLAN Port Assignment**
4. Choose **Edit**
5. Find the port to assign
6. Press **space** on **Default VLAN** until it shows **No**
7. Move over to the column for the VLAN to which this port will be assigned
8. Press **space** until it shows **Untagged**.

35.65.5 Netgear Managed Switches

This example is on a GS108Tv1, but other Netgear models are all very similar if not identical. There are also several other vendors including Zyxel who sell switches made by the same manufacturer, using the same web interface with a different logo. Log into the web interface of the switch to start.

Planning the VLAN configuration

Before configuring the switch, several items are required:

1. The number of VLANs to be configured
2. The IDs to use for the VLANs
3. How each switch port needs to be configured

For this example, an 8 port GS108Tv1 is used, and it will be configured as shown in Table *Netgear GS108T VLAN Configuration*.

Table 24: Netgear GS108T VLAN Configuration

Switch port	VLAN mode	VLAN assigned
1	trunk	*10* and *20*, tagged
2	access	*10* untagged
3	access	*10* untagged
4	access	*10* untagged
5	access	*20* untagged
6	access	*20* untagged
7	access	*20* untagged
8	access	*20* untagged

Enable 802.1Q VLANs

To configure the switch to use 802.1Q VLAN trunking:

- Navigate to the **System** menu on the left side of the page
- Click **VLAN Group Setting**, as indicated in Figure *VLAN Group Setting*.



Fig. 101: VLAN Group Setting

- Select IEEE 802.1Q VLAN (Figure *Enable 802.1Q VLANs*).

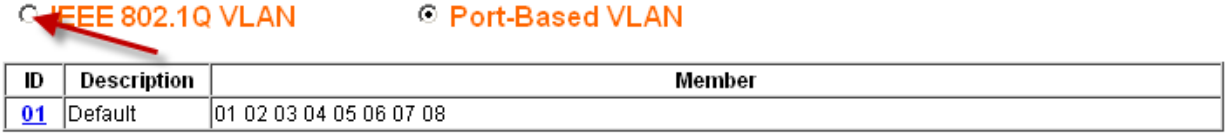


Fig. 102: Enable 802.1Q VLANs

- Click **OK** to confirm the switch to 802.1Q trunking, as shown in Figure *Confirm change to 802.1Q VLAN*.

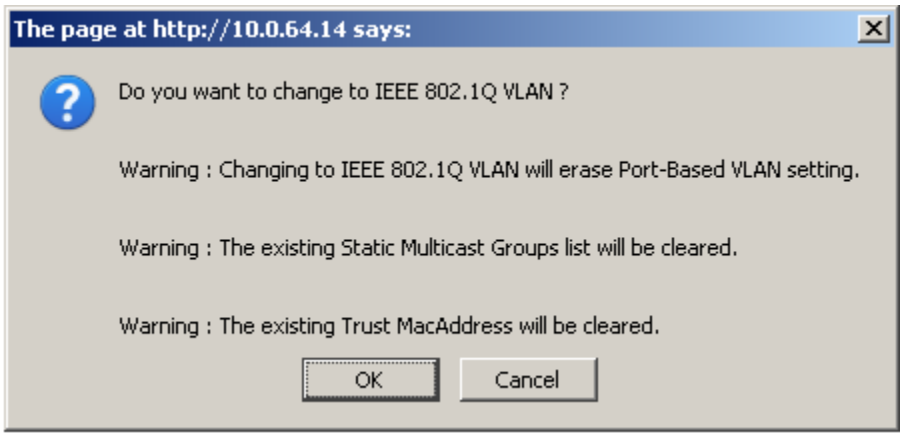


Fig. 103: Confirm change to 802.1Q VLAN

After clicking OK, the page will refresh with the 802.1Q VLAN configuration as shown in Figure *Default 802.1Q Configuration*.

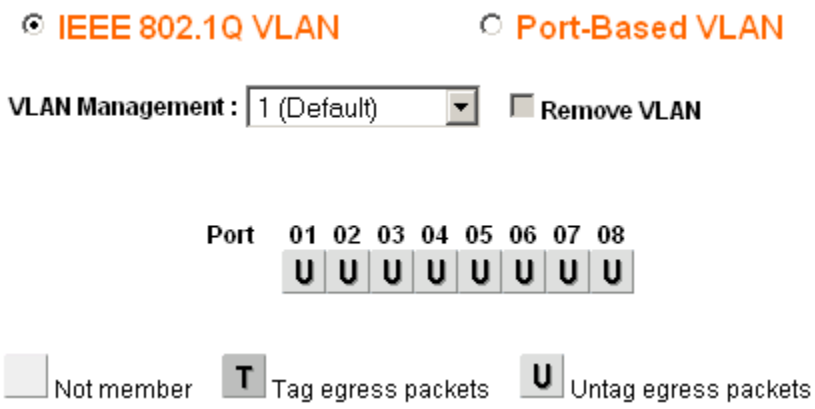


Fig. 104: Default 802.1Q Configuration

Add VLANs

For this example, two VLANs are added with IDs `10` and `20`.

To add a VLAN:

- Click the **VLAN Management** select box
- Click **Add New VLAN** as shown in Figure *Add New VLAN*.
- Enter the VLAN ID for this new VLAN, such as `10`
- Click **Apply**. The VLAN screen is now ready to configure VLAN *10* (Figure *Add VLAN 10*).
- Click **Add New VLAN** again as shown in Figure *Add New VLAN* to add VLAN *20* (Figure *Add VLAN 20*).

Add as many VLANs as needed, then continue to the next section.

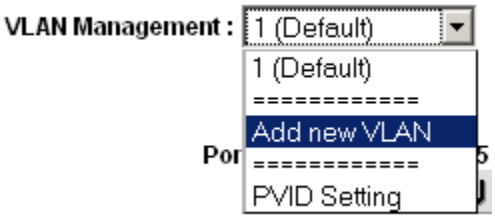


Fig. 105: Add New VLAN

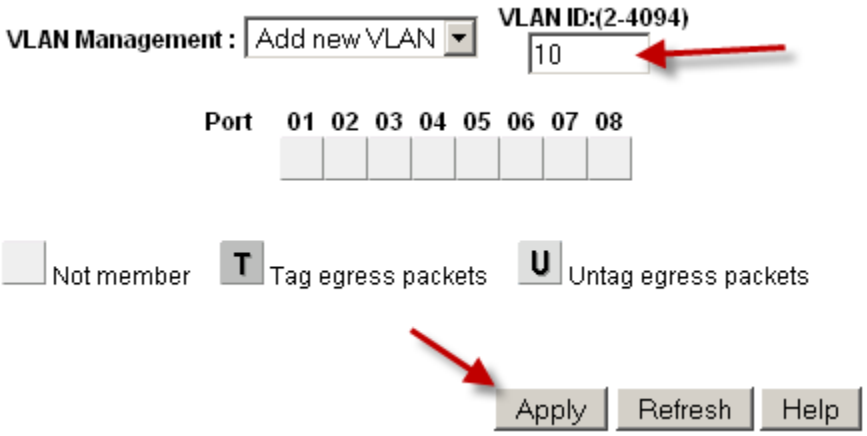


Fig. 106: Add VLAN 10

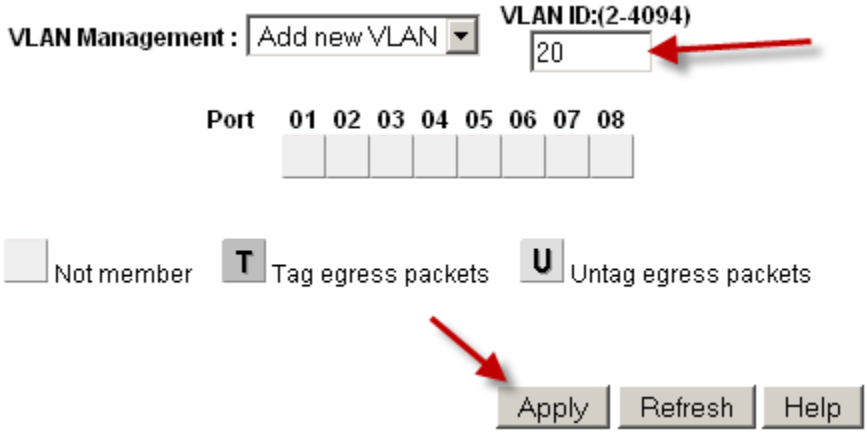


Fig. 107: Add VLAN 20

Configure VLAN tagging

When a VLAN is selected from the **VLAN Management** select box, it shows how that VLAN is configured on each port:

- A **blank** box means the port is not a member of the selected VLAN.
- A box containing **T** means the VLAN is sent on that port with the 802.1Q tag.
- **U** indicates the port is a member of that VLAN and it leaves the port untagged.

The trunk port must have both VLANs added and tagged.

> **Warning:** Do not change the configuration of the port being used to access the web interface of the switch! This will lock the administrator out of the switch. The only means of recovery on the GS108Tv2 is using the **reset to factory defaults** button since it does not have a serial console. For the switches that have serial consoles, keep a null modem cable handy in case network connectivity with the switch is lost. Configuring the management VLAN is covered later in this section.

Click in the boxes beneath the port number as shown in Figure ref:*figure-toggle-vlan-membership* to toggle between the three VLAN options.

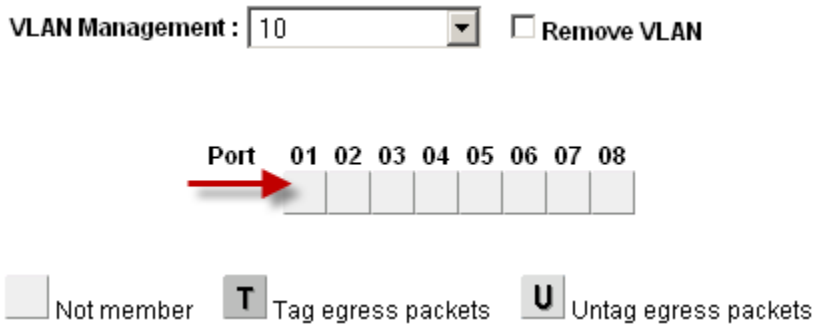


Fig. 108: Toggle VLAN Membership

Configure VLAN 10 membership

Figure *Configure VLAN 10 Membership* shows VLAN **10** configured as outlined in Table *table-netgear-gs108t-vlan-configuration*. The access ports on this VLAN are set to **untagged** while the trunk port is set to **tagged**.

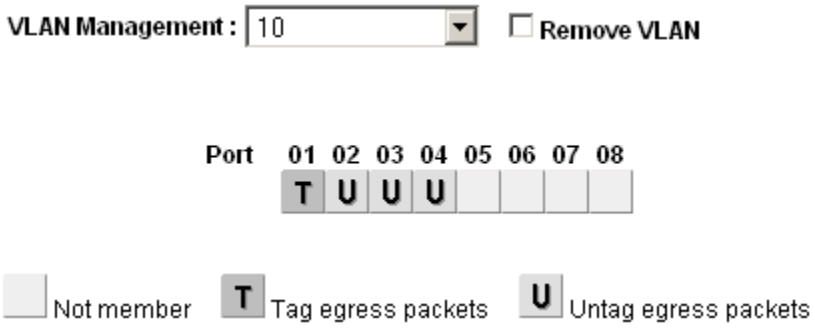


Fig. 109: Configure VLAN 10 Membership

Configure VLAN 20 membership

Select **20** from the VLAN Management select box to configure the port memberships for VLAN **20**.

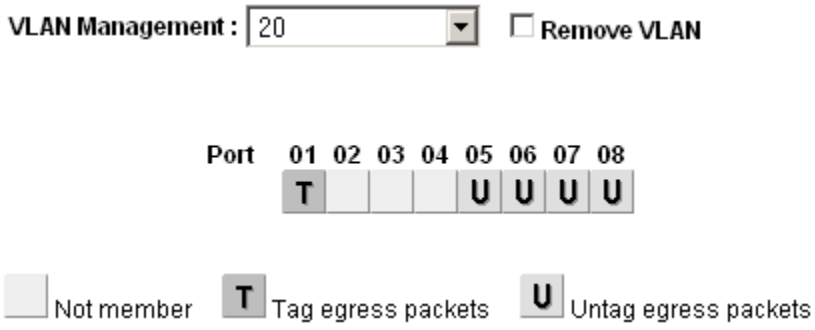


Fig. 110: Configure VLAN 20 Membership

Change PVID

On Netgear switches, in addition to the previously configured tagging settings, the PVID must also be configured to specify the VLAN used for frames entering a port:

- Select **PVID** from the VLAN Management select box as shown in Figure *PVID Setting*.

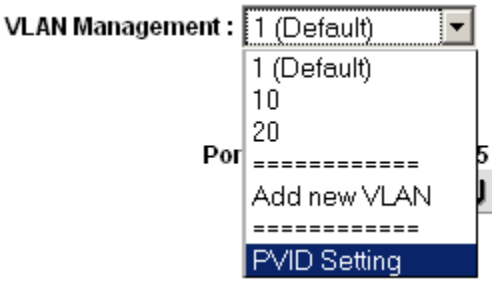


Fig. 111: PVID Setting

The default PVID setting is VLAN 1 for all ports as shown in Figure *Default PVID Configuration*.

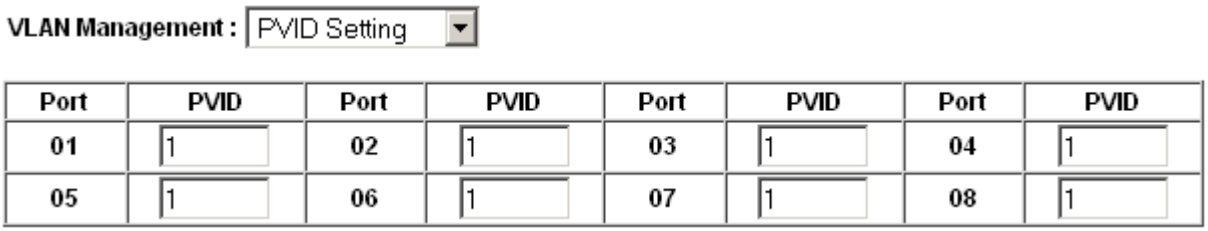


Fig. 112: Default PVID Configuration

- Change the PVID for each access port, but leave the trunk port and port used to access the switch management interface set to 1.

Figure *VLAN 10 and 20 PVID Configuration* shows the PVID configuration matching the port assignments shown in Table *Netgear GS108T VLAN Configuration*, with port 8 being used to access the switch management interface.

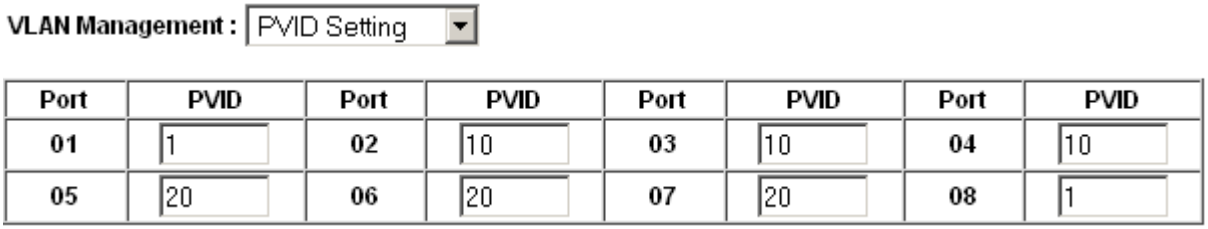


Fig. 113: VLAN 10 and 20 PVID Configuration

- Apply changes when finished

Remove VLAN 1 configuration

By default, all ports are members of VLAN 1 with untagged egress frames. To remove VLAN 1 from the other ports:

- Select *1 (Default)* from the **VLAN Management** select box
- Remove VLAN 1 from all ports except the one used to manage the switch and the trunk port, to avoid being disconnected.

  In this example, port 8 is used to manage the switch. When finished, the screen will look like Figure *Remove VLAN 1 Membership*.

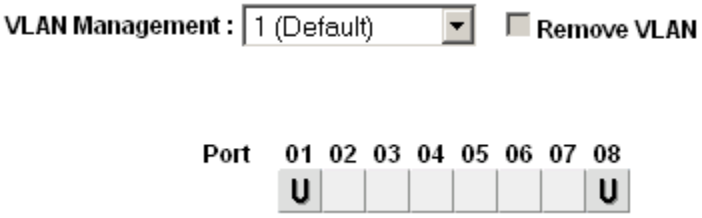


Fig. 114: Remove VLAN 1 Membership

- Apply changes when finished

Verify VLAN functionality

Configure VLANs on pfSense, including the DHCP server on the VLAN interfaces if needed. Plug systems into the configured access ports and test connectivity. If everything works as desired, continue to the next step. If things do not work as intended, review the tagging and PVID configuration on the switch, and the VLAN configuration and interface assignments on pfSense software.

35.65.6 Dell PowerConnect managed switches

The management interface of Dell switches varies slightly between models, but the following procedure will accommodate most models. The configuration is quite similar in style to Cisco IOS.

First, create the VLANs:

```
console# config
console(config)# vlan database
console(config-vlan)# vlan 10 name dmz media ethernet
console(config-vlan)# vlan 20 name phones media ethernet
console(config-vlan)# exit
```

Next, set up a trunk port:

```
console(config)# interface ethernet 1/1
console(config-if)# switchport mode trunk
console(config-if)# switchport allowed vlan add 1-4094 tagged
console(config-if)# exit
```

Finally, add ports to the VLANs:

```
console(config)# interface ethernet 1/15
console(config-if)# switchport allowed vlan add 10 untagged
console(config-if)# exit
```

35.66 Using the Shaper Wizard to Configure ALTQ Traffic Shaping

The easiest way to get started with traffic shaping is by using the wizard for the first time, which guides administrators through the shaper configuration process.

Tip: Due to the complexity of the shaper queues and rules, starting from scratch is quite complicated. If a firewall needs custom rules, step through the wizard and approximate the requirements, then make custom rules afterward.

Each step of the wizard sets up unique queues and rules that control what traffic is assigned into those queues. To configure everything manually, specify the WAN speed at the first screen, then click **Next** through all the remaining steps. The wizard requires options to be enabled on at least one step, but it does not matter which step.

Note: Completing the wizard and clicking **Finish** at the end will replace **all** existing shaper queues and floating rules created by the wizard, including those cloned from wizard rules, with the queues and rules from the new wizard configuration.

35.66.1 Choosing a Wizard

To get started with the Traffic Shaping Wizard, navigate to **Firewall > Traffic Shaper** and click the **Wizards** tab. This page displays a list of available traffic shaper wizards, including:

Multiple LAN/WAN
: Used when the firewall has one or more WANs and one or more LANs. This is the most common wizard and it covers almost every scenario.

Dedicated Links
: Used when specific LAN+WAN pairings should be accounted for in the shaper configuration.

35.66.2 Starting the Wizard

Each wizard name is followed by the filename of the wizard, which is a link. Click the link to start the wizard. This example uses the **Multiple LAN/WAN** wizard, so click `traffic_shaper_wizard_multi_all.xml`.

Next, the wizard starts and the first step prompts for the number of WAN and LAN type connections on the firewall, as in Figure *Entering the Interface Count*.

- Enter the number of WAN-type connections on the firewall. These are connections with a gateway configured on the interface, or dynamic WAN type interfaces such as DHCP or PPPoE
- Enter the number of LAN type connections. These are local network interfaces without a gateway on the interface
- Click **Next** to proceed with the next step

In this example the firewall only has one WAN and one LAN interface.

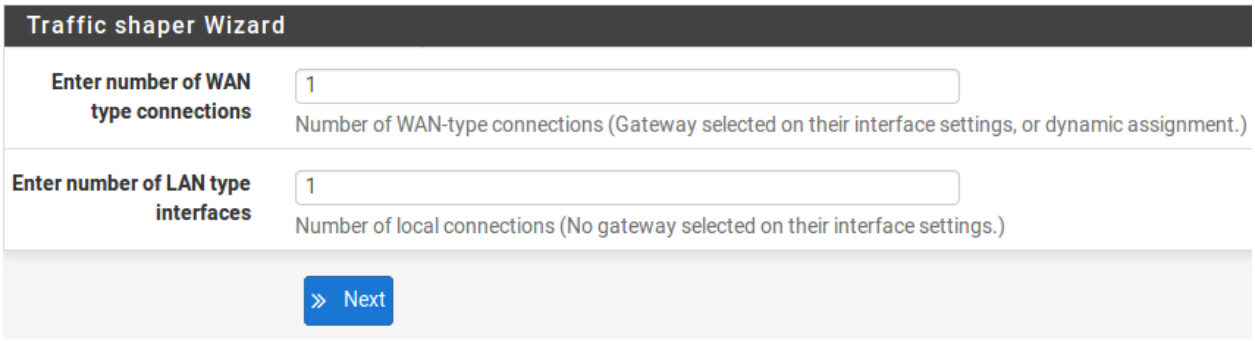


Fig. 115: Entering the Interface Count

35.66.3 Networks and Speeds

This step, shown in Figure *Shaper Configuration*, defines the network interfaces that will be the inside and outside from the point of view of the shaper, along with the **Download** and **Upload** speeds for a given WAN. When the firewall has more than one interface of a given type, the wizard displays multiple sections on the page to handle each one individually.

In addition to the interfaces and their speeds, select an ALTQ **Scheduler** (*ALTQ Scheduler Types*) for the WAN(s) and LAN(s). Use the same scheduler on every interface.

Depending on the connection type, the true link speed may not be the actual usable speed. In the case of PPPoE, the circuit has not only PPPoE overhead, but also overhead from the underlying ATM network link being used in most PPPoE deployments. By some calculations, between the overhead from ATM, PPPoE, IP, and TCP, the circuit may lose as much as 13% of the advertised link speed. When in doubt of what to set the speed to, be conservative. Reduce

by 10-13% and work it back up to larger values. If the firewall has a 3Mbit/s line, set it for about `2.7` Mbit/s and then test. The speed on the resulting parent queue can be edited later to adjust the bandwidth. If it has a low value, the connection will be maxed out at exactly the defined speed. Nudge it up higher until the firewall no longer sees any performance gains.

Interface speeds can be specified in *Kbit/s*, *Mbit/s*, or *Gbit/s*, but use the same units on every page.

- Choose an **Interface** and **Scheduler** for each LAN-type interface (e.g. *LAN*, *PRIQ*)
- Choose an **Interface** and **Scheduler** for each WAN-type interface (e.g. *WAN*, *PRIQ*)
- Define the **Upload** speed and units for each WAN-type interface (e.g. `1`, *Mbit/s*)
- Define the **Download** speed and units for each WAN-type interface (e.g. `10`, *Mbit/s*)
- Click **Next** to proceed with the next step

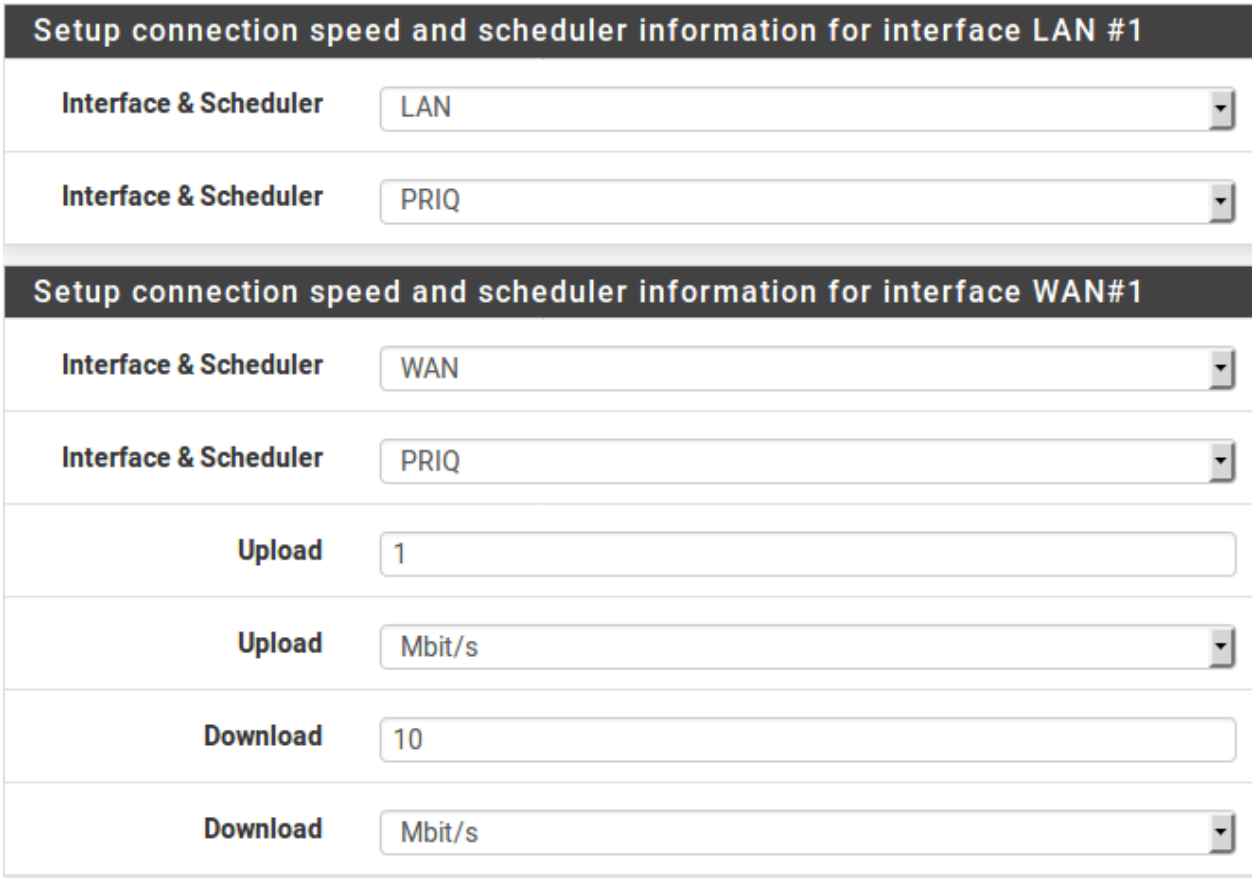


Fig. 116: Shaper Configuration

35.66.4 Voice over IP

The wizard contains several options for handling VoIP call traffic, shown in Figure *Voice over IP*. Prioritizing Voice over IP traffic sets up queues and rules to give priority to VoIP calls and related traffic. This behavior can be fine-tuned by the other settings on this step of the wizard.

Enable
A checkbox to enable the VoIP settings on this step. When unchecked, the options are disabled and these queues and rules will not be added by the wizard.

Provider
There are a few well-known providers including *Vonage*, *Voicepulse*, *PanasonicTDA*, and *Asterisk* servers. If the VoIP provider for this site is not in the list, choose *Generic*. This choice sets up rules based on the ports and protocols known to be used by these providers, rather than matching by address.

Note: This choice matches based on SIP and RTP ports, among others, therefore it can match traffic from other sources as well if they use the same ports as the selected service.

Upstream SIP Server
The IP of the upstream PBX or SIP trunk, or an alias containing the IP addresses or networks for the SIP trunk(s). When set, this overrides the **Provider** field and will instead match traffic based on these addresses.

Note: This choice matches all UDP traffic to and from the specified address(es). In most cases this is OK, but if there are other Non-VoIP UDP-based services on the same remote address, it could match that traffic as well. Such cases are rare, however, so this option tends to be more reliable than matching by port.

WAN Connection Upload
The amount of upload bandwidth to guarantee for VoIP devices. This will vary based on how many VoIP devices are on the network and how much bandwidth each session requires. This setting is used by HFSC and CBQ, and should be left blank for PRIQ.

Note: The bandwidth reservation for a service such as VoIP cannot exceed 30% of the available bandwidth on the link. For example, on a 10Mbit/s link, the shaper cannot reserve more than 3Mbit/s.

LAN Connection Download
The amount of download bandwidth to guarantee for VoIP devices. This setting is used by HFSC and CBQ, and should be left blank for PRIQ.

Note: The best practice is to use the **remote** SIP trunk or PBX address because otherwise the shaper may not be able to match traffic properly. For example, using the IP addresses of phones the shaper may only match traffic in one direction, or not at all. This is due to the way the shaper matches traffic with floating rules in an outbound direction. NAT applies before traffic is matched when exiting a WAN, so the shaper rules cannot match outbound connections based on local private IP addresses.

To use these options:

- Check **Prioritize Voice over IP traffic**
- Pick **ONE** of the following:

  - Choose a **Provider** from the list **OR**
  - Enter an **Upstream SIP Server** address or alias containing a **remote** SIP trunk or PBX
- Leave **Upload** and **Download** blank if using PRIQ, otherwise enter an appropriate **Upload** or **Download** value for each connection
- Click **Next** to proceed with the next step

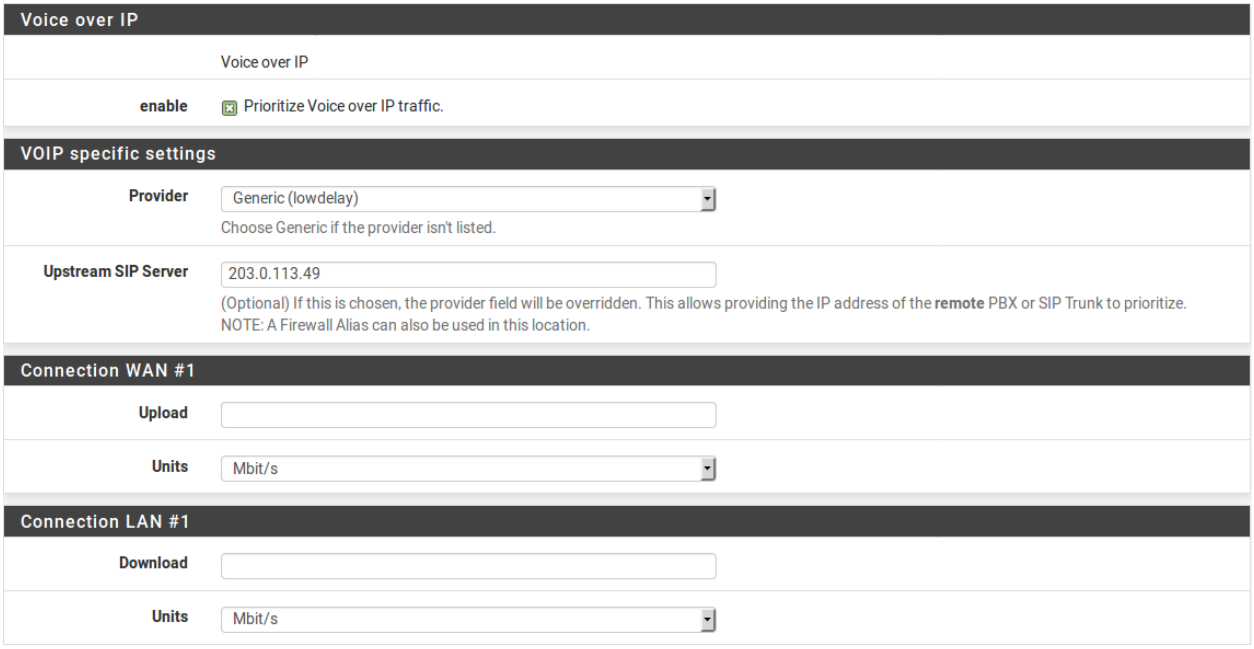


Fig. 117: Voice over IP

35.66.5 Penalty Box

The penalty box, depicted in Figure *Penalty Box*, is a place to relegate misbehaving users or devices that would otherwise consume undesirable amounts of bandwidth. These devices are assigned a hard bandwidth cap which they cannot exceed.

Enable
A checkbox to enable the Penalty Box settings on this step. When unchecked, the options are disabled and these queues and rules will not be added by the wizard.

Address
The IP address to penalize, or an alias containing multiple addresses to penalize.

Bandwidth
The amount of bandwidth that **Address** can consume, at most.

To use these options:

- Check **Penalize IP or Alias**
- Enter an IP address or Alias in the **Address** box
- Enter the **Bandwidth** limit
- Choose the correct units for the **Bandwidth** limit
- Click **Next** to proceed with the next step

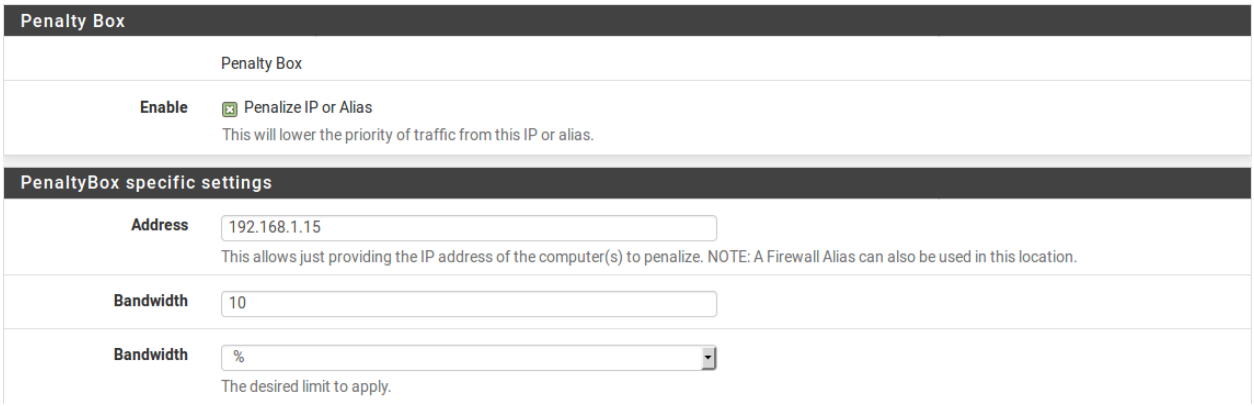


Fig. 118: Penalty Box

35.66.6 Peer-to-Peer Networking

The next step, shown in Figure *Peer-to-Peer Networking*, sets controls for many Peer-to-Peer (P2P) networking protocols. By design, P2P protocols will utilize all available bandwidth unless limits are put in place. If P2P traffic will be present on a network, the best practice is to ensure it will not degrade other traffic.

Note: P2P protocols deliberately attempt to avoid detection. BitTorrent is especially guilty of this behavior. It often utilizes non-standard or random ports, or ports associated with other protocols. Identifying all P2P traffic can be difficult or impossible.

Enable
A checkbox to enable the P2P traffic settings on this step. When unchecked, the options are disabled and these queues and rules will not be added by the wizard.

Peer-to-Peer Catch All
Causes any unrecognized traffic to be assumed as P2P traffic, and such traffic will have its priority lowered accordingly.

Bandwidth
The amount of bandwidth that unclassified traffic can consume, at most, when P2P Catch All is active.

> **Warning:** This option effectively takes over the **Default** traffic shaping queue and lowers its priority. When this option is active, it is critical for all legitimate traffic to be matched by rules that set a priority higher than the priority of the P2P catch all queue.
>
> The **Raise / Lower Other Applications** step of the wizard can help here, but ultimately accomplishing this task frequently requires additional manual rules.

Enable/Disable specific P2P protocols
These options identify various known P2P protocols. The firewall will assign ports and protocols associated with each enabled option as P2P traffic.

To use the options in this step:

- Check **Lower priority of Peer-to-Peer traffic**
- Optionally enable the **p2p Catch All** feature

  - Enter the **Bandwidth** limit for **p2p Catch all**, if enabled
  - Choose the correct units for the **Bandwidth** limit
- Select protocols for the firewall to classify as P2P traffic
- Click **Next** to proceed with the next step

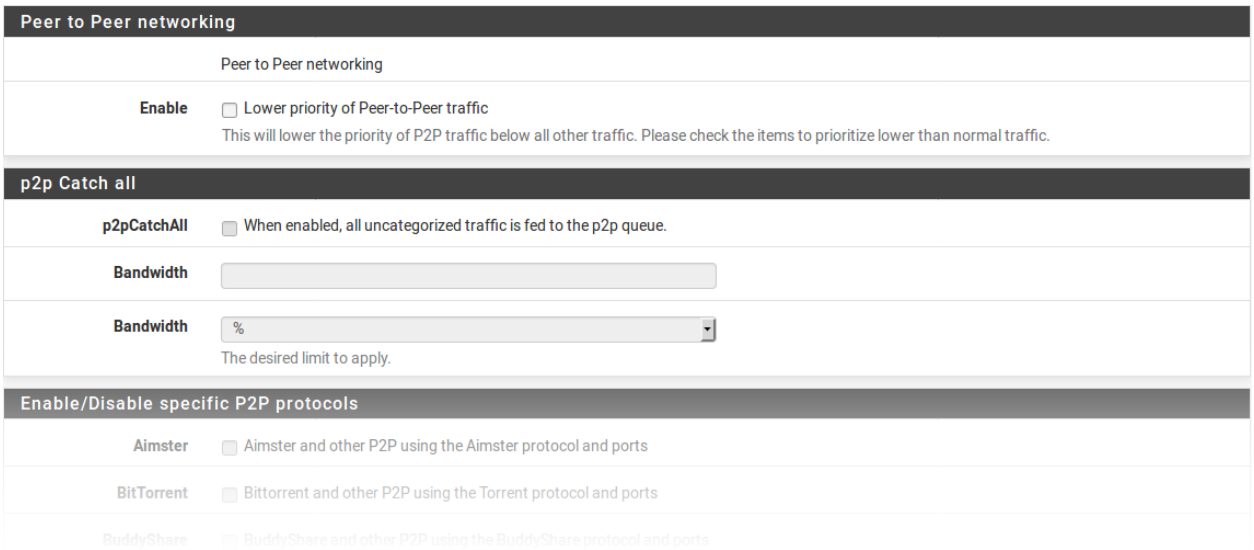


Fig. 119: Peer-to-Peer Networking

35.66.7 Network Games

Online games typically rely on low latency for acceptable player experiences. If a user on the network attempts to download large files or game patches while playing, that traffic can easily drown out the packets associated with the game itself and cause lag or disconnections. If the firewall gives gaming traffic priority, it can ensure that traffic will be delivered first and fastest.

Enable
A checkbox to enable the gaming traffic settings on this step. When unchecked, the options are disabled and these queues and rules will not be added by the wizard.

Enable/Disable specific game consoles and services
These options match traffic for entire game consoles or online services which use common ports and protocols across all, or at least a majority, of their games.

Enable/Disable specific games
These options match traffic for specific games which deviate from the general categories in the previous section.

Tip: To prioritize a game that is not listed, check any other game from the list so that the wizard will create the queues and rules to use as a reference base. After completing the wizard, edit the resulting rules to match the unlisted game.

To use the options in this step:

- Check **Prioritize network gaming traffic**
- Select any games consoles on the network from the list in **Enable/Disable specific game consoles and services**

- Select any games on the network from the list in **Enable/Disable specific games**
- Click **Next** to proceed with the next step

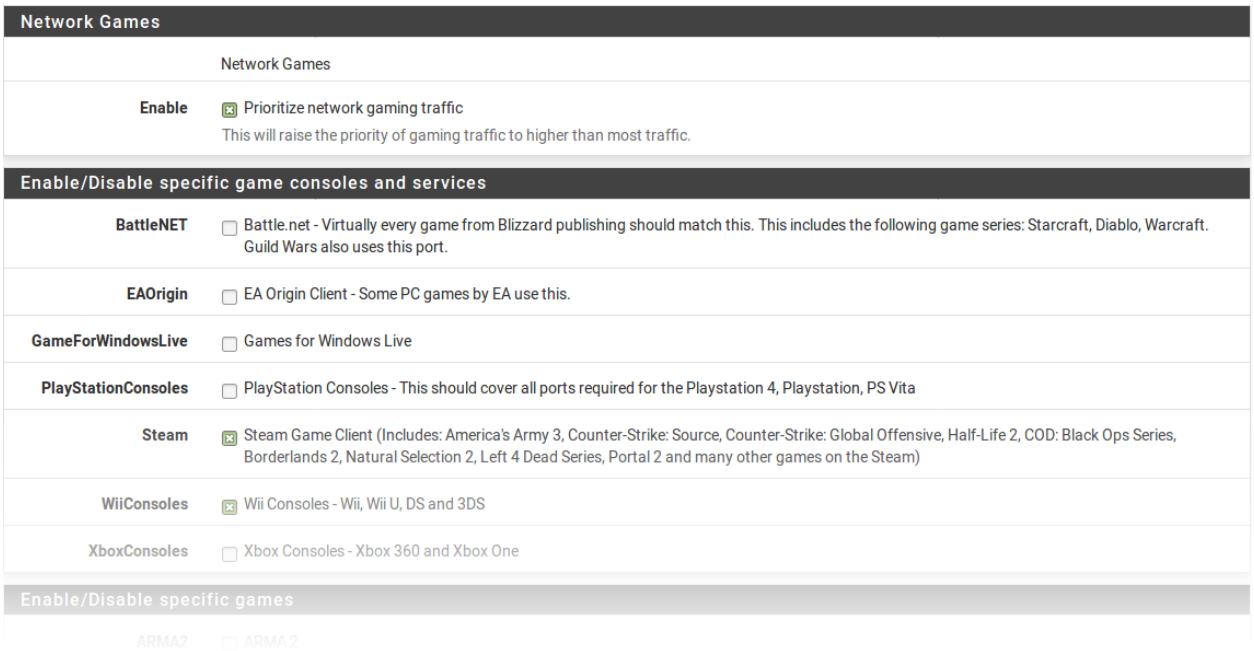


Fig. 120: Network Games

35.66.8 Raising or Lowering Other Applications

The last configuration screen of the shaper wizard, seen in Figure *Raise or Lower Other Applications*, lists a number of other commonly available applications and protocols.

The needs of a particular network dictate how the firewall should handle each protocol. For example, in a corporate environment management may want to lower the priority of non-interactive traffic such as e-mail where a reduction in speed is not usually noticed by users, and they may also want to raise the priority of interactive services like RDP where poor performance is an impediment for employees. In a home, multimedia streaming may be more important, and other services can have their priority lowered by the shaper.

Tip: As with other steps of this shaper wizard, if a protocol is not listed, select a similar protocol and then adjust the rules after completing the wizard.

Enable
A checkbox to enable the settings on this step. When unchecked, the options are disabled and these queues and rules will not be added by the wizard.

Protocol Categories
Each section contains well-known protocols, grouped by their general function.

There are more than 40 protocols to choose from, and each can be given a *Higher priority*, *Lower priority*, or left at the *Default priority*.

Tip: If **p2pCatchAll** is active, the best practice is to use this step to ensure that these other protocols

are recognized and treated normally, rather than penalized by the default p2pCatchAll rule.

To use the options in this step:

- Check **Other networking protocols**
- Locate specific protocols in the list to alter priority.
- For each protocol, choose one of *Higher priority*, *Lower priority*, or leave it at the *Default priority*.
- Click **Next** to proceed with the next step

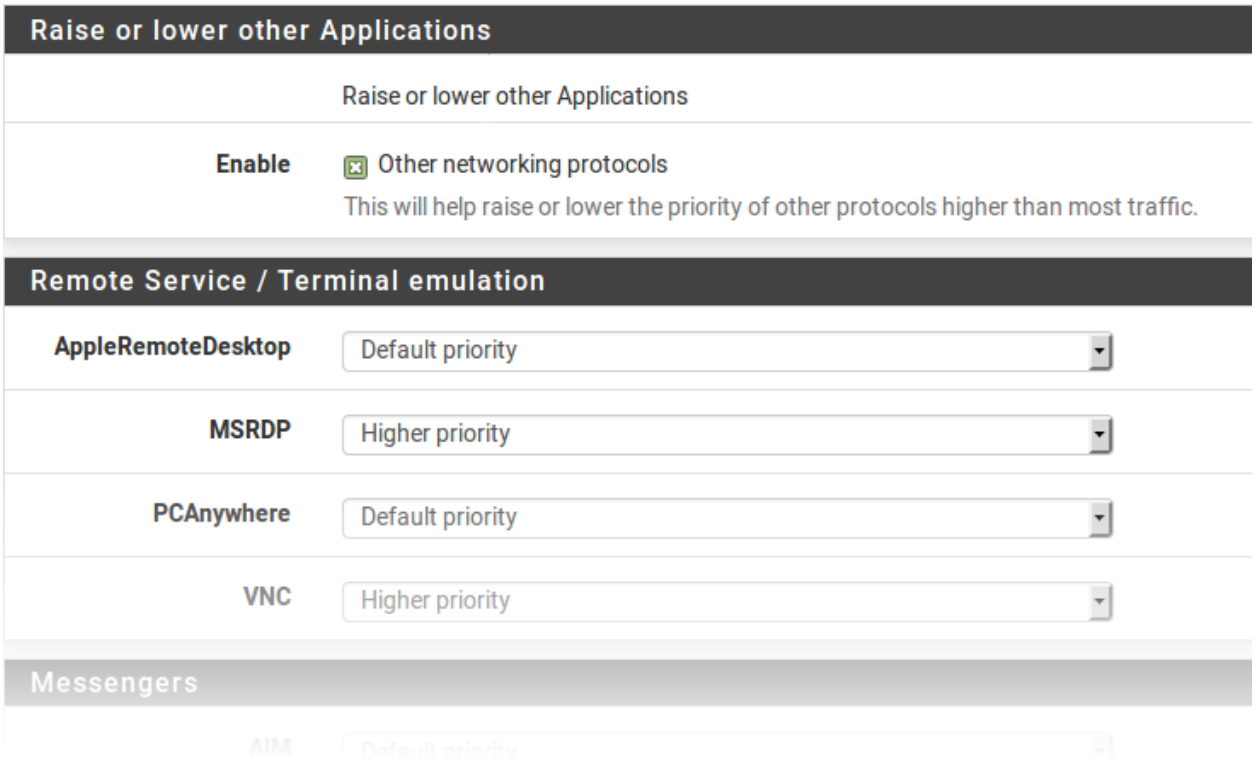


Fig. 121: Raise or Lower Other Applications

35.66.9 Finishing the Wizard

Click **Finish** to complete the wizard. The firewall will then create all the rules and queues for enabled options, and then it will reload the ruleset to activate the new traffic shaper settings.

Due to the firewall operating in a stateful manner, the firewall can only apply changes in traffic shaping to new connections. In order for the new traffic shaping settings to be fully active on all connections, clear the states.

To reset the state table contents:

- Navigate to **Diagnostics > States**
- Click the **Reset States** tab
- Check **Reset the firewall state table**
- Click **Reset**

35.66.10 Shaper Wizard and IPv6

The shaper wizard creates rules for IPv4 traffic only. Rules can be manually adjusted or cloned and set for IPv6.

35.67 Configuring CoDel Limiters for Bufferbloat

The FQ_CODEL limiter scheduler can help alleviate the effects of Bufferbloat. The CoDel algorithm and bufferbloat are discussed in the ALTQ chapter at *CoDel Active Queue Management* and the same concepts apply to FQ_CODEL with limiters as well.

Before starting, use a Bufferbloat Test Site to determine if changes are necessary. If the firewall already receives a high score the circuit may not be prone to bufferbloat and thus may not require these limiters.

This configuration requires a limiter and queue for both download and upload, plus a floating rule to apply the limiters to outgoing traffic.

35.67.1 Create Download Limiter and Queue

The first task is to create a download limiter and queue:

- Navigate to **Firewall > Traffic Shaper**, **Limiters** tab
- Click  **New Limiter**
  - Configure the limiter with the following settings:

    Enable
    Checked

    Name
    `WANDown`

    Bandwidth
    Set equal to WAN download bandwidth (confirm via speed test first)

    Mask
    None

    Description
    `WAN Download`

    Queue Management Algorithm
    Tail Drop

    Scheduler
    FQ_CODEL

    The page will display FQ_CODEL options and their default values after saving this limiter, but leave them at defaults.

    Queue Length
    Can vary depending on the speed of the link, but `1000` should be a safe default for most high speed WANs (100Mbit/s). For very high speed WANs (e.g. 1Gbit/s+), consider increasing further to `3000-5000`.

    ECN
    Checked

  - Click **Save**

- Click  **Add New Queue** under WANDown
  - Configure the queue with the following settings:

    Enable
    Checked

    Name
    `WANDownQ`

    Mask
    None

    Description
    `WAN Download Queue`

    Queue Management Algorithm
    Tail Drop
  - Leave the other fields at their default values
  - Click **Save**

35.67.2 Create Upload Limiter and Queue

- Navigate to **Firewall > Traffic Shaper**, **Limiters** tab

- Click  **New Limiter**
  - Configure the limiter with the following settings:

    Enable
    Checked

    Name
    `WANUp`

    Bandwidth
    Set equal to WAN upload bandwidth (confirm via speed test first)

    Mask
    None

    Description
    `WAN Upload`

    Queue Management Algorithm
    Tail Drop

    Scheduler
    FQ_CODEL

    The page will display FQ_CODEL options and their default values after saving this limiter, but leave them at defaults.

    Queue Length
    Can vary depending on the speed of the link, but `1000` should be a safe default for most high speed WANs (100Mbit/s). For very high speed WANs (e.g. 1Gbit/s+), consider increasing further to `3000-5000`.

ECN
Checked

– Click **Save**

- Click  **Add New Queue** under WANUp
  – Configure the queue with the following settings:

Enable
Checked

Name
`WANUpQ`

Mask
None

Description
`WAN Upload Queue`

Queue Management Algorithm
Tail Drop

  – Leave the other fields at their default values
  – Click **Save**
- Click **Apply Changes**

35.67.3 Create Floating Rule

- Navigate to **Firewall > Rules**, **Floating** tab
- Click  **Add** to create a new rule at the bottom of the list to match IPv4 traffic from clients
  – Configure the rule as follows:

Action
Pass

Quick
Checked

Interface
WAN

Direction
Out

Address Family
IPv4

Protocol
Any

Source
WAN Address

> **Warning:** It is important not to match too loosely on the source, especially when a firewall has multiple WANs. In certain cases with multiple WANs, if traffic meant to exit an alternate non-default WAN matches this kind of floating rule, the traffic will end up dropped as PF may still process that traffic outbound on the default WAN before redirecting out the correct interface.

Destination
Any

Description
`CoDel Limiters`

Gateway
Must be set to the gateway for this WAN interface

In / Out Pipe
WANUpQ / *WANDownQ*

> **Note:** On WAN floating rules in the outbound direction, "in" traffic is upload, and "out" traffic is download, from the perspective of LAN clients.

- Save

- Create additional rules for IPv6 client prefixes if this WAN also carries IPv6 traffic

  When creating IPv6 rules, the parameters are mostly the same as the previous rule example except:

  Address Family
  IPv6

  Source
  LAN subnets

  > **Note:** In most cases this is the IPv6 prefix used by local clients, so choosing an option such as *LAN subnets* is ideal.
  >
  > If the firewall is performing NPt or other outbound NAT on IPv6 traffic from local clients, then set this to the post-NAT address or prefix instead.

- Apply Changes
- Reset states to force all traffic to use new limiters

35.67.4 Test Again

Use a Bufferbloat Test Site again and compare score now to the score before the test was run. In most cases, the new score should be an A or higher.

If the score does not improve, or gets worse, there is likely a problem with the configuration. First, go back and compare all the settings with the suggested values on this document.

If the configuration matches, the settings may need further adjustment. For example, the bandwidth values may be higher than the circuit is capable of delivering, the queue sizes may need increased, or the CoDel parameters may need changed. Post on the Netgate Forum for assistance with diagnosing the problem.

35.67.5 Notes

Certain configurations may require alterations to the suggested procedure above.

Multiple WANs

For multiple WANs make a complete set of queues for **each** WAN and make a **separate** floating rule for each WAN. Ensure the rules do not match the source IP address(es) of the other WANs.

For example:

- Pass quick out WAN1 from WAN1 Address to any, gateway WAN1GW, In/Out Pipe WAN1UpQ/WAN1DownQ
- Pass quick out WAN2 from WAN2 Address to any, gateway WAN2GW, In/Out Pipe WAN2UpQ/WAN2DownQ

Multiple Addresses/VIPs

If there are multiple IP addresses on a WAN (e.g. VIPs, routed subnets), create an alias with all the necessary addresses and use it as the source of the floating rule.

35.68 Copy Files to a USB Drive

On occasion, it may be necessary to copy files to or from the firewall using a USB flash drive.

Note: This procedure assumes the drive is formatted with FAT or FAT32 (Also known as "DOS") partitions.

At this time it is not possible to use drives formatted as exFAT or NTFS.

This procedure assumes the user is *connected to the firewall console* or *using SSH* and *connected using an SSH client*.

35.68.1 Locate Drive and Partition Name

Mounting the drive requires knowing the device name of the USB flash drive and the name of the FAT partition.

There are a couple different methods to determine these values.

Using Device Labels

The most convenient way to mount a drive is by its device label, if it is known. This is the name given by the user when formatting the drive or by altering the drive properties.

With the drive connected, look at the list of available device labels for DOS partitions:

```
: ls -l /dev/msdosfs/
crw-r-----  1 root  operator  0x93 Jul  8 13:56 MYDRIVE
crw-r-----  1 root  operator  0x65 Jul  8 11:30 EFISYS
```

In this example, the drive is named `MYDRIVE` and when mounting the full name of `/dev/msdosfs/MYDRIVE` is valid for use by `mount`.

Warning: On UEFI systems the `EFISYS` label is the system EFI boot partition, do not mount or alter the content of that partition!

Using `gpart`

The most definitive way to locate the correct drive and partition is to use `gpart` and look for a `fat32` entry.

This example system has two multiple disks, but only one of them is a USB thumb drive with a FAT32 partition:

```
: gpart list | egrep 'Name:| type:'
1. Name: mmcsd0p1
   type: efi
2. Name: mmcsd0p2
   type: freebsd-boot
3. Name: mmcsd0p3
   type: freebsd-swap
4. Name: mmcsd0p4
   type: freebsd-zfs
1. Name: mmcsd0
1. Name: da0s1
   type: fat32
1. Name: da0
```

In this example the target partition is `da0s1` as it's the name corresponding to the `fat32` type entry.

This next example system has a USB drive containing multiple partitions, but only one of them is FAT32:

```
: gpart list | egrep 'Name:| type:'
1. Name: da0s1
   type: efi
2. Name: da0s2
   type: fat32
3. Name: da0s3
   type: freebsd
1. Name: da0
1. Name: da0s3a
   type: freebsd-zfs
1. Name: da0s3
```

Based on the above output the target partition is `da0s2`.

System Log and Device List

This method is not as accurate but may be good enough for the majority of use cases.

Monitor the console or watch the system log when inserting the USB drive (e.g. `tail -F /var/log/system.log`). This will contain output similar to the following:

```
ugen0.2: <USB Flash Disk> at usbus0
umass0 on uhub0
umass0: <USB Flash Disk, class 0/0, rev 2.00/11.00, addr 1> on usbus0
da0 at umass-sim0 bus 0 scbus0 target 0 lun 0
```

(continues on next page)

(continued from previous page)

```
da0: <USB Flash Disk 1100> Removable Direct Access SPC-2 SCSI device
da0: Serial Number FBG1204030507369
da0: 40.000MB/s transfers
da0: 1912MB (3915776 512 byte sectors)
da0: quirks=0x2<NO_6_BYTE>
```

This output indicates the correct drive is `da0`, but it does not help determine the correct partition on that drive.

Next, look at the list of devices in `/dev/` to see which partitions are present on the drive:

```
: ls -l /dev/da0*
crw-r-----  1 root  operator  0x91 Jul  8 13:56 /dev/da0
crw-r-----  1 root  operator  0x92 Jul  8 13:56 /dev/da0s1
```

For a USB drive containing only a single FAT32 partition, `da0s1` is likely the correct partition.

35.68.2 Mount the Partition

Before mounting, create a directory to use as the mount point. The directory `/mnt` can be used for this purpose, but a safer practice is to create a custom directory:

```
: mkdir -p /root/usb
```

The next step is to mount the drive using the full path to the label or partition and the mount point:

Label example:

```
: mount -t msdosfs /dev/msdosfs/MYDRIVE /root/usb
```

Partition device example:

```
: mount -t msdosfs /dev/da0s1 /root/usb
```

Warning: Remember to *unmount the drive* before removing it from USB!

35.68.3 Copy the Files

With the drive mounted, files can be copied to or from the drive using the mount point directory, `/root/usb` in this example.

Copy/Move files from the firewall to the USB drive:

```
: cp /conf/config.xml /root/usb/config-backup.xml
: mv /tmp/status_output.tgz /root/usb/
```

Copy files from the USB drive to the firewall:

```
: cp /root/usb/myscript.sh /root/bin/
```

35.68.4 Unmount and Clean Up

After copying files, the drive **must** be unmounted:

```
: umount /root/usb/
```

With the drive unmounted it is now safe to remove the USB device from the firewall.

Warning: Failing to unmount the drive before removing the USB device can result in a kernel panic and reboot!

Next, remove the mount point directory if it is no longer necessary:

```
: rmdir /root/usb/
```

Note: This is optional. The mount point directory may be left in place for future use.

Warning: Do not use `rm -rf` or similar on the mount point! If the device was still mounted, this would destroy files on the device. Using `rmdir` ensures the operation will only have an effect if the directory is empty.

35.68.5 Full Example

```
# Insert the USB drive

# Find the label
: ls -l /dev/msdosfs/
crw-r-----  1 root  operator  0x93 Jul  8 13:56 MYDRIVE

# Create the mountpoint
: mkdir /root/usb

# Mount the drive
: mount -t msdosfs /dev/msdosfs/MYDRIVE /root/usb

# Copy files
: cp /conf/config.xml /root/usb/config-backup.xml

# Unmount the drive
: umount /root/usb/

# Remove the mountpoint
: rmdir /root/usb/

# Remove the USB drive
```

35.69 Virtualizing pfSense Software with VMware vSphere / ESXi

This article is about building a pfSense® virtual machine on vSphere / ESXi. Article explains how to install any major pfSense software version on VMware vSphere.

> **Warning:** Depending on the version of pfSense software in use, there may be specific vSphere / ESX version requirements as well. For example, versions of pfSense software based on FreeBSD 12.x require ESX 6.7 or later while versions based on FreeBSD 14.x require ESX 7.0 or later. Typically, the VM hardware version must also be at that level or greater. While older versions may work, they may be unstable.
>
> Refer to both *Versions of pfSense software and FreeBSD* and the VMware Guest OS Compatibility Guide to find the base OS and requirements for each version.

The article does not cover how to install vSphere or how to configure pfSense software to do any of the many amazing things it can. A basic, working, virtual machine running pfSense software will exist by the end of this document.

Note: If the pfSense software instance will be running as a perimeter firewall for an organization and the "attack surface" should be minimized, many will say it is preferable to run it on stand-alone hardware without virtualization. That is a decision for the user and/or organization to make, however.

This guide starts at the point with a vanilla ESXi install connected using the vSphere client. If other VMs are already running on ESXi, then it is not likely necessary to follow the networking steps too closely. However, skim through it to see what is suggested before building the pfSense software virtual machine.

35.69.1 Assumptions

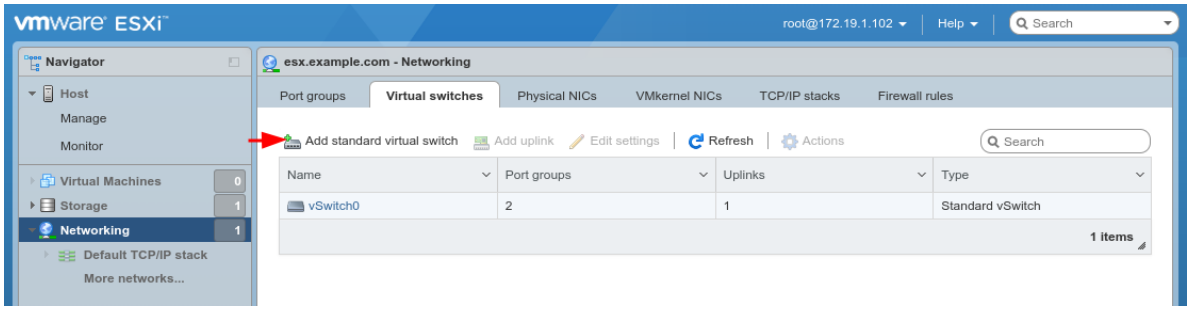
- vSphere host is running, and the user can log into the web interface through its management network adapter.
- The reader has an understanding of network addressing.
- vSphere host has a working datastore.
- The pfSense software installation `.iso` image is present in a datastore.

The following steps include the necessary vSphere web client configuration required to get a VM for pfSense software running.

35.69.2 Basic vSphere web client networking setup

Before creating a new VM in vSphere web client, create two virtual switches and two port groups. First, create Virtual switches for WAN and LAN and after that two port groups for the WAN and LAN. If there are existing virtual switches in the environment which can be used for this VM, skip this step.

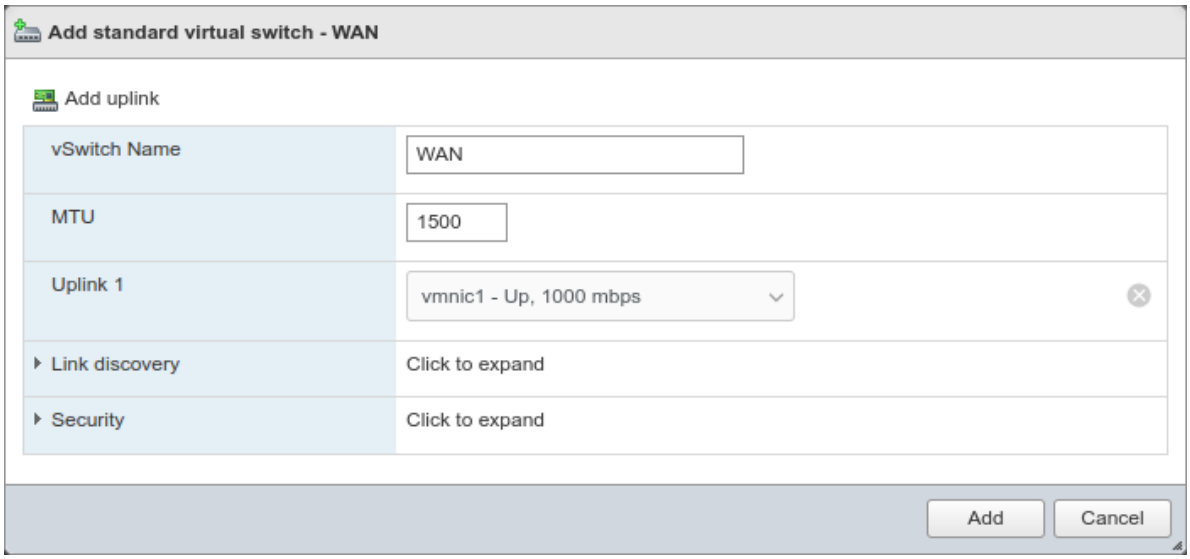
- Open the vSphere web interface
- Click **Networking**, **Virtual switches** tab
- Click **Add standard virtual switch**



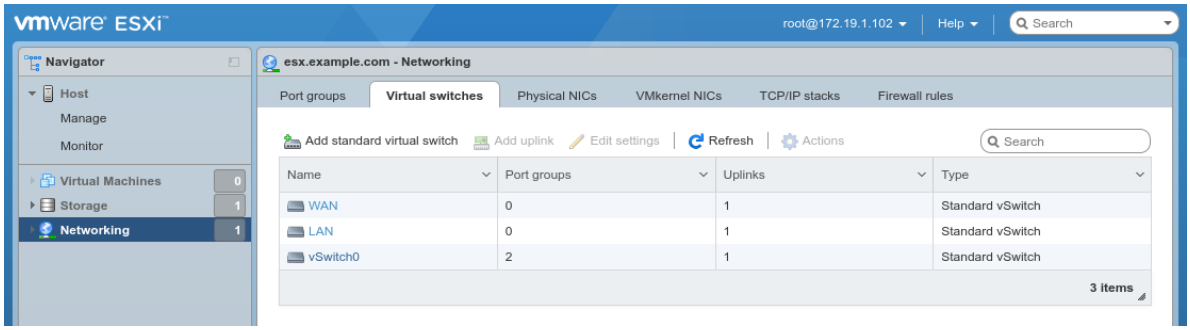
- Configure the vSwitch as follows:

  vSwitch Name
  `WAN`

  Uplink 1
  vmnic1



- Click **Add**
- Repeat the process and add another vSwitch named `LAN` for *vmnic2*



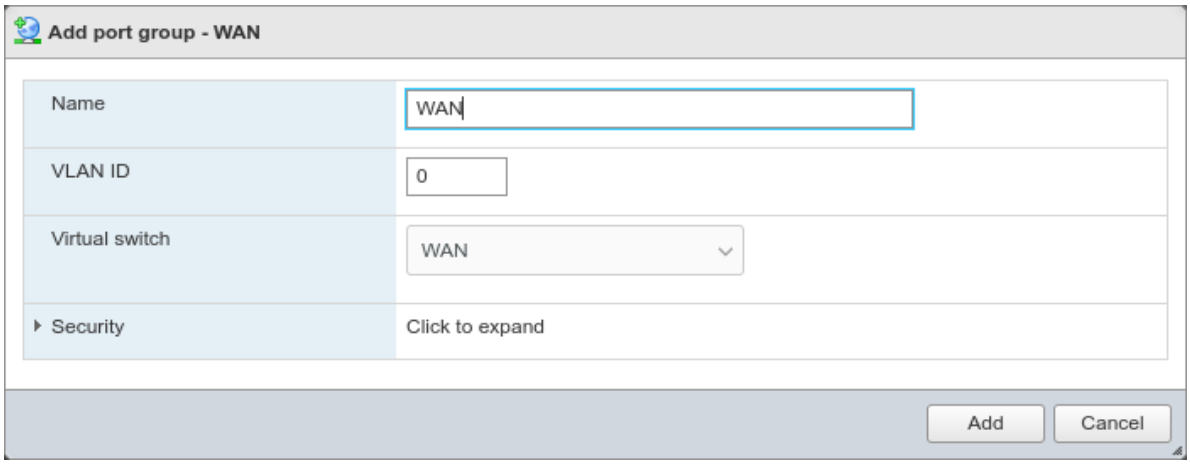
35.69.3 Creating port groups

After creating Virtual switches, now create port groups. If there are existing port groups in the environment which can be used for this VM, skip this step.

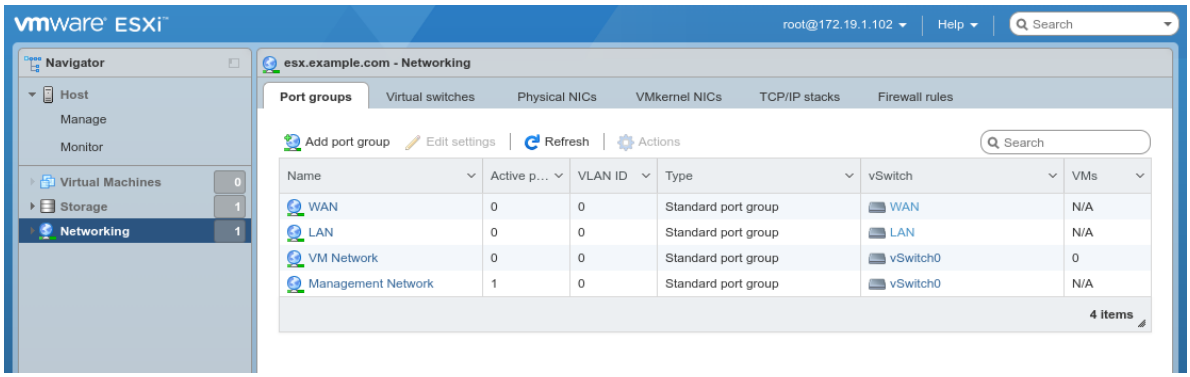
- Click **Networking**, **Port groups** tab
- Click **Add port group**
- Configure the port group as follows:

  Name
  `WAN`

  Virtual switch
  WAN



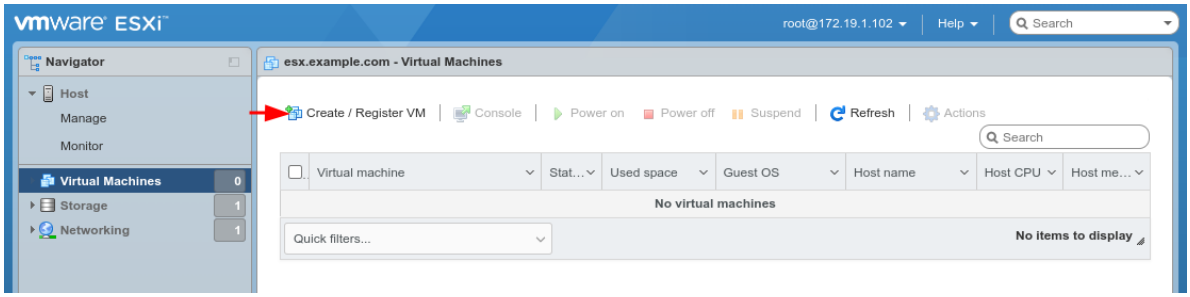
- Click **Add**
- Repeat the process and add another port group named `LAN` for the *LAN* vSwitch.



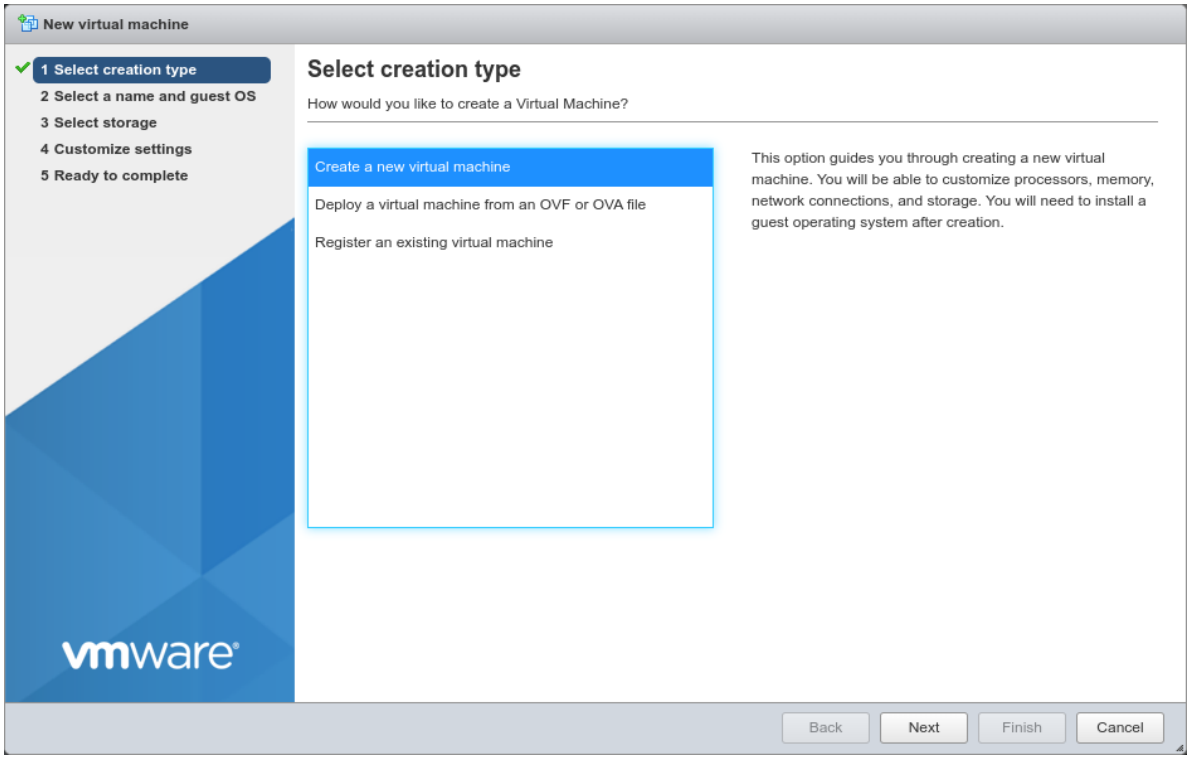
35.69.4 Creating a Virtual Machine

With the required networking configured, the next step is to create a virtual machine.

- Click **Virtual Machines** on the left Navigator pane
- Click **Create/Register VM**



- Select **Create a new virtual machine** on the first wizard



- Click **Next**
- Configure the **Select a name and guest OS** screen of the wizard as follows:

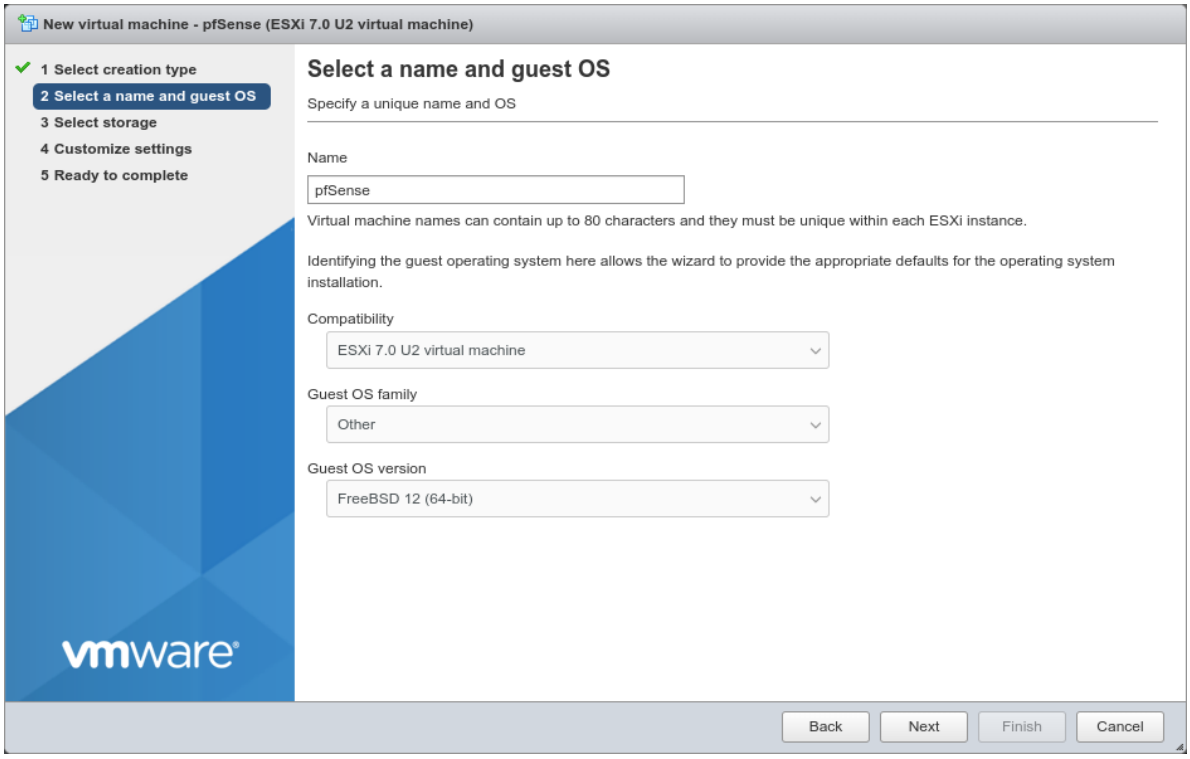
  Name
  `pfSense` or another meaningful name, such as `firewall`.

  Compatibility
  The latest version available (e.g. *ESXi 7.0 U2 virtual machine*)

  Guest OS Family
  Other

  Guest OS Version

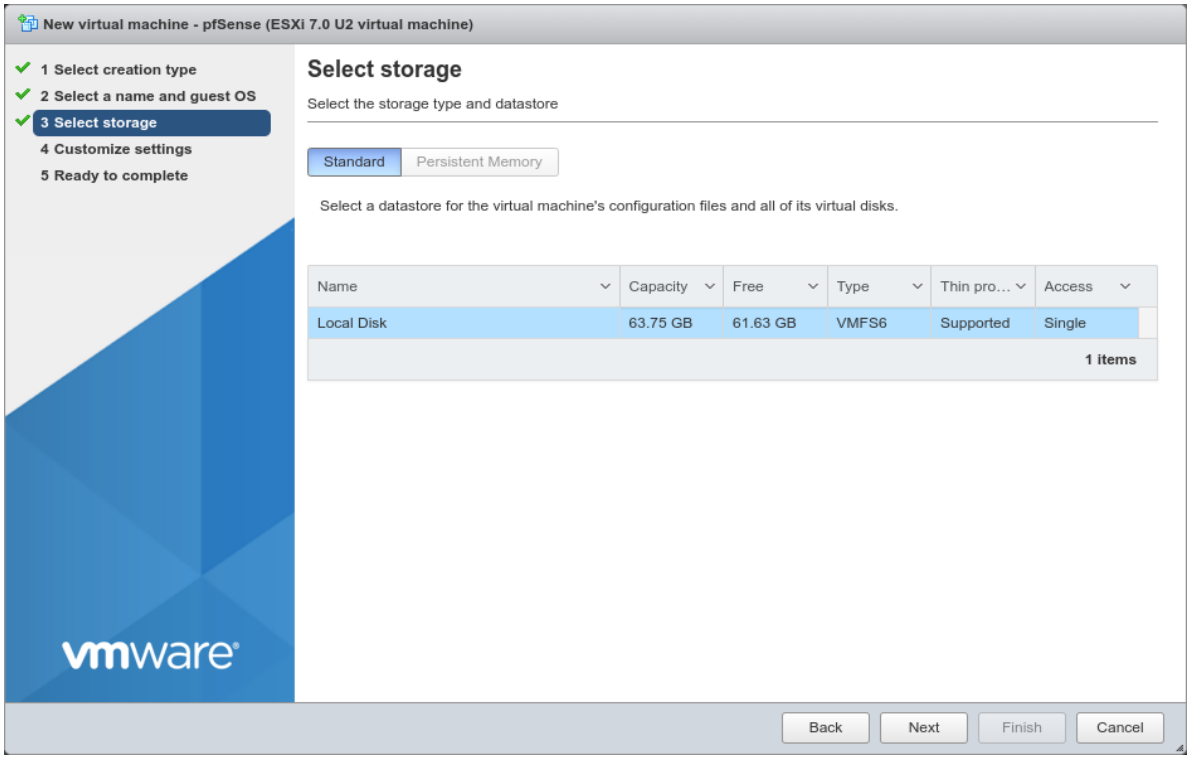
FreeBSD 12 (64-bit) or whichever version best matches the version of FreeBSD used by the chosen version of pfSense software. See *Versions of pfSense software and FreeBSD* for a list.



- Click **Next**
- Select the datastore where the VM disk will live

  This is where ESX will allocate storage to hold the configuration and operating files for the virtual machine. There may be multiple datastores available to ESX, local or even remote NFS volumes.

  Local disks are faster and more reliable, but pfSense software does not require a fast disk in most use cases, so in those environments it can run off an NFS disk if necessary.



- Click **Next**
- Click **Add network adapter** to create a second NIC
- Configure the items on the **Customize settings** screen as follows:

  CPU
  Use a single CPU socket.

  If the hypervisor host has sufficient cores available, click to expand the CPU options and set a higher value for **Cores per socket**.

  Memory
  Depending on the number and type of packages that will be installed on the pfSense software, a basic firewall VM should run comfortably in 1024MB of RAM. For deployments which require more or larger packages, increase the RAM as needed.

  Hard Disk 1
  Give the VM at least 16 GB of space, more for larger packages.

  SCSI Controller 0
  The default *LSI Logic SAS* is compatible, leave it as-is.

  Network Adapter 1
  Select the *WAN* port group.

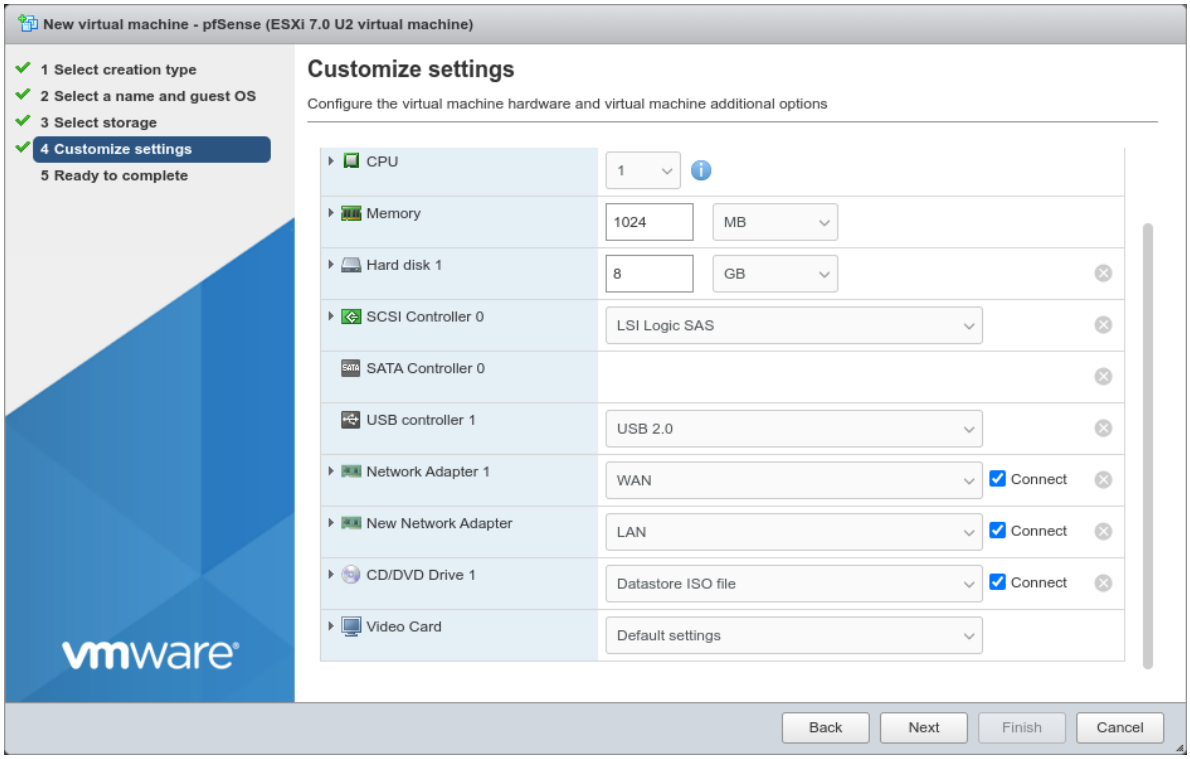
  For best performance, use VMXNET 3 type of adapters which is the current default in vSphere 7.x. Click to expand the interface options and ensure it's set to VMXNET 3.

  New Network Adapter
  Select the *LAN* port group.

  Click to expand the interface options and ensure it's set to VMXNET 3.

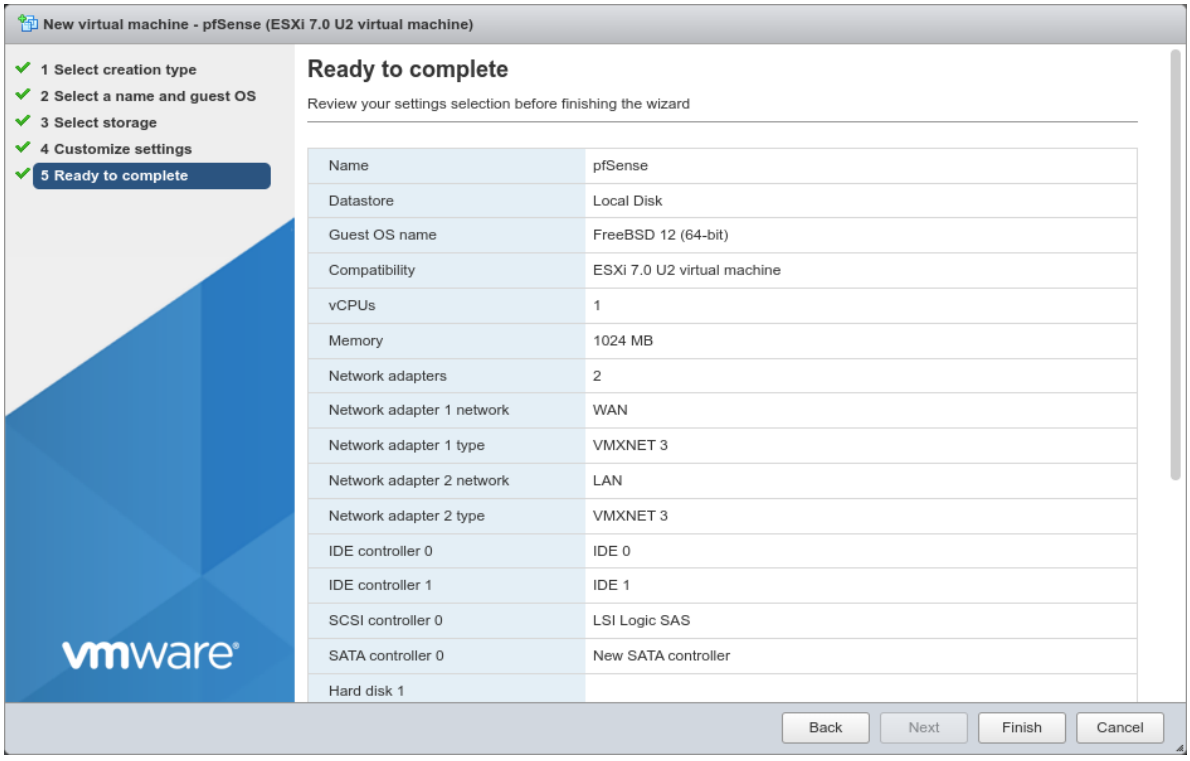
  CD/DVD Drive 1
  Select *Datastore ISO file** and then browse to and pick the pfSense software installer ISO.

The remaining options can remain at their default values, or change to suit the needs of the environment.



- Click **Next**
- Review the settings for the VM

  If anything is incorrect, go back to the previous screens and correct it.

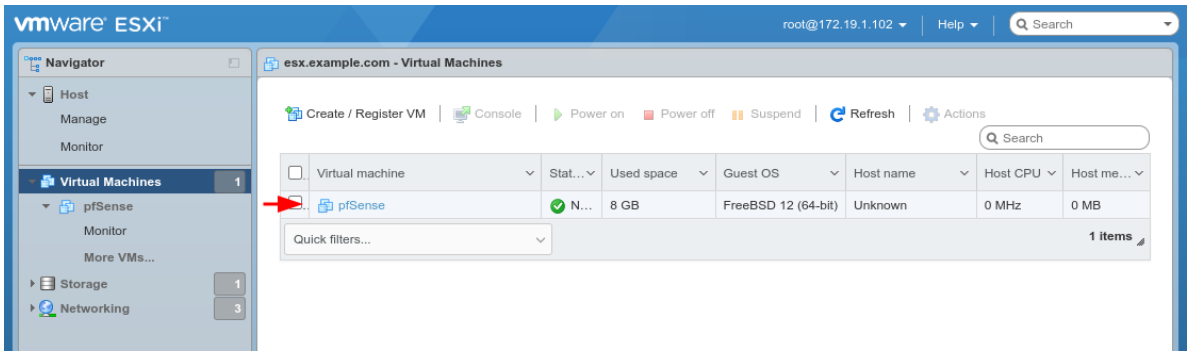


- Click **Finish**

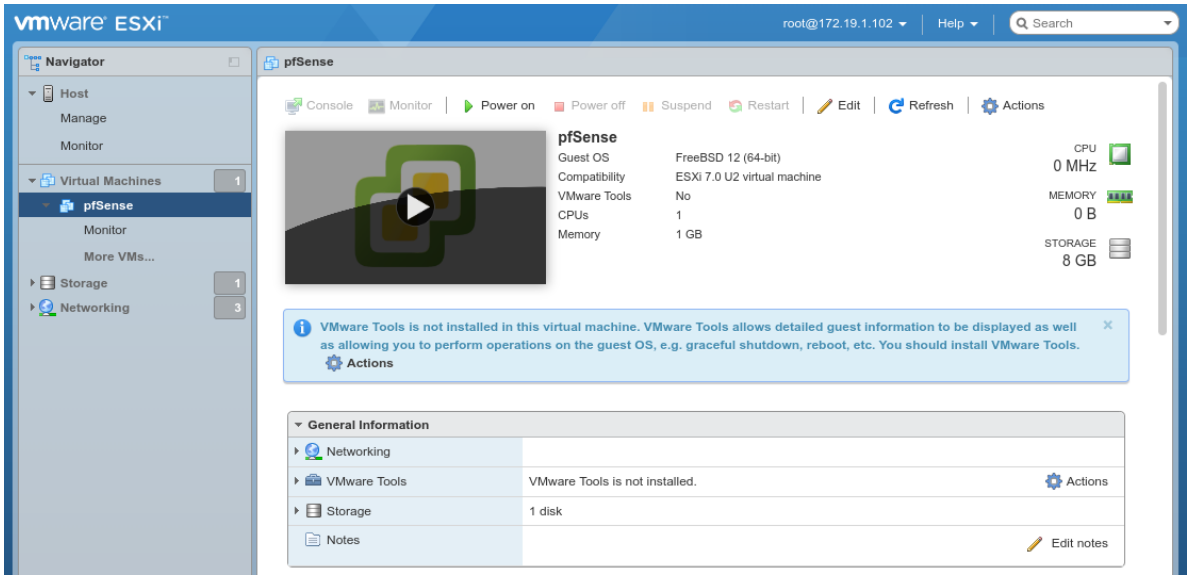
35.69.5 pfSense software installation

The vSphere web interface will now have an entry for the new VM.

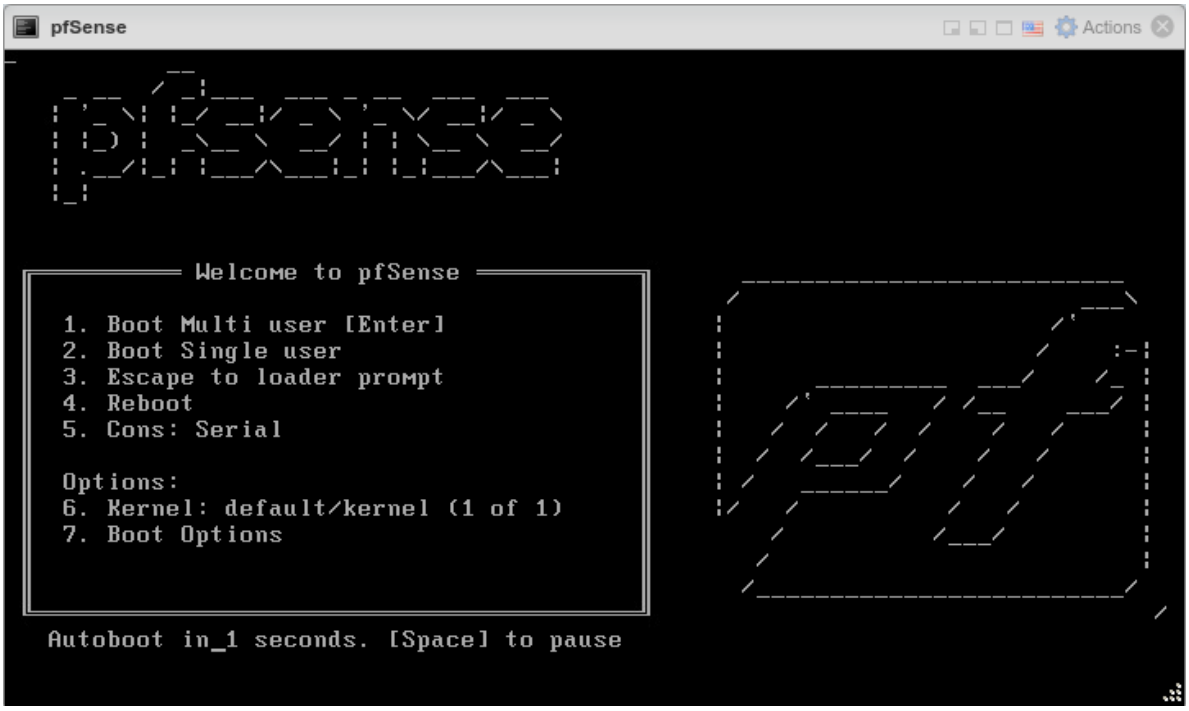
- Click **Virtual Machines** in the **Navigator** panel on the left



- Click the name of the VM in the list to open it



- Click **Power on** to start the VM
- Click inside the console window to open the console view to continue the installation.



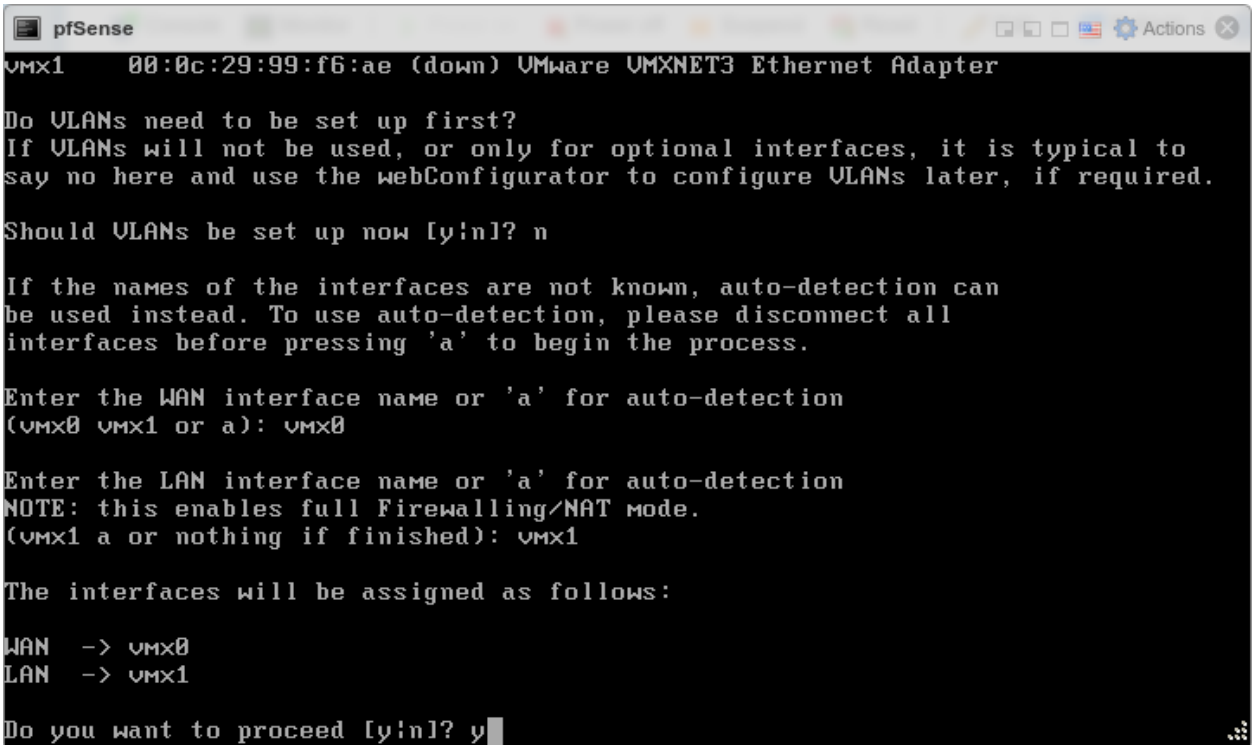
When the VM starts it will boot into the installer automatically. From there, follow the installation steps as usual, and reboot when finished.

See also:

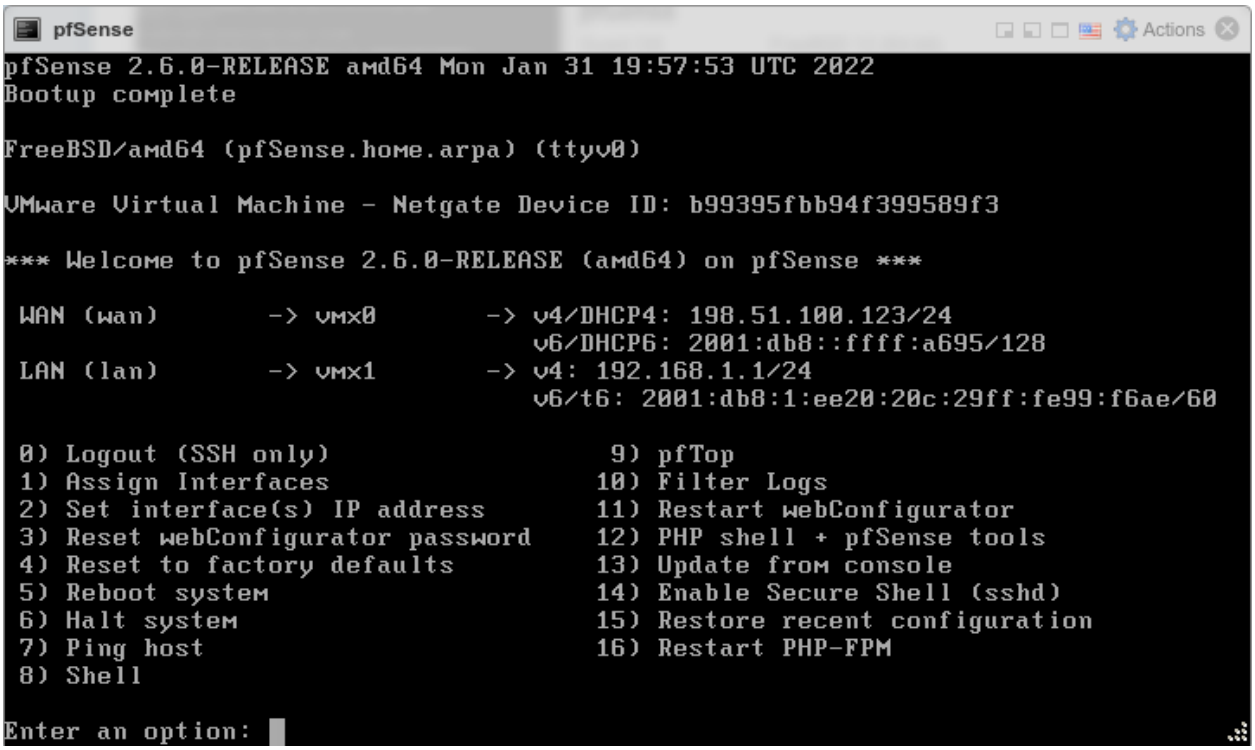
See *Installation Walkthrough* for a detailed walkthrough of the installation process.

After the virtual machine boots back up, the console will stop at a prompt to assign interfaces.

- Type `n` and press `Enter` to skip VLAN configuration
- Enter `vmx0` for WAN
- Enter `vmx1` for LAN
- Press `Enter` if prompted for additional interfaces
- Type `y` and press `Enter` to complete the interface assignment



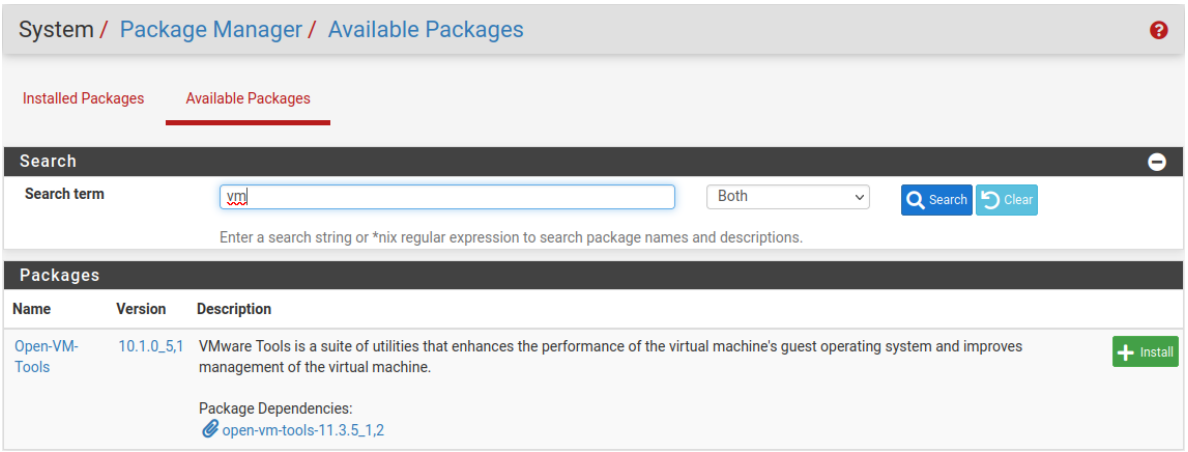
After assigning the interfaces the VM will complete the boot process. It is now ready to configure like any other firewall running pfSense software.



Installing Open-VM-Tools

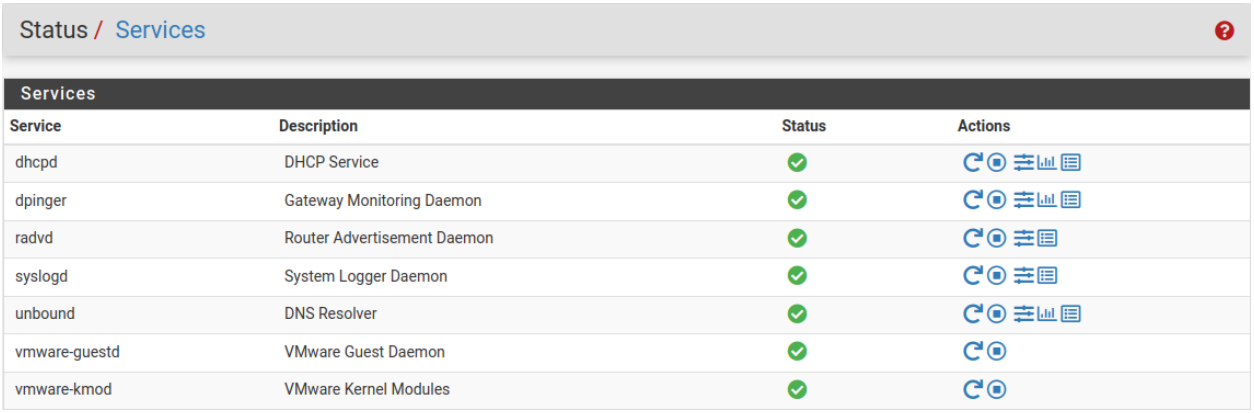
Once the pfSense software installation is complete, upon first boot install the Open-VM-Tools.

- Navigate to **System > Packages**, **Available Packages** tab
- Find **Open-VM-Tools** in the list or search for it



- Click  **Install**
- Confirm the installation

Make sure the Open-VM-Tools service is running under **Status > Services**.



Congratulations, the installation of pfSense software on ESXi is complete!

35.69.6 Additional Information and Tips

Dedicated Management Network

The best practice is to separate the ESXi Management network from other networks. The example in this recipe uses a dedicated management network, which is common in well-designed networks. Separation can be accomplished using VLANs or an additional NIC on the ESXi host dedicated only for ESXi management. The vSphere client PC may need additional routing or networking connections to reach the dedicated management network.

Identifying Interfaces

If multiple physical interfaces are available in the ESXi host, it can be a bit of a struggle to work out which one has been identified as `vmnic1`, `vmnic2` and so on. If the MAC address of each NIC is noted down along with the slot it occupied when it was installed in the machine, look at the Network Adapters screen under the Configuration tab to match up the MAC addresses. However, having that foresight is rare, so lacking that information the easiest way to match physical NICs to `vmnic` entries is to plug a PC or switch into them, one at a time. The speed and duplex on the Networking or Network Adapters screens should change as the interface comes up. Click **Refresh** to update the list.

35.70 Virtualizing pfSense Software with Hyper-V

This article is about running pfSense® software in a virtual machine under Microsoft Hyper-V. The guide applies to any Hyper-V version, desktop or server (this includes the standalone Hyper-V Server). The guide explains how to install any major pfSense software version under Hyper-V. Article covers the Hyper-V networking setup and pfSense software virtual machine setup process. The guide does not cover how to install Hyper-V or Windows Server. A basic, working, pfSense software virtual machine will exist by the end of this article.

Note: If pfSense software will be used as a perimeter firewall for an organization and the attack surface should be minimized, the best practice is typically to run the firewall non-virtualized on standalone hardware. That is a decision for the user and/or organization to make, however.

This guide starts at a point with a Windows and the Hyper-V role installed. If other VMs are already running on Hyper-V, then it is not likely necessary to follow the networking steps too closely. However, skim through it to see what is suggested before building the pfSense software virtual machine part.

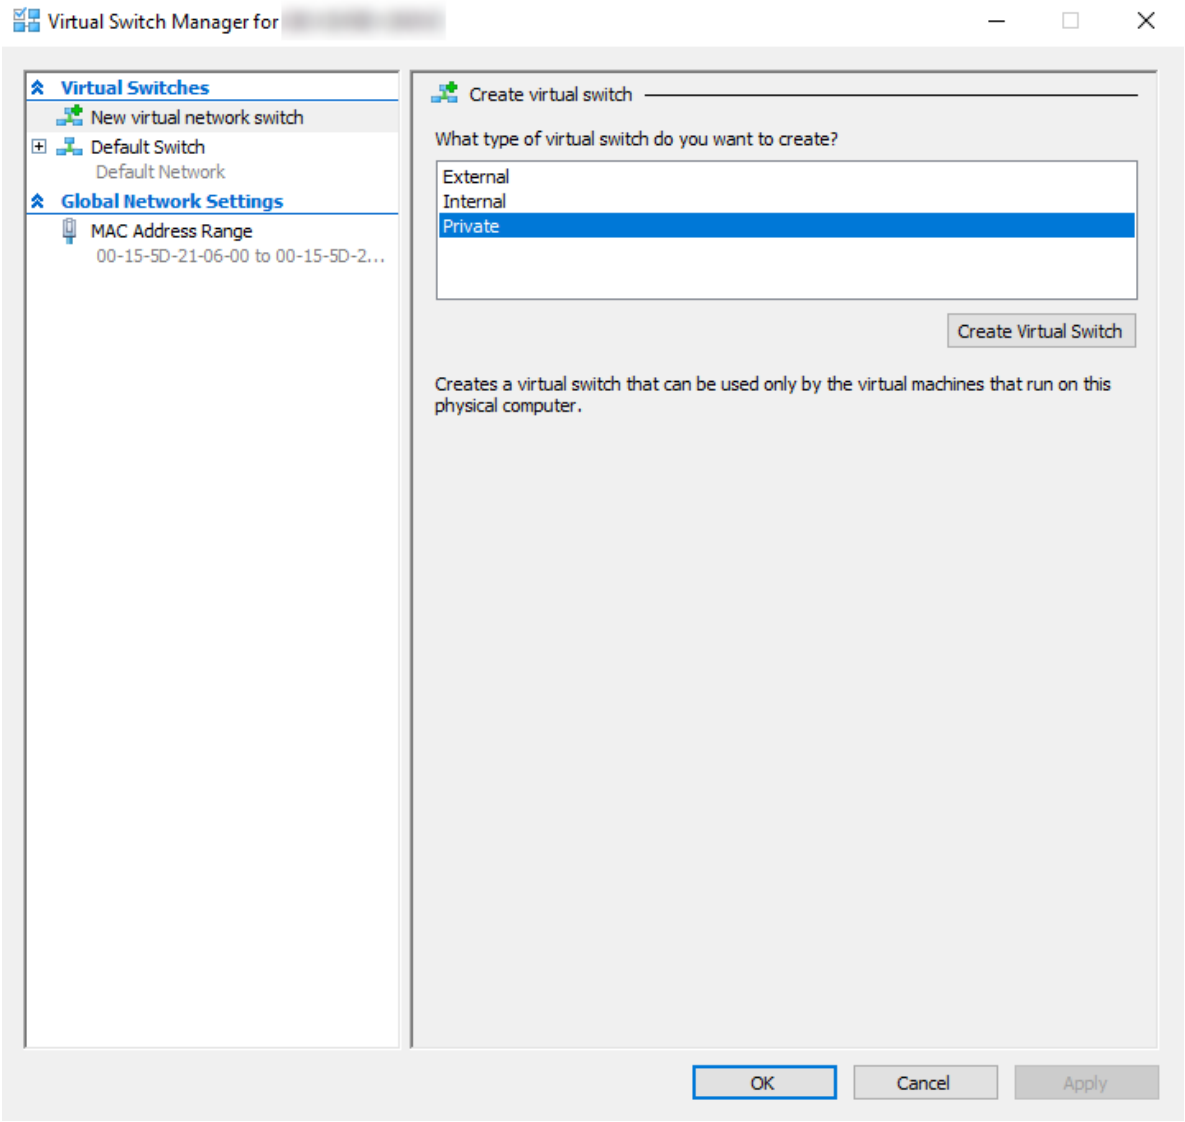
35.70.1 Assumptions

- Hyper-V host is up and Hyper-V role/feature has been installed
- The reader has a basic understanding of networking and Hyper-V virtualization

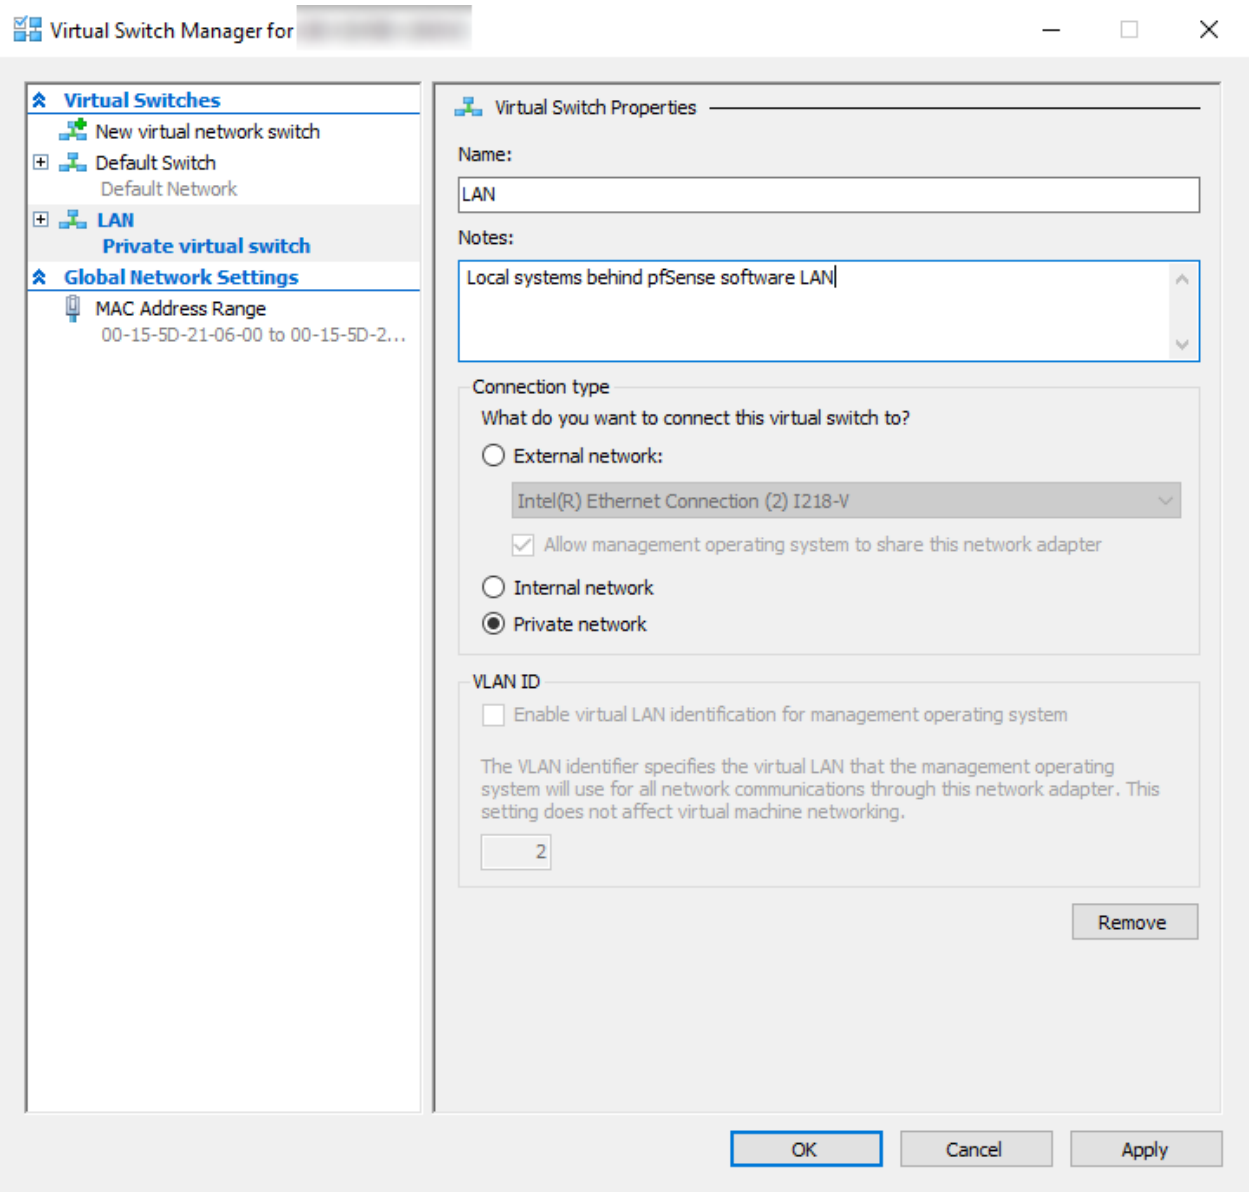
35.70.2 Basic Hyper-V Networking

To virtualize pfSense software, first create two **Virtual Switches** via Hyper-V Manager.

- Open the Hyper-V Manager
- Click **Virtual Switch Manager** from the **Actions** menu
- Select **Private** for the type of virtual switch
- Click **Create Virtual Switch**



- Set the **Name** for the newly added switch to `LAN`
- Set an appropriate description in the **Notes** field
- Ensure the **Connection type** is set to *Private network*
- Click **Apply**



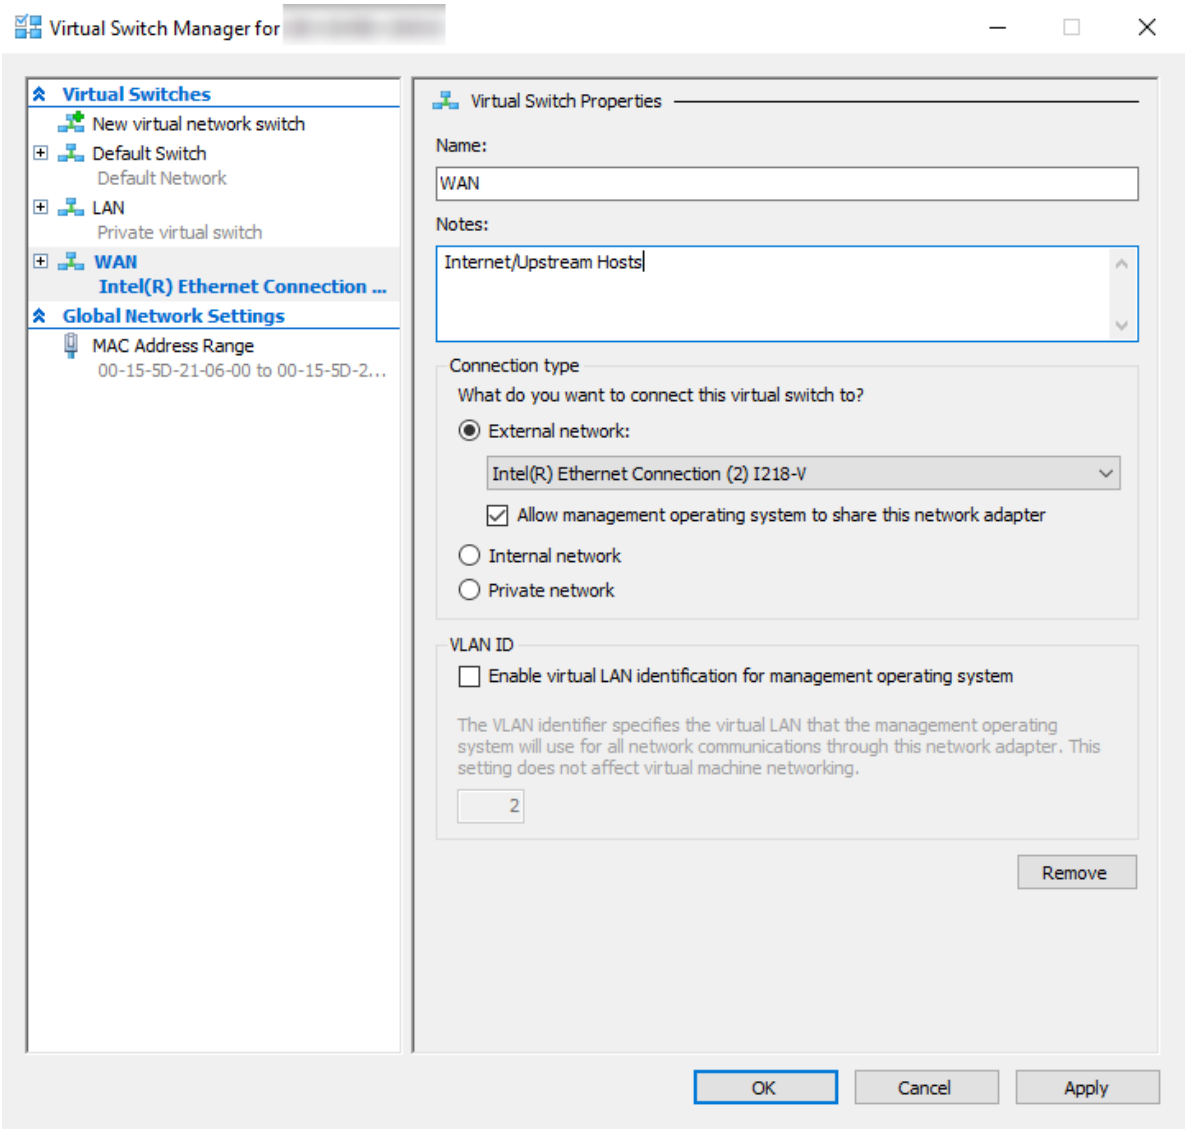
Now create a switch for the WAN/Upstream networks:

- Click **New virtual network switch**
- Select **External** for the type of virtual switch
- Click **Create Virtual Switch**
- Set the **Name** for the newly added switch to `WAN`
- Set an appropriate description in the **Notes** field
- Select the appropriate interface for the **External network**

  This is the interface on the Windows host which connects to the upstream/WAN switch/CPE or similar uplink.
- Uncheck **Allow management operating system to share this network adapter** if the hypervisor host has a dedicated interface for WAN.

  For the purpose of this guide the management was allowed, however production use requires a separate NIC for

WAN.



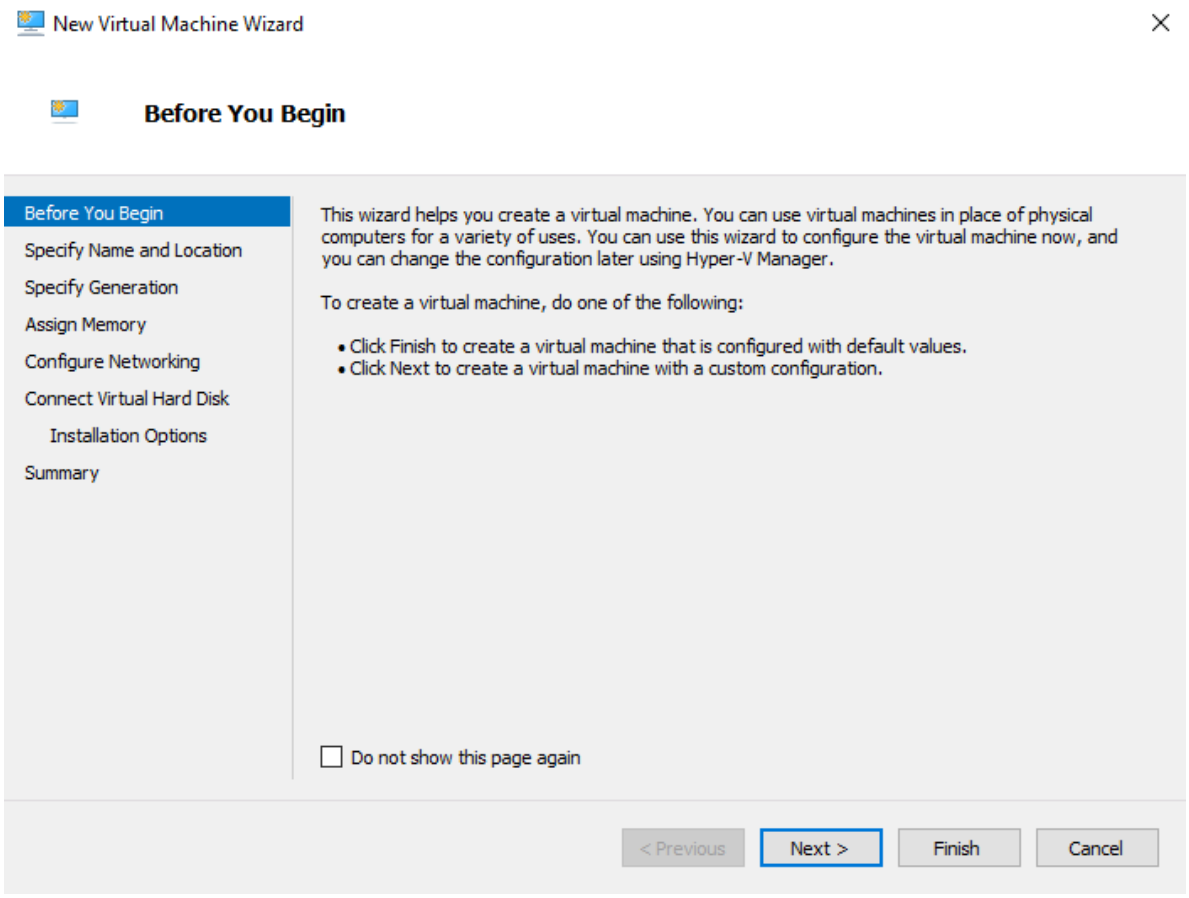
- Click **OK** to complete the switch setup

35.70.3 Creating the virtual machine

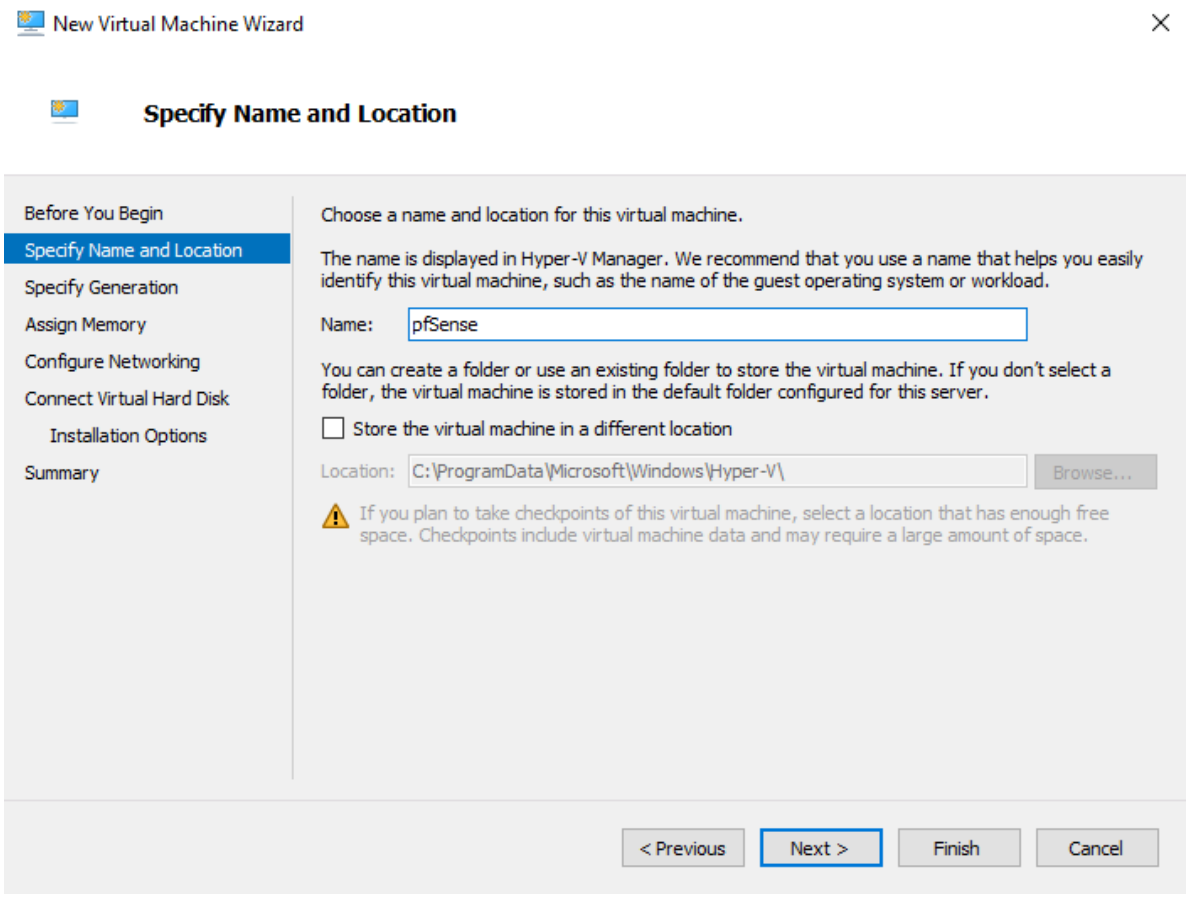
After creating WAN and LAN switches, move to virtual machine creation.

- Click **New > Virtual Machine** from the **Actions** list

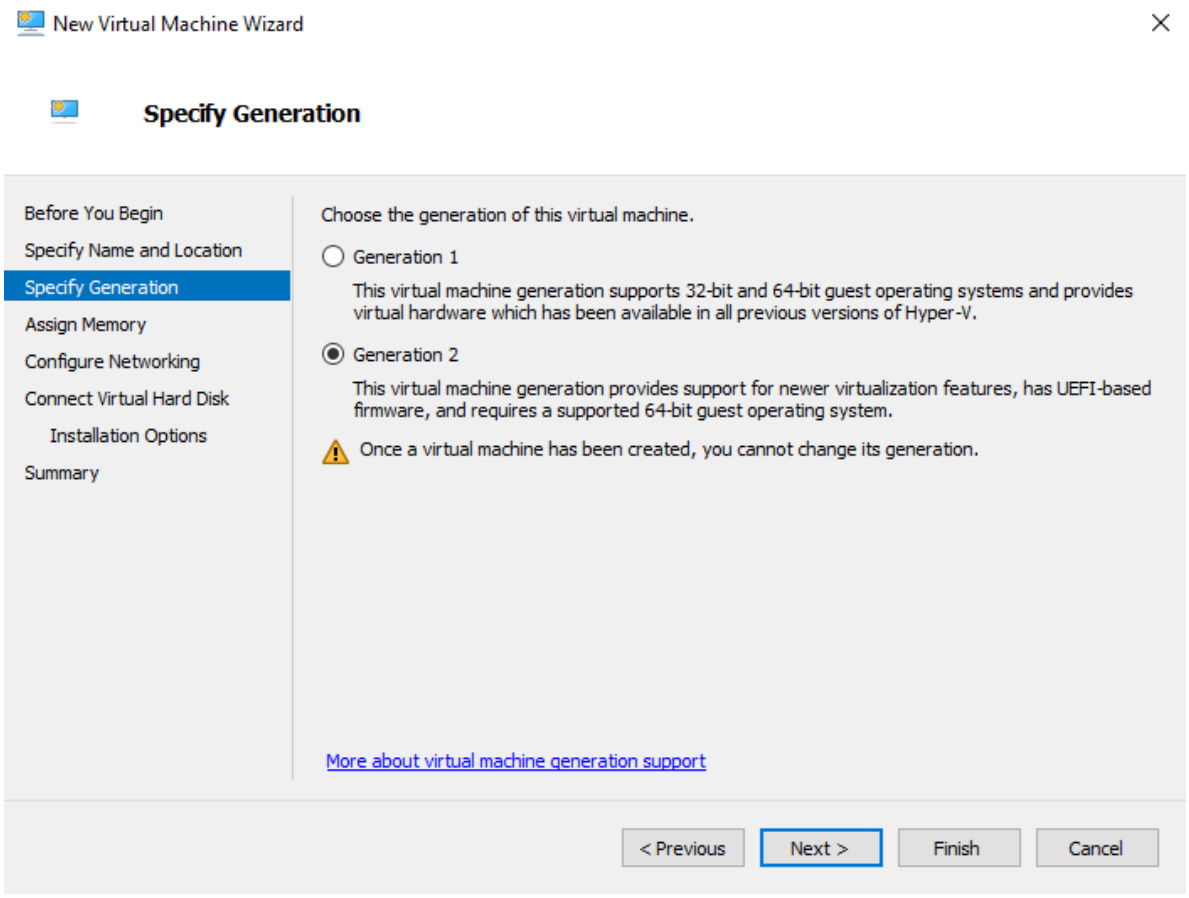
  This starts the new virtual machine wizard.



- Click **Next** and proceed to the **Specify Name and Location** step
- Enter a **Name** for the virtual machine, such as `pfSense`

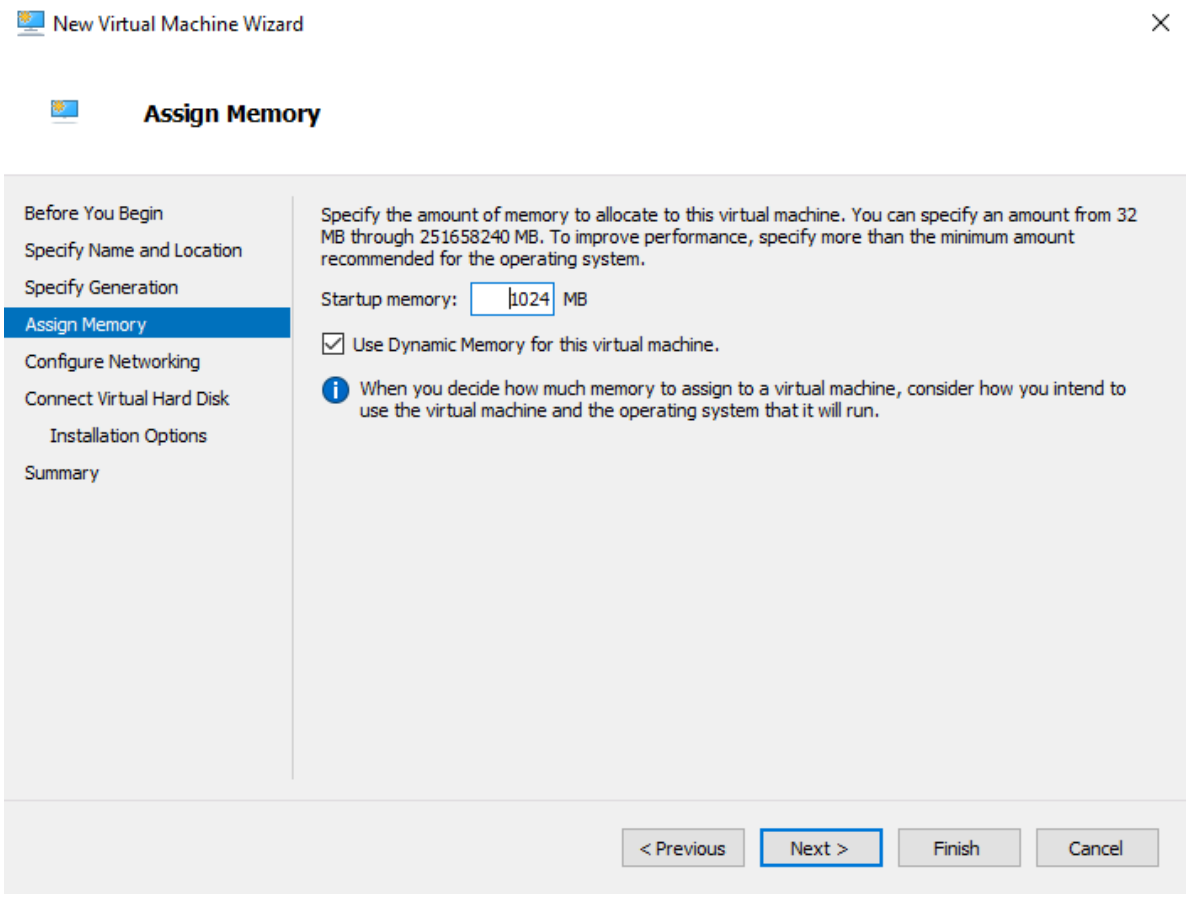


- Click **Next** and proceed to the **Specify Generation** step
- Select the appropriate virtual machine generation: *Generation 2*



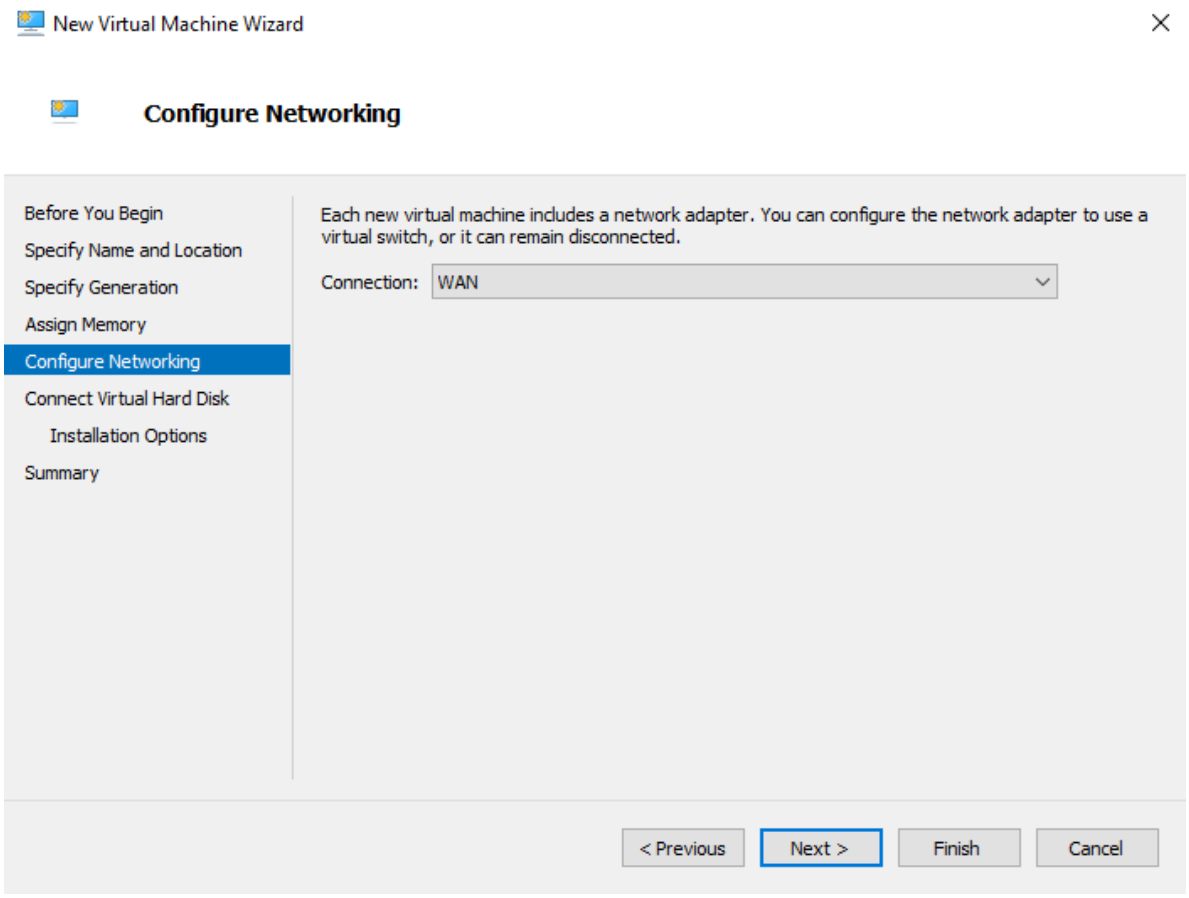
- Click **Next** and proceed to the **Assign Memory** step
- Add enough RAM to meet the requirements of this environment

  This guide uses 1GB (`1024` MB). 2GB is better if this VM will run multiple packages.



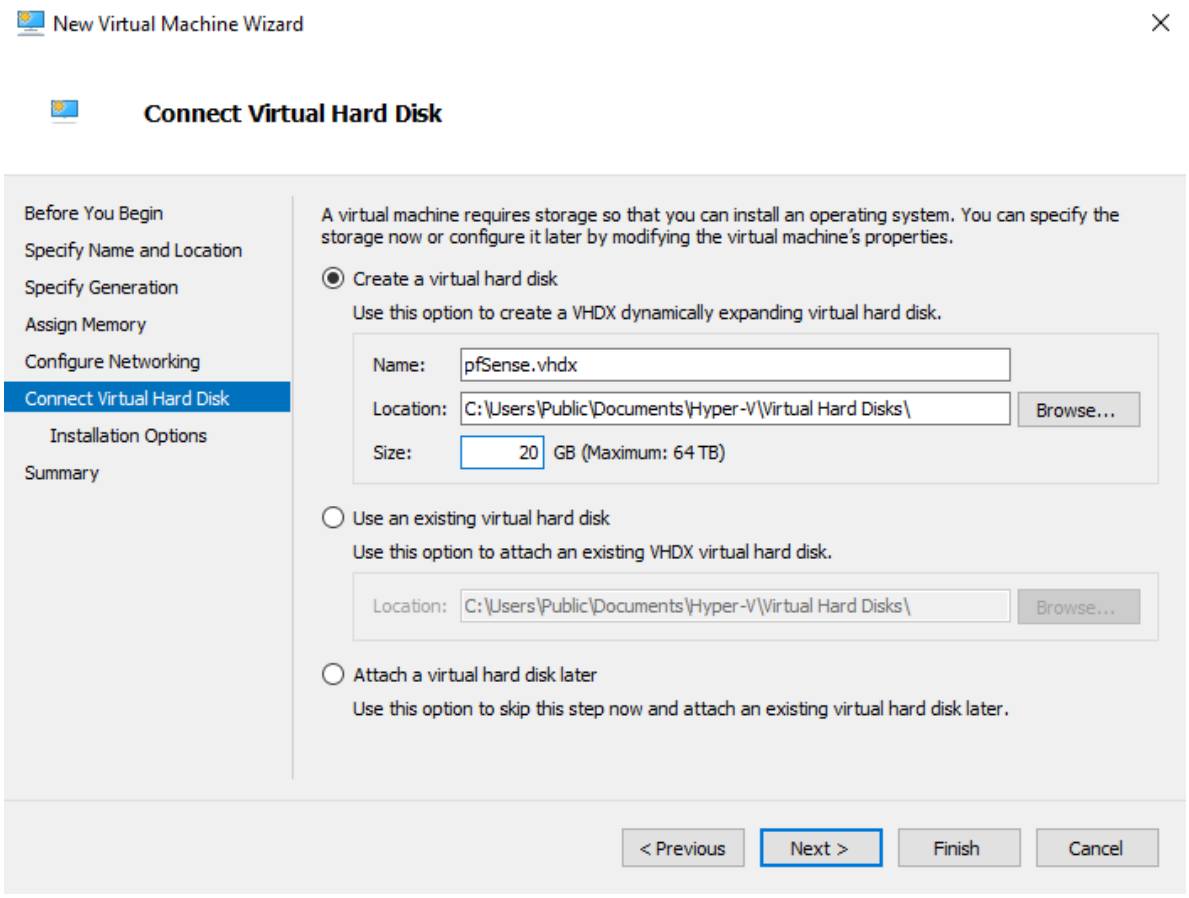
- Click **Next** and proceed to the **Configure Networking** step
- Select WAN from **Connection** drop-down menu

  The LAN will be added later after completing the wizard.

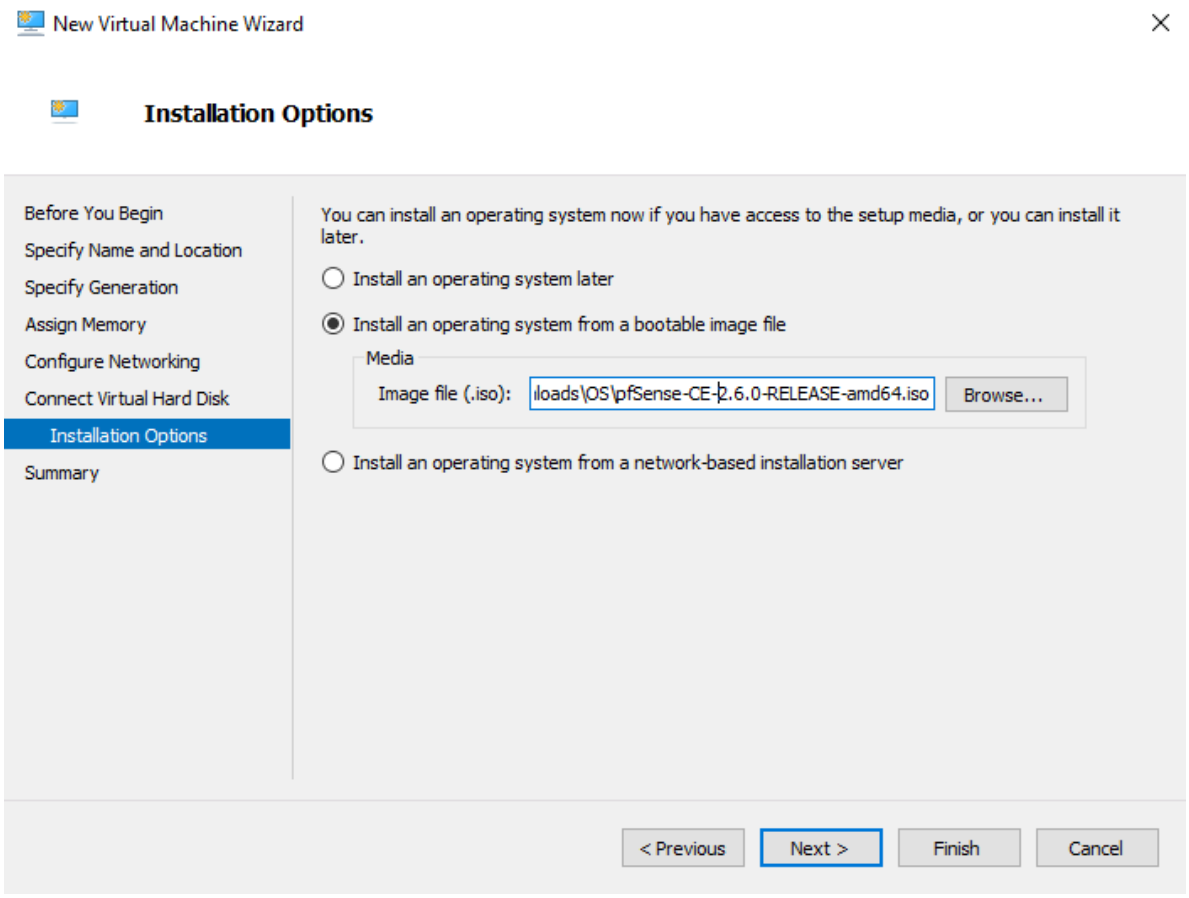


- Click **Next** and proceed to the **Connect Virtual Hard Disk** step
- Select **Create a virtual hard disk**
- Assign `10` to `20` GB for the VM disk

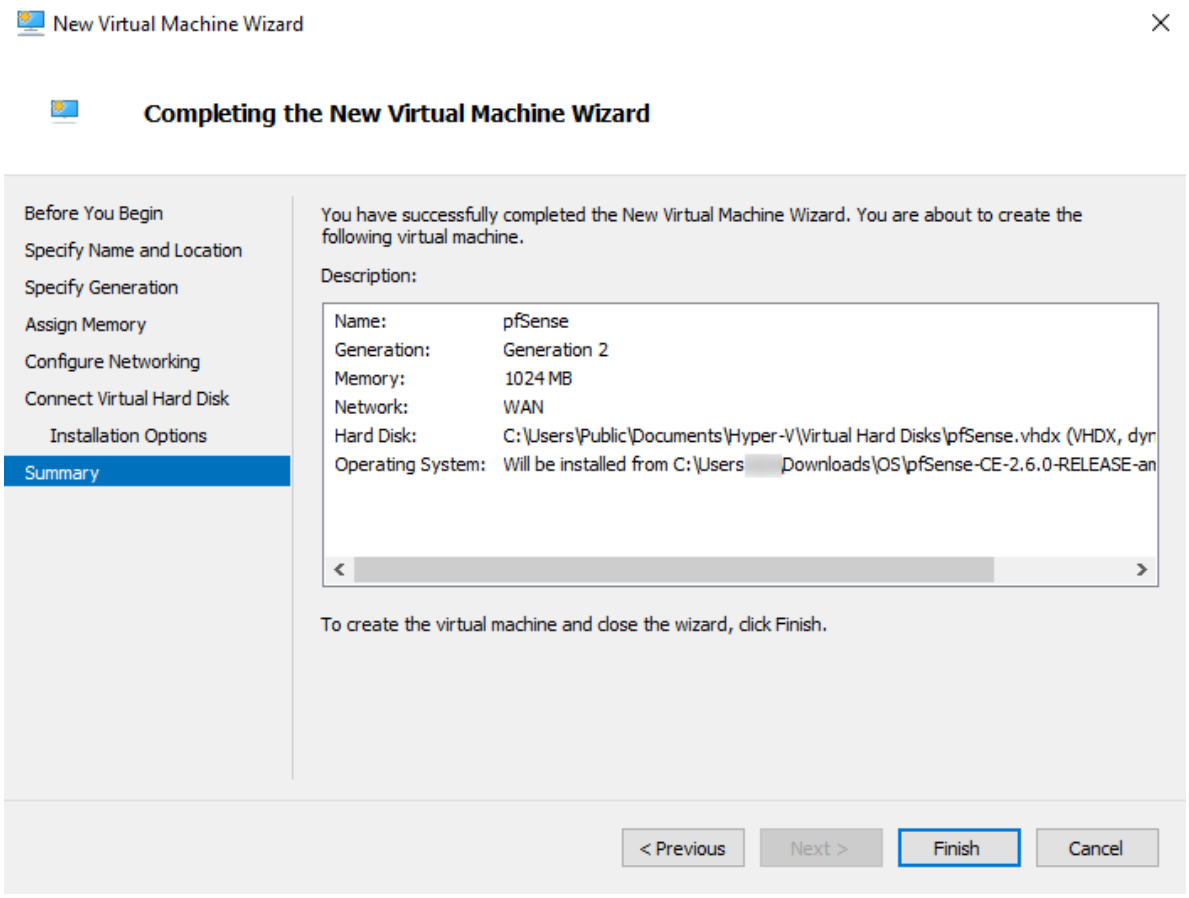
  Disk-intensive tasks such as packages for IDS/IPS or proxies may require larger disk sizes.



- Click **Next** and proceed to the **Installation Options** step
- Select **Install an operating system from a bootable image file**
- Browse to the pfSense software installer ISO image



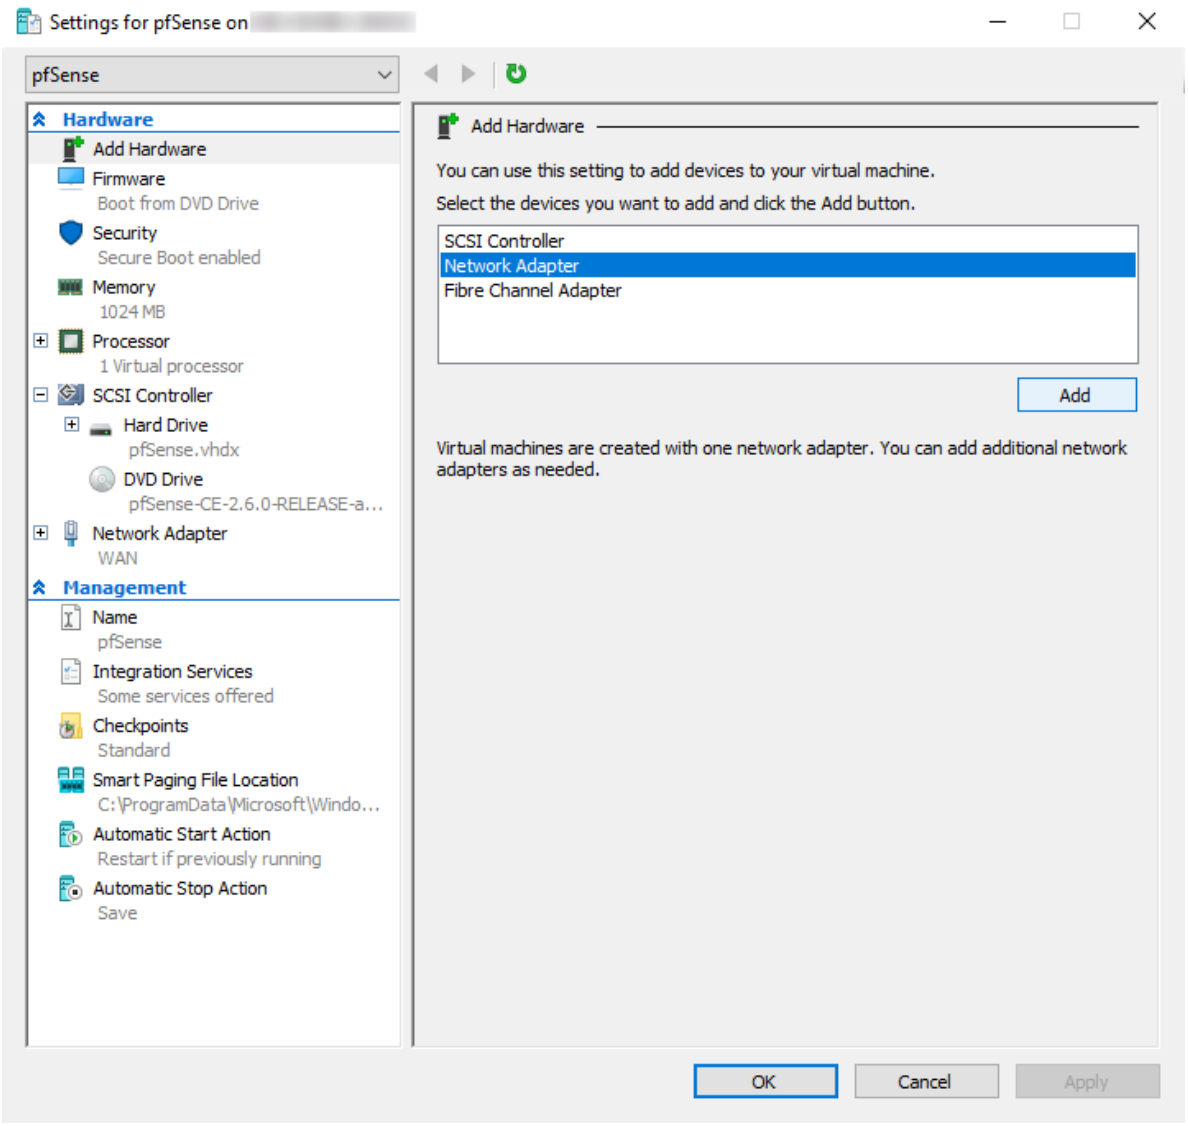
- Click **Next** to display the summary at the end of the wizard
- Review the virtual machine information



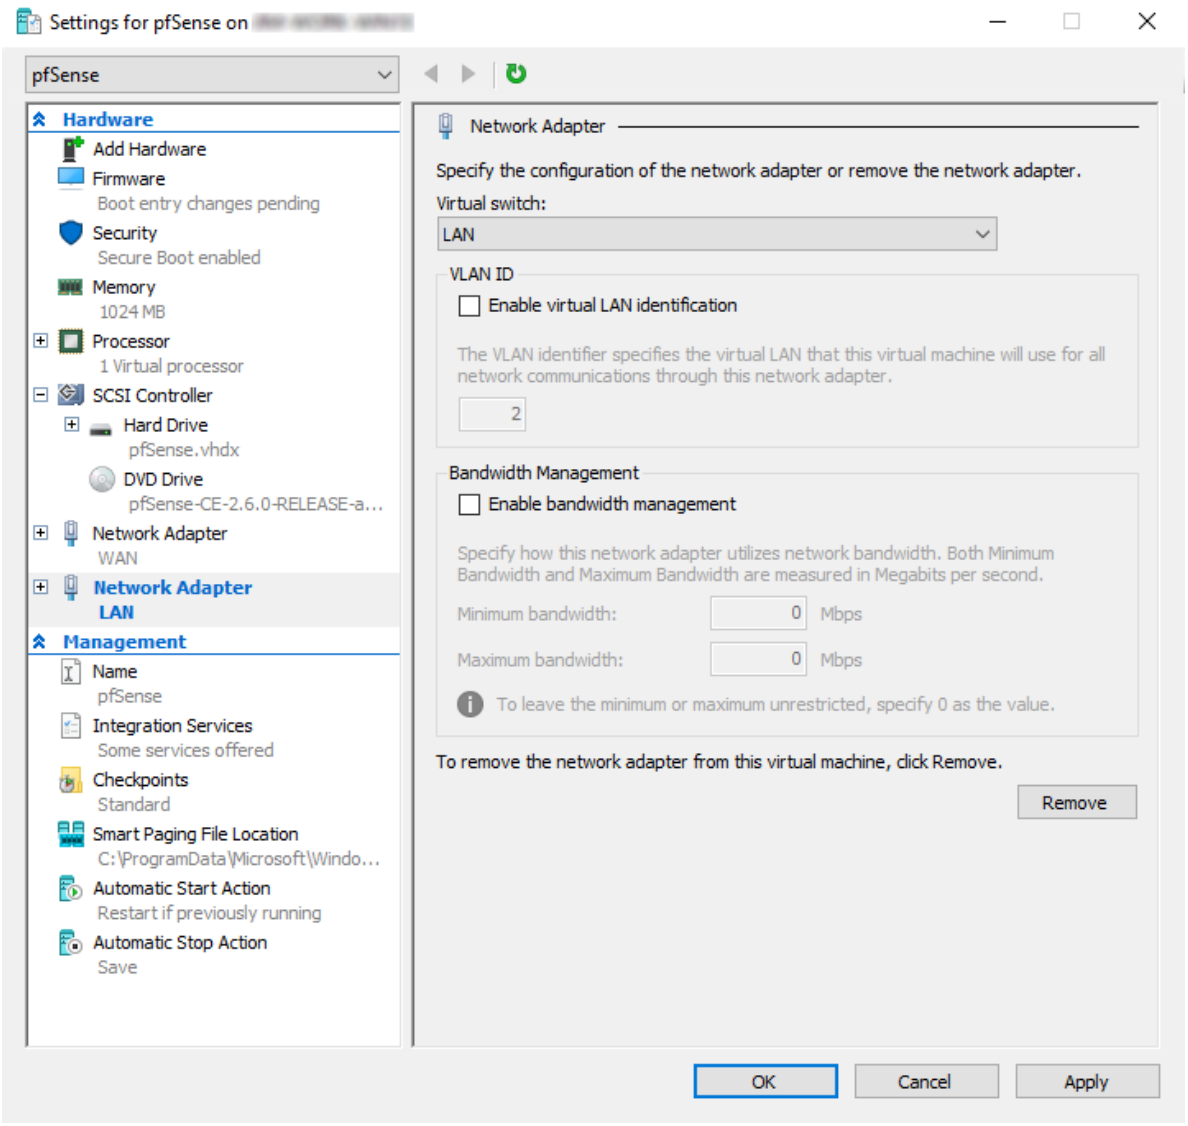
- Click **Finish** if all the information is correct

This completes the wizard, but there are several items which must be set on the VM for it to successfully install and boot pfSense software.

- Select the VM in the **Virtual Machines** list in the Hyper-V Manager
- Click **Settings** on the **Actions** panel for this VM
- Select **Add Hardware** under **Hardware** in the panel on the left side
- Select **Network Adapter**

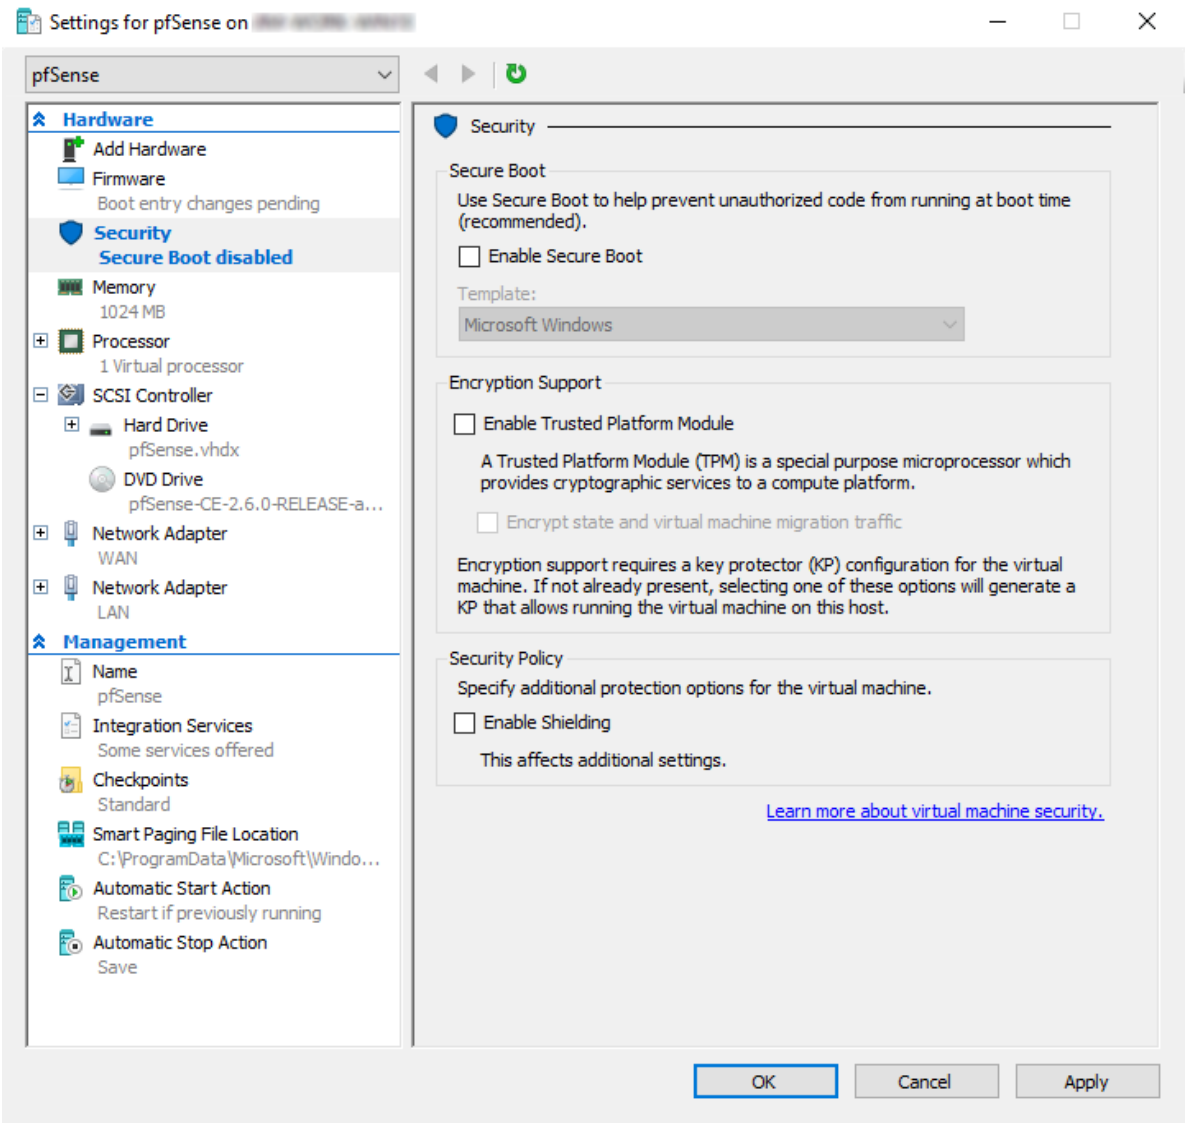


- Set the **Virtual Switch** to the *LAN* switch created earlier

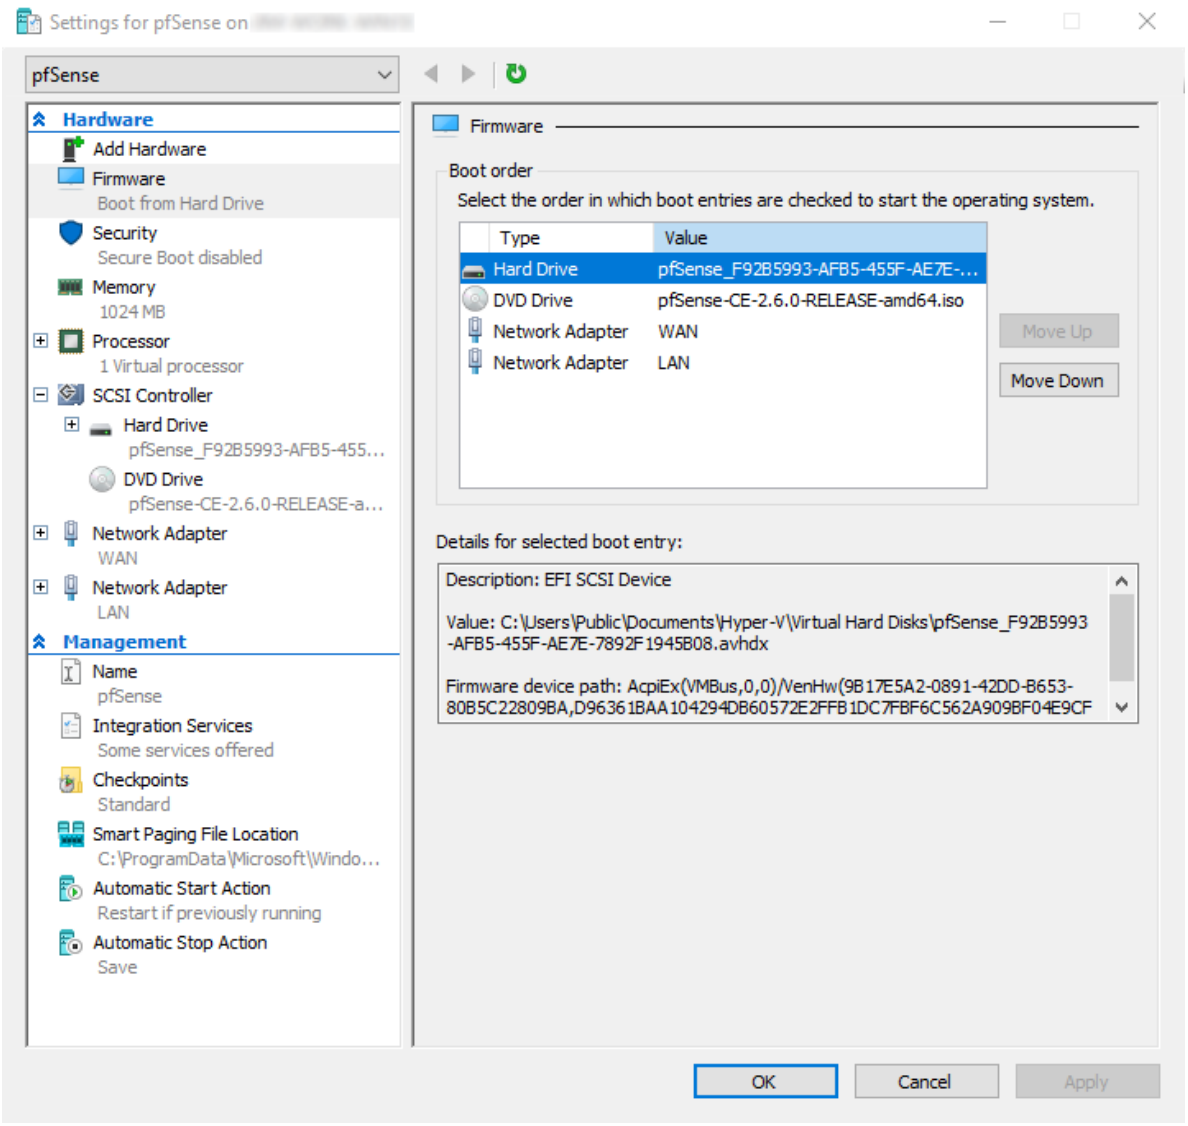


- Click **Apply**
- Select **Security** under **Hardware** in the panel on the left side
- Uncheck **Enable Secure Boot**

> **Warning:** Secure boot must be disabled for the VM to boot pfSense software.



- Click **Apply**
- Select **Firmware** under **Hardware** in the panel on the left side
- Select the **Hard Drive** entry in the **Boot Order** list
- Click **Move Up** until the **Hard Drive** entry is at the *top* of the list

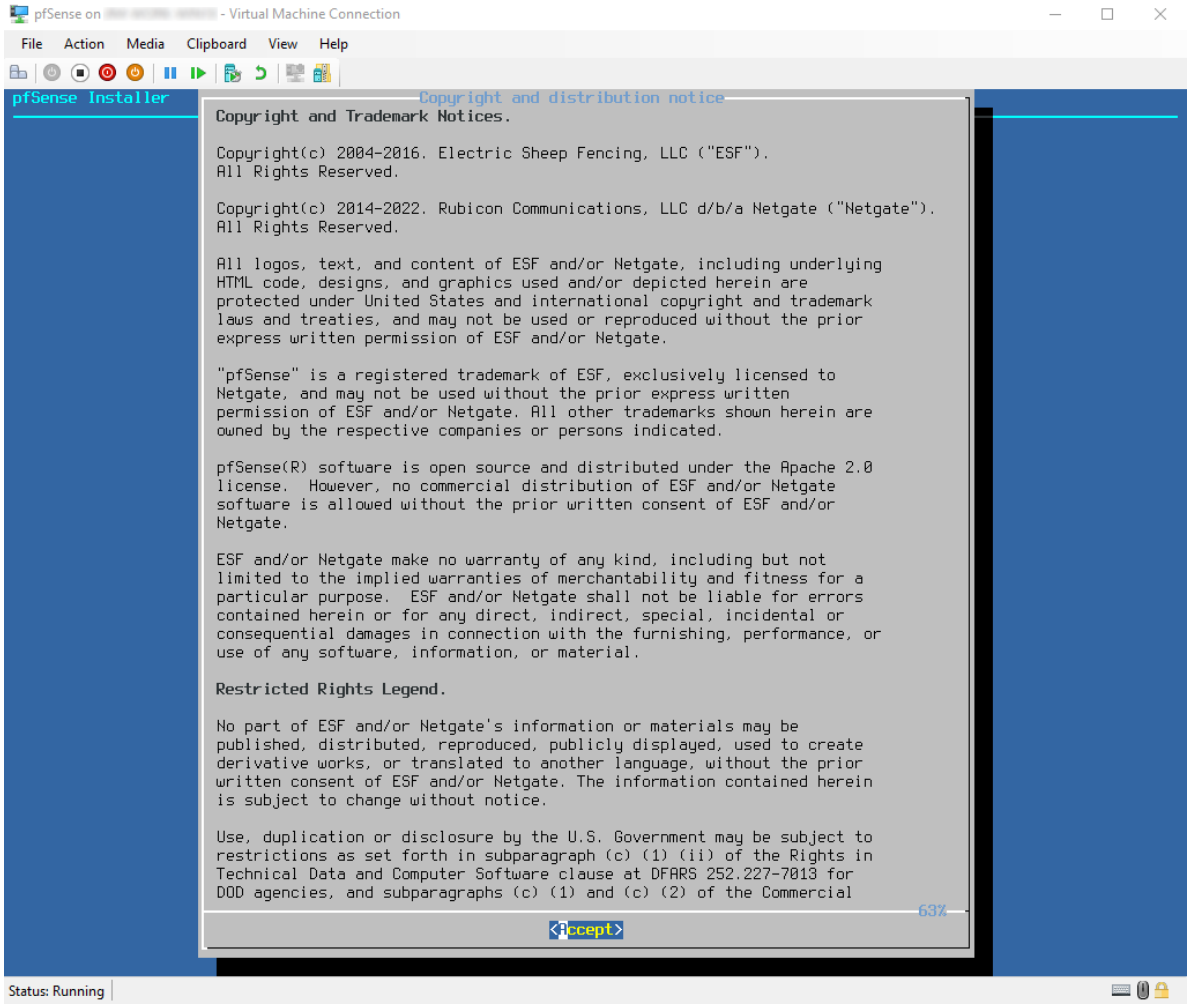


- Click **Apply**
- Review the other VM settings and make the WAN and LAN switches are selected under the respective network adapters
- Click **OK**

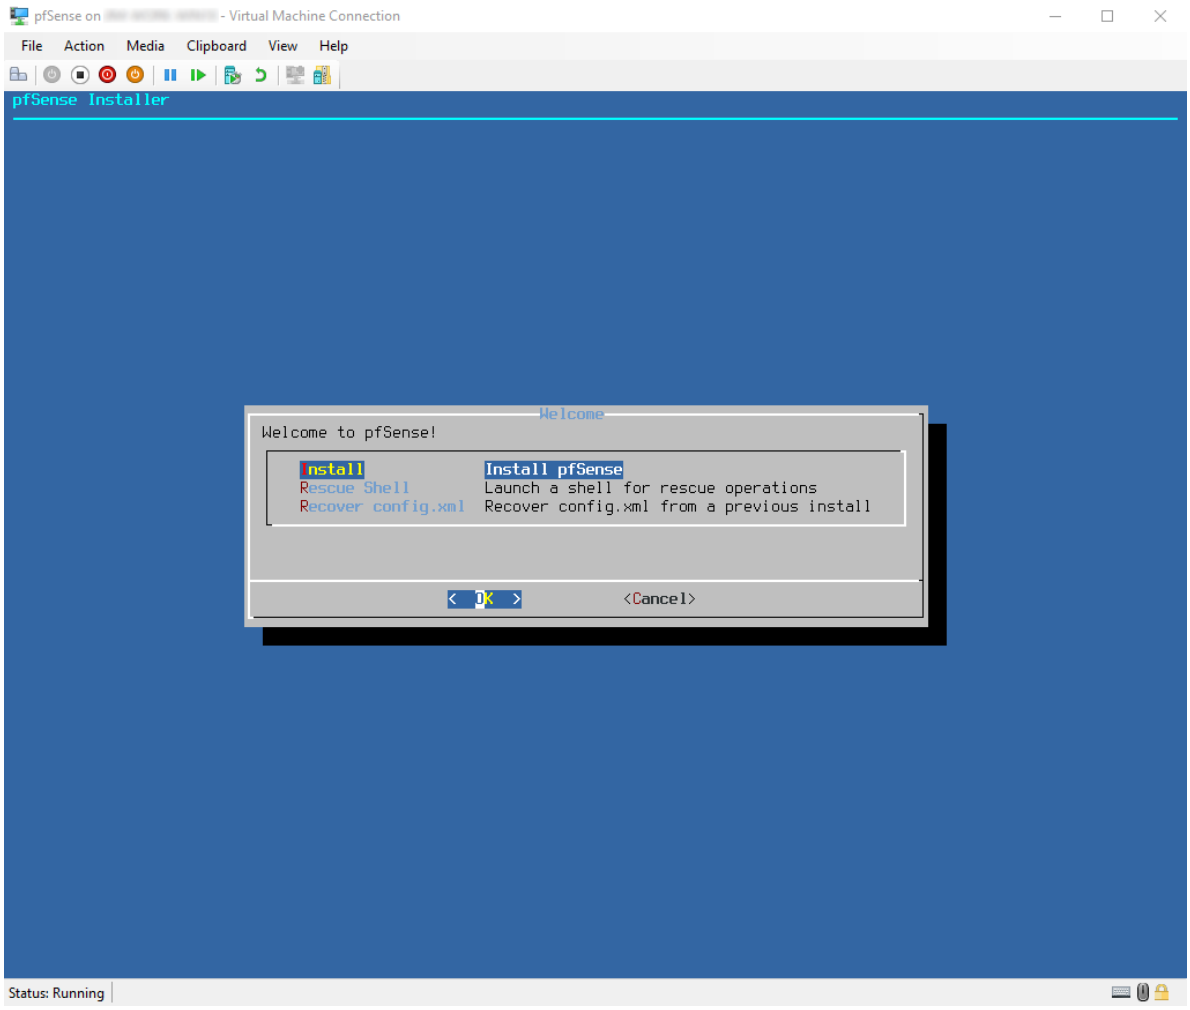
35.70.4 Installing pfSense Software

After successfully creating and configuring the pfSense software virtual machine, it's time to start it.

- Select the VM in the **Virtual Machines** list in the Hyper-V Manager
- Click **Start** from the VM menu in the **Actions** panel
- Click **Connect…** from the VM menu to open a console for the VM
- Wait for the virtual machine to boot and launch the installer



- Read and accept the EULA to display the installation menu



- Proceed through the installation as usual.

  See also:

  See *Installation Walkthrough* for a detailed walkthrough of the installation process.

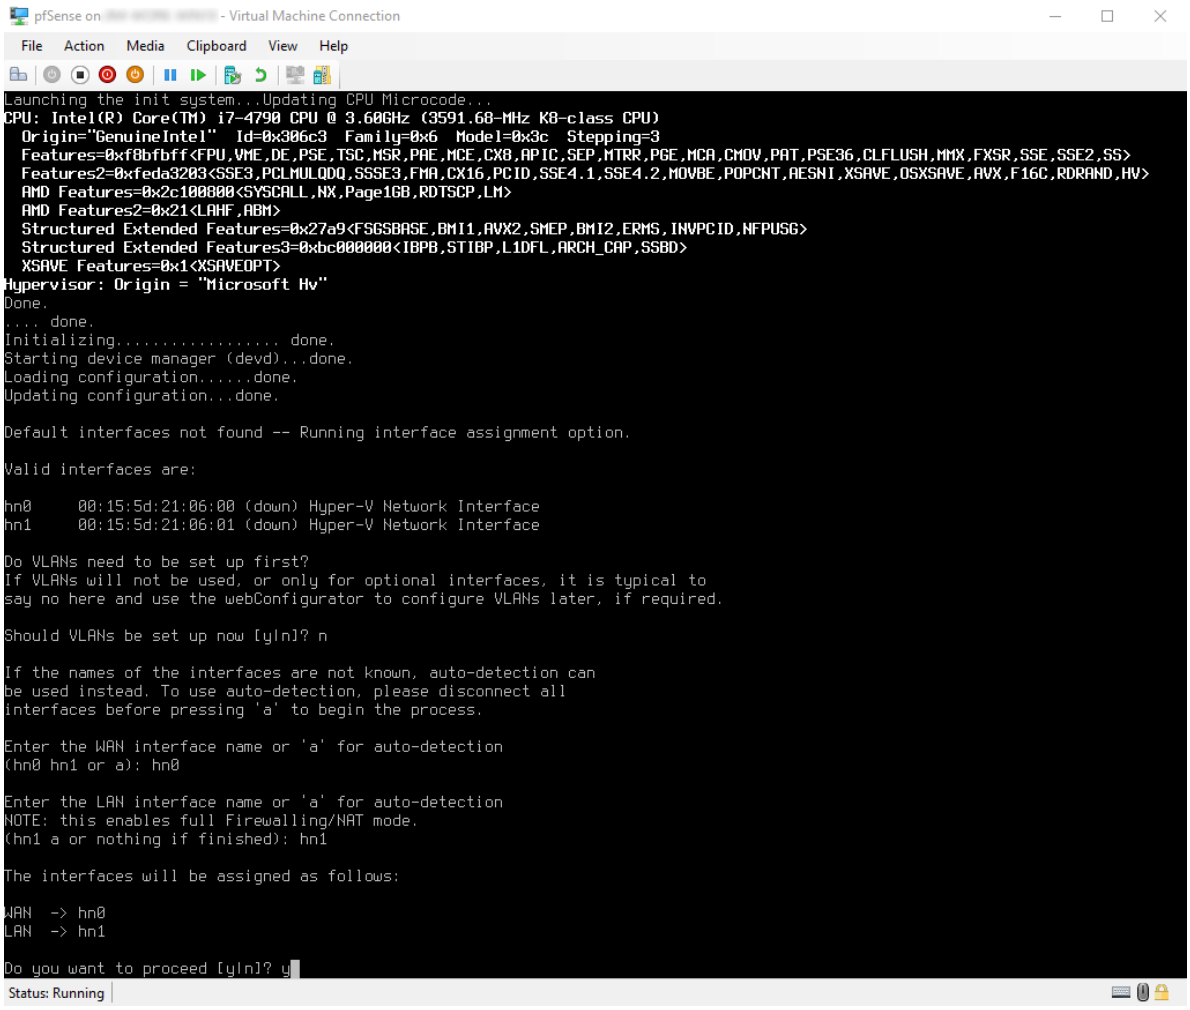
- Finish the installation, select reboot, and eject the ISO from the **Media** menu of the VM console

The VM will restart and begin its first boot.

35.70.5 First boot and interfaces assignment

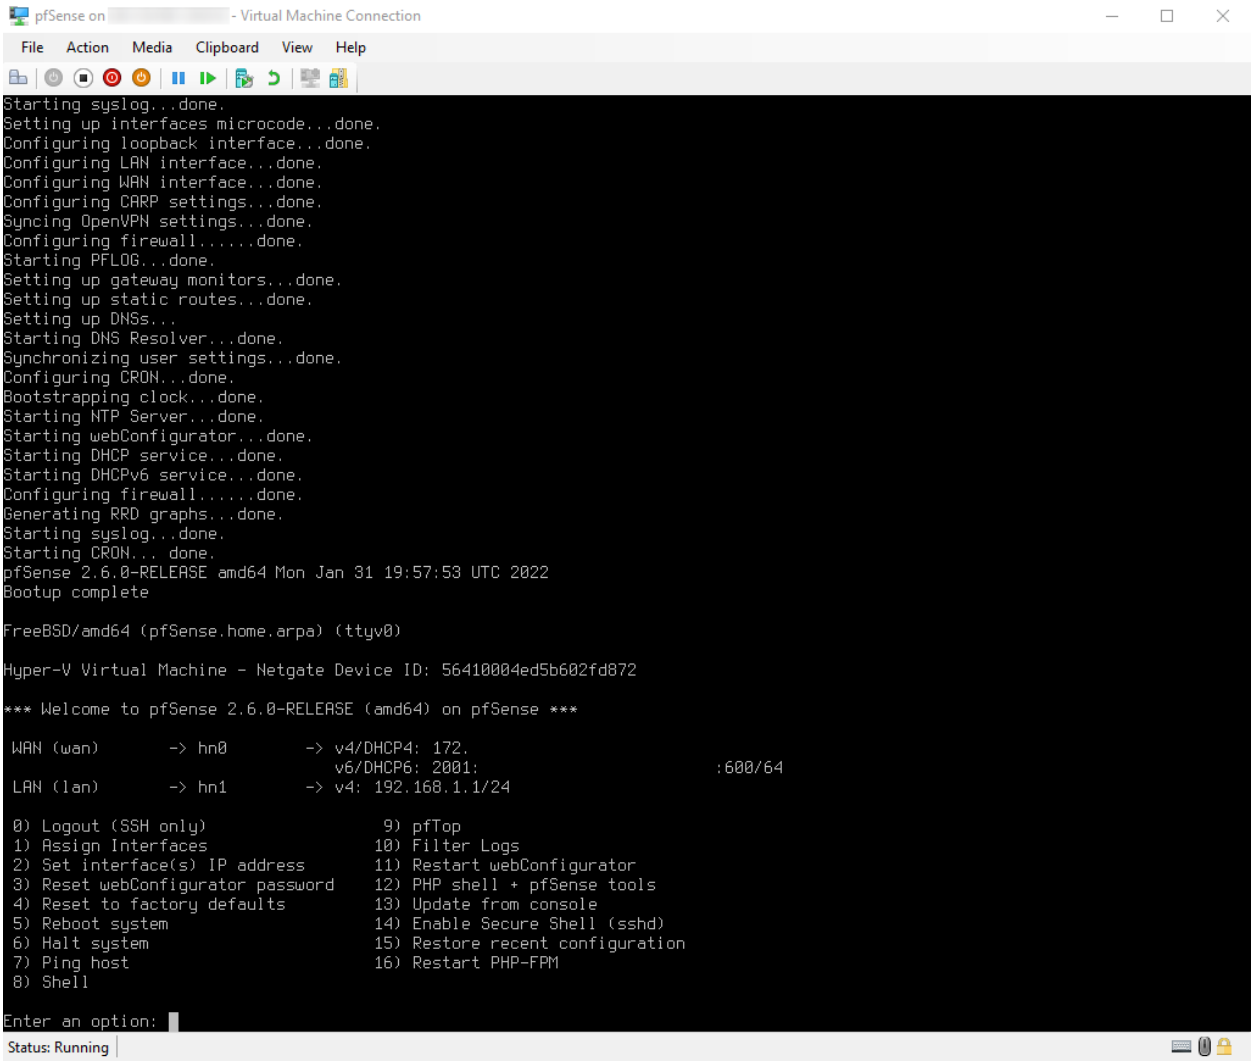
The pfSense software virtual machine should boot up quickly and prompt for interface assignments.

- Enter `n` and press the `Enter` key to skip the VLAN setup
- Enter `hn0` and press the `Enter` key when prompted for the name of the WAN interface
- Enter `hn1` and press the `Enter` key when prompted for the name of the LAN interface
- Enter `y` and press the `Enter` key to proceed



Tip: The MAC addresses printed on the console can be verified against the virtual machine settings to confirm which interface is which.

After assigning interfaces, pfSense software will finish the boot-up. Verify both interfaces have the correct IP addresses.



Congratulations! The virtual machine is now running pfSense software on Microsoft Hyper-V.

From here, proceed through the configuration process for pfSense software as usual. See *Configuration* for details.

35.71 Virtualizing with Proxmox® VE

This following article is about building and running pfSense® software on a virtual machine under Proxmox Virtual Environment (VE). The guide also applies to any newer Proxmox VE version. Article covers Proxmox VE networking setup and firewall virtual machine setup process. The guide does not cover how to install Proxmox VE.

A basic, working, virtual machine will exist by the end of this article.

35.71.1 Assumptions

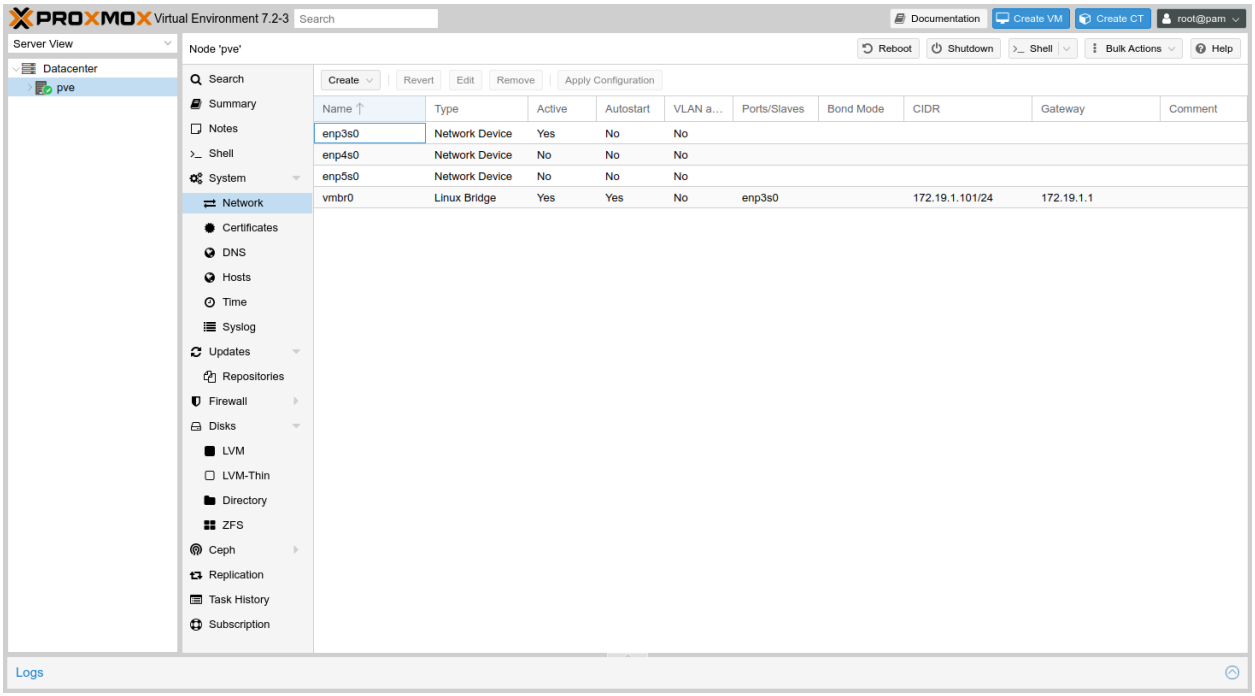
- Proxmox VE host is up and running
- Host has at least two network interfaces available for WAN and LAN.
- pfSense software ISO image is present on the Proxmox VE host

35.71.2 Basic Proxmox VE networking

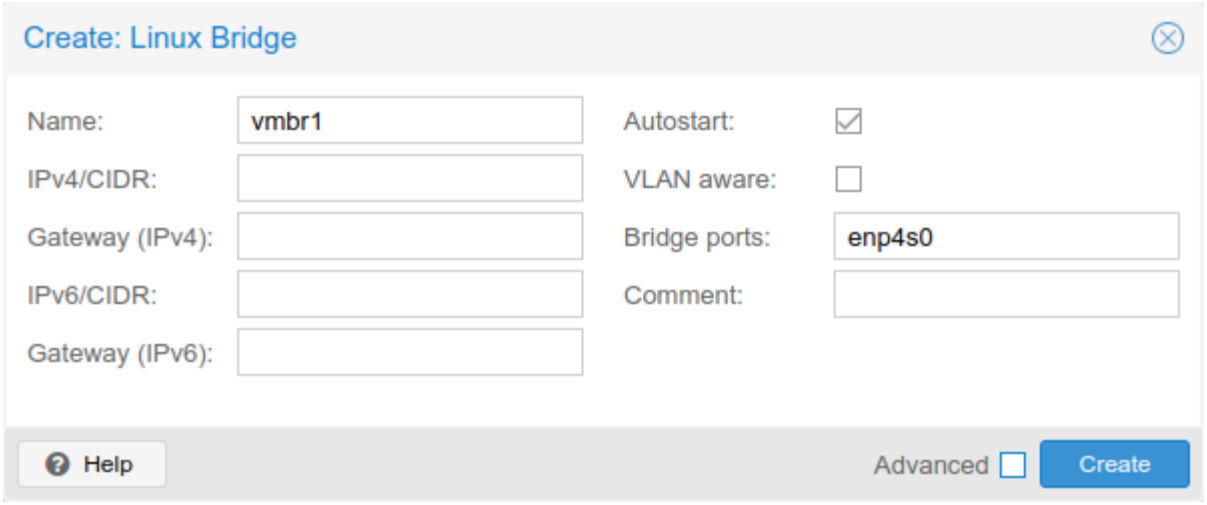
First create two Linux Bridges on Proxmox VE, which will be used for LAN and WAN on the firewall VM.

- Select the host from the server view
- Navigate to **System > Network**

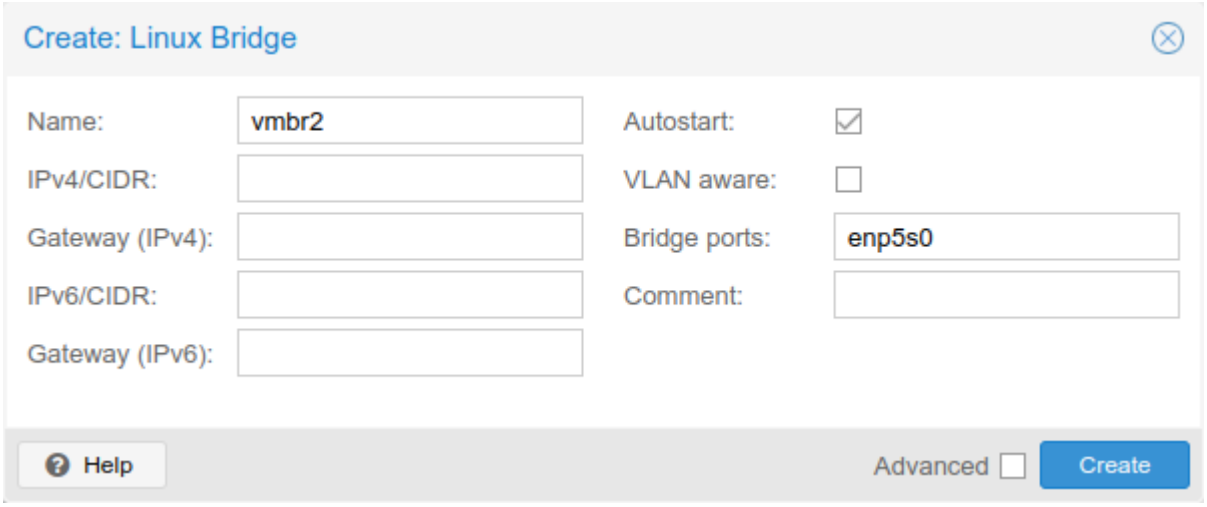
This example uses `enp4s0` and `enp5s0` interfaces for the firewall, while `enp3s0` is for Proxmox VE management. The naming of interfaces will vary depending on the hardware involved (interface type, bus location, etc.).



- Click **Create**
- Select Linux Bridge
- Enter `enp4s0` under **Bridge ports**



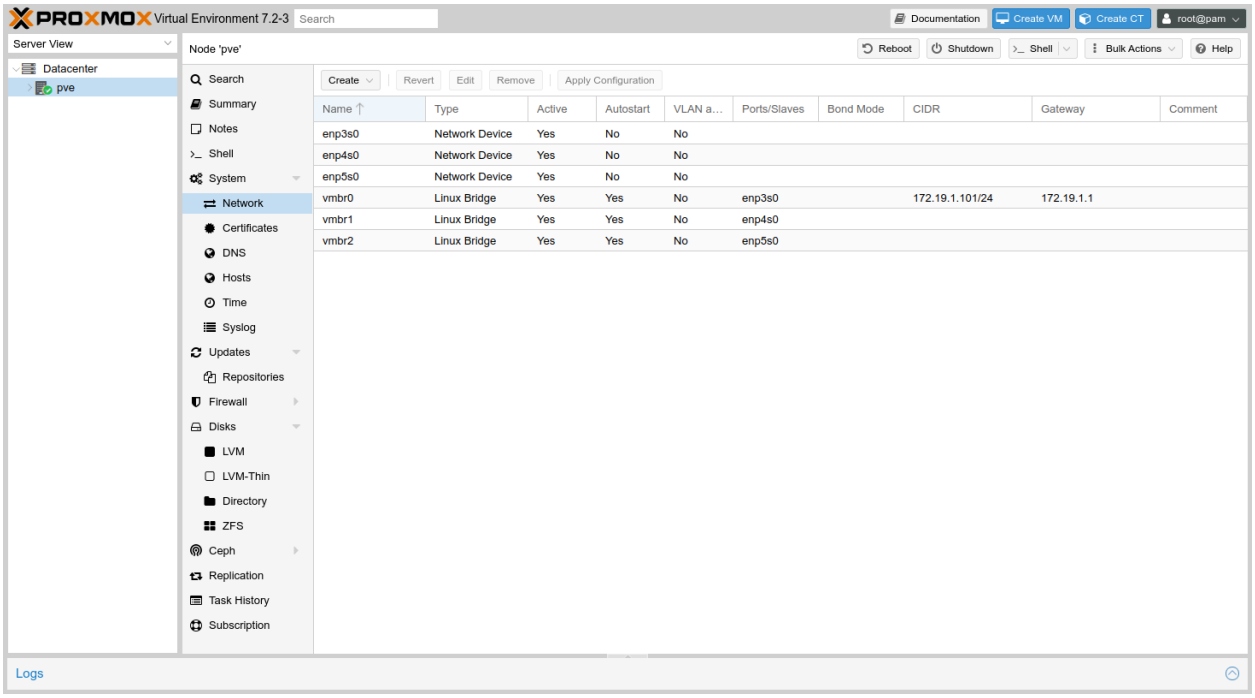
Repeat the process to add another Linux Bridge, this time add `enp5s0` under **Bridge ports**.



- Click **Apply Configuration** to configure the new interfaces in the OS
- Click **Yes** to confirm the action

Proxmox VE networking should now display two Linux bridges like on the following screenshot.

Note: If the interfaces do not show as **Active**, reboot the Proxmox VE host.



35.71.3 Creating a Virtual Machine

After creating WAN and LAN Linux bridges, now proceed to create a new virtual machine.

- Click **Create VM** from the top right section to display the new virtual machine wizard
- Navigate to the **General** tab
- Enter a **Name** for the VM (e.g. `firewall`)
- Navigate to the **OS** tab
- Set the following options:

  Use CD/DVD disc image file
  Selected

  Storage
  local

  ISO image
  Select the previously uploaded ISO image

  Guest OS Type
  Other

- Navigate to the **System** tab
- Set the following options:

  Graphic card
  SPICE

  Note: The SPICE console uses less CPU when idle and implements more advanced console features than the default console. It is compatible with the VNC Proxmox VE console as well as

the more advanced virt-viewer console application.

- Navigate to the **Hard Disk** tab
- Set the following options:

  Bus/Device
  VirtIO Block

  Disk Size
  Enter an appropriate disk size, no less than `8` GB.

- Navigate to the **CPU** tab
- Set the following options:

  Socket
  `1`

  Cores
  `1` or more cores as needed

  Type
  Host to match the CPU on the hypervisor hardware.

  Extra CPU Flags
  These settings adjust the CPU capabilities and behavior of the guest. If using **Host** for **Type** these can likely be left at the default.

  When setting a CPU type other than **Host**, consider setting the **AES** flag to **+** (**On**) which allows the guest to use AES-NI (*Cryptographic Accelerators*).

- Navigate to the **Memory** tab
- Set the following options:

  Memory
  At least `1024` MB

- Navigate to the **Network** tab
- Set the following options:

  Bridge
  vmbr1

  Model
  VirtIO (paravirtualized)

- Navigate to the **Confirm** tab
- Review the settings and make any final corrections if necessary
- Click **Finish**
- Wait for the VM creation process to finish

Now add another network adapter to the VM:

- Expand the **Server View** list on the left to show the contents under **Datacenter** and the name of this hypervisor node (e.g. **pve**, **proxmox**, etc.)
- Select the newly created virtual machine from list
- Click **Hardware** in the right pane

- Click **Add**
- Click **Network Device**
- Set the following options:

  Bridge
  vmbr2

  Model
  VirtIO (paravirtualized)

- Click **Add**

Review the hardware list for the VM and confirm it now contains two network interfaces.

35.71.4 Starting and configuring the virtual machine

After creating a new virtual machine and adding network interfaces, it is time to start the virtual machine.

- Expand the **Server View** list on the left to show the contents under **Datacenter** and the name of this hypervisor node (e.g. **pve**, **proxmox**, etc.)
- Select the newly created virtual machine from list
- Click **Start**
- Click **Console** on the left, under **Summary**

Note: The **Console** button at the top will launch the console in a new window, which depending on the settings may require an additional client installation such as virt-viewer.

When the VM starts it will boot into the installer automatically. From there, follow the installation steps as usual, and reboot when finished.

See also:

See *Installation Walkthrough* for a detailed walkthrough of the installation process.

After the virtual machine reboots, the console will stop at a prompt to assign interfaces.

- Type `n` and press `Enter` to skip VLAN configuration
- Enter `vtnet0` for WAN
- Enter `vtnet1` for LAN
- Press `Enter` if prompted for additional interfaces
- Type `y` and press `Enter` to complete the interface assignment

After interfaces have been assigned, the VM will complete the boot process.

35.71.5 Disable Hardware Checksums with Proxmox VE VirtIO

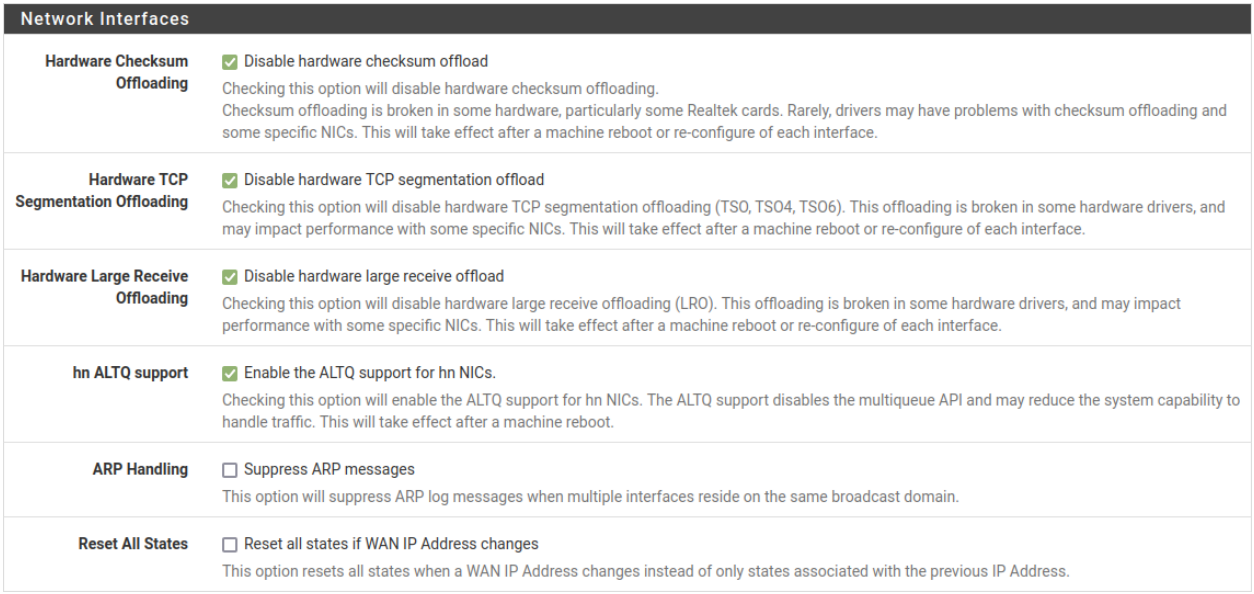
When using VirtIO interfaces in Proxmox VE, network interface hardware checksum offloading **must** be disabled. Current versions of pfSense software attempt to disable this automatically for `vtnet` interfaces, but the best practice is to double-check the setting in case changes in Proxmox VE result in the automatic process failing.

> **Warning:** Do not skip this step, otherwise the virtual machine will not properly pass traffic. Accessing the firewall may be sluggish at first, but changing this setting will correct that as well.

After the installation and interfaces assignment processes are complete, connect to the assigned LAN port from another computer or VM on the LAN-side bridge.

To disable hardware checksum offload:

- Navigate to **System > Advanced**, **Networking** tab
- Locate the **Networking Interfaces** section
- Check **Disable hardware checksum offload**
- Click Save
- Reboot the firewall from **Diagnostics > Reboot** or the console menu



Congratulations, the virtual machine installation and configuration on Proxmox VE is now complete.

35.71.6 Booting UEFI

pfSense software can boot UEFI in a Proxmox VE guest but doing so requires a few extra steps.

When creating the VM:

- Set **Machine** to **q35**
- Set **BIOS** to **OVMF (UEFI)**
- Add an EFI disk when prompted

- Pick the storage for the EFI disk, other settings can remain at defaults

---

Note: An existing non-UEFI VM can be reconfigured to boot UEFI with these settings on its **Hardware**, but the process is more error-prone. For example, the EFI disk is a separate manual process and not semi-automated as it is when creating a VM.

After creating the VM:

- Edit the VM Hardware and add a serial port device

  Note: On some versions of pfSense software the EFI boot process for a Proxmox VE VM works more reliably with a serial port present in the VM hardware, even if the OS is not actively using the port.

On the first boot, go into the boot settings and disable secure boot:

- Hit `Esc` while the boot splash screen is visible
- Select **Device Manager**
- Select **Secure Boot Configuration**
- Uncheck **Attempt Secure Boot**
- Press `F10` to save
- Press `Esc` to exit
- Reset the VM

With secure boot disabled the VM can now boot with UEFI from the ISO as well as after installation.

35.72 General

- *Basic Firewall Configuration Example*
- *Blocking Websites*
- *Using External Wireless Access Points*
- *Using Software from FreeBSD*
- *Using NAT and FTP without a Proxy*
- *Configuring pfSense Software for Online Gaming*
- *Migrating an Assigned LAN to LAGG*
- *Accessing a CPE/Modem from Inside the Firewall*
- *Exporting NetFlow with softflowd*
- *Configuring Switches with VLANs*
- *Using the Shaper Wizard to Configure ALTQ Traffic Shaping*
- *Configuring CoDel Limiters for Bufferbloat*
- *Copy Files to a USB Drive*
- *Diagnostic Data*
- *Changing Credentials and Keys*

- *WAN Connectivity with 802.1X Authentication Bridging and VLAN 0 PCP Tagging*

35.73 DNS

- *Configuring DNS over TLS*
- *Configuring BIND as an RFC 2136 Dynamic DNS Server*
- *Blocking External Client DNS Queries*
- *Redirecting Client DNS Requests*

35.74 Authentication

- *Accessing Files on the Firewall with SCP*
- *Granting Users Access to SSH*
- *External User Authentication Examples*
- *Authenticating Users with Google Cloud Identity*
- *Using EAP and PEAP with FreeRADIUS*
- *Using Mobile One-Time Passwords with FreeRADIUS*
- *Authenticating from Active Directory using RADIUS/NPS*

35.75 Firewall/NAT

- *Allowing Remote Access to the GUI*
- *Preventing RFC 1918 Traffic from Exiting a WAN Interface*
- *Accessing Port Forwards from Local Networks*
- *Configuring NAT for a VoIP PBX*
- *Configuring NAT for VoIP Phones*
- *Configuring NAT64 for IPv6-only Clients*

35.76 Routing

- *Dynamic Routing Protocol Basics*
- *Configuring IPv6 Through A Tunnel Broker Service*
- *Configuring Multi-WAN for IPv6*
- *Routing Public IP Addresses*

35.77 High Availability

- *High Availability Configuration Example*
- *Converting High Availability DHCP from ISC to Kea*
- *High Availability Configuration Example with Multi-WAN*
- *High Availability Configuration Example without NAT*

35.78 IPsec

- *IPsec Site-to-Site VPN Example with Pre-Shared Keys*
- *IPsec Site-to-Site VPN Example with Certificate Authentication*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS*
- *Configuring IPsec IKEv2 Remote Access VPN Clients*
- *IPsec Remote Access VPN Example Using IKEv1 with Xauth*
- *IPsec Remote Access VPN Example Using IKEv1 with Pre-Shared Keys*
- *Routing Internet Traffic Through a Site-to-Site IPsec Tunnel*

35.79 L2TP/IPsec

- *L2TP/IPsec Remote Access VPN Configuration Example*
- *Connecting to L2TP/IPsec from Android*

35.80 OpenVPN

- *OpenVPN Site-to-Site Configuration Example with SSL/TLS*
- *OpenVPN Site-to-Site Configuration Example with SSL/TLS and DCO*
- *OpenVPN Site-to-Site Configuration Example with Shared Key*
- *OpenVPN Remote Access Configuration Example*
- *Adding OpenVPN Remote Access Users*
- *Installing OpenVPN Remote Access Clients*
- *Authenticating OpenVPN Users with FreeRADIUS*
- *Authenticating OpenVPN Users with RADIUS via Active Directory*
- *Connecting OpenVPN Sites with Conflicting IP Subnets*
- *Routing Internet Traffic Through A Site-To-Site OpenVPN Tunnel*
- *Bridging OpenVPN Connections to Local Networks*

- *OpenVPN Site-to-Site with Multi-WAN and OSPF*

35.81 WireGuard

- *WireGuard Remote Access VPN Configuration Example*
- *WireGuard Site-to-Site VPN Configuration Example*
- *WireGuard Site-to-Multisite VPN Configuration Example*
- *WireGuard VPN Client Configuration Example*

35.82 Virtualization

- *Virtualizing pfSense Software with VMware vSphere / ESXi*
- *Virtualizing pfSense Software with Hyper-V*
- *Virtualizing with Proxmox® VE*

CHAPTER

THIRTYSIX

MENU GUIDE

36.1 System

The **System** menu contains choices for the firewall itself, general and advanced options, updates, add-on packages, users, and routing.

Advanced
Advanced settings for the firewall, hardware, SSH, notifications, tunables, and many others.

See *Advanced Configuration Options*.

Certificates
Manage Certificate Authorities, Certificates, and Certificate Revocation Lists (x.509).

See *Certificate Management*.

General Setup
General settings such as hostname, domain, and DNS servers.

See *General Configuration Options*.

High Availability
Controls how nodes running pfSense® software in a High Availability (HA) cluster synchronize states and configuration.

See *High Availability Synchronization Settings*.

Logout
Logs out of the GUI, returning the user to the login screen.

See *User Management and Authentication*.

Package Manager
Additional software add-ons for pfSense software to expand its functionality.

See *Packages*.

Routing
Manage gateways, static routes, and gateway groups for using multiple WANs.

See *Routing*.

Setup wizard
The Setup Wizard performs the basic initial configuration.

See *Setup Wizard*.

Update
Upgrade pfSense software to the latest version.

See *Upgrading using the GUI*.

User Manager
Manage users, groups, and authentication servers (RADIUS or LDAP) for GUI access, VPN access, etc.

See *User Management and Authentication*.

User Password Manager
Self-service user password manager. Allows GUI users to change their own password.

See *User Password Manager*

User Settings
If per-user settings are enabled, this page provides a way for users to override default behavior options found under **General Setup**.

36.2 Interfaces

The **Interfaces** menu contains an entry for assigning interfaces along with entries for each currently assigned interface. Menu entries for assigned interfaces use the configured names, or the standard names if they have not been changed (e.g. WAN, LAN, OPTx)

Assignments
Assign interfaces to logical roles (e.g. WAN, LAN, OPT), and create/maintain VLANs and other types of virtual interfaces.

See *Interface Configuration*, *Interface Types and Configuration*, and *Virtual LANs (VLANs)*.

WAN
Configure the WAN interface.

See *Interface Configuration*.

LAN
Configure the LAN interface.

See *Interface Configuration*.

OPTx
Configure any additional optional interfaces.

See *Interface Configuration*.

36.3 Firewall

The **Firewall** menu entries configure firewall rules, NAT rules, and related functionality.

Aliases
Manages collections of IP addresses, networks, or ports to simplify rule creation and management.

See *Aliases*.

NAT
Manages NAT rules that control port forwards, 1:1 NAT, and Outbound NAT behavior.

See *Network Address Translation*.

Rules
Configures firewall rules. This page contains one tab for each configured interface, plus tabs for groups and different VPN types, when enabled.

See *Introduction to the Firewall Rules screen*.

Schedules
Manages time-based rule schedules.

See *Time Based Rules*.

Traffic Shaper
Manages traffic shaping/Quality of Service (QoS) settings.

See *Traffic Shaper*.

Virtual IPs
Configure Virtual IP addresses which allow pfSense® software to handle traffic for more than one IP address per interface, typically for NAT rules or High Availability.

See *Virtual IP Addresses*.

36.4 Services

The **Services** menu contains items which control services provided by daemons running on the firewall.

See also:

Services

Auto Config Backup
Configures the automatic configuration backup service which optionally uploads encrypted configuration backups for secure off-site backups.

See *Automatic Configuration Backup Service*.

Captive portal
Controls the Captive Portal service which directs users to a web page for authentication before permitting Internet access.

See *Captive Portal*.

DHCP relay
Configures the DHCP relay service which proxies DHCP requests from one network segment to another.

See *DHCPv4 & DHCPv6 Relay*.

DHCP server
Configures the DHCP service which provides automatic IP address configuration for clients.

See *DHCPv4 Server*.

DHCPv6 Relay
Configures the DHCP relay service for IPv6 which proxies DHCPv6 requests from one network segment to another.

See *DHCPv4 & DHCPv6 Relay*.

DHCPv6 Server
Configures the DHCP service for IPv6 which provides automatic IPv6 address configuration for clients.

See *DHCPv6 Server*.

DNS Forwarder
Configures the built-in caching DNS forwarder.

See *DNS Forwarder*.

DNS Resolver
Configures the built-in caching DNS resolver.

See *DNS Resolver*.

Dynamic DNS
Configures Dynamic DNS services ("DynDNS") which updates a remote name server when the WAN IP address of this firewall changes.

See *Dynamic DNS*.

IGMP Proxy
Configures the Interior Group Management Protocol proxy for passing multicast traffic between interfaces.

See *IGMP Proxy*.

NTP
Configures the Network Time Protocol server daemon.

See *NTPD*.

PPPoE Server
Configures the PPPoE server which accepts and authenticates connections from PPPoE clients on local networks.

See *PPPoE Server*.

Router Advertisement
Configures IPv6 Router Advertisement behavior which provides IPv6 routing and address configuration for clients.

See *IPv6 Router Advertisements*.

SNMP
Configures the Simple Network Management Protocol (SNMP) daemon to allow network-based collection of statistics from this firewall.

See *SNMP*.

UPnP IGD & PCP
Configures the Universal Plug and Play (UPnP IGD) & Port Control Protocol (PCP) service which automatically configures NAT and firewall rules for devices capable of using the UPnP IGD or PCP standards. This menu entry only appears when the firewall contains more than one assigned interface.

See *UPnP IGD & PCP*.

Wake on LAN
Configures Wake on LAN entries which remotely wake up local client devices.

See *Wake on LAN*.

36.5 VPN

The **VPN** menu contains items pertaining to Virtual Private Networks (VPNs), including IPsec, OpenVPN and L2TP.

See also:

Virtual Private Networks

IPsec
Configure IPsec VPN tunnels, mobile IPsec, and IPsec settings.

See *IPsec*.

L2TP
Configure L2TP services and users.

See *L2TP VPN*.

OpenVPN
Configure OpenVPN servers and clients, as well as client-specific configuration.

See *OpenVPN*.

36.6 Status

The **Status** menu entries display status information and logs for various system components and services.

Captive Portal
When Captive Portal is enabled, this entry shows user and voucher status.

See *Captive Portal Status*.

CARP (failover)
Shows the status of CARP IP addresses on this firewall, such as MASTER/BACKUP state for each CARP VIP. Also has controls for HA maintenance mode.

See *CARP Status*.

Dashboard
A shortcut back to the main page of the firewall GUI, which displays general system information.

See *Dashboard*.

DHCP leases
Shows a list of all IPv4 DHCP leases assigned by this firewall and provides controls based on those leases, such as adding static mappings.

See *DHCPv4 Leases*.

DHCPv6 leases
Shows a list of all IPv6 DHCP leases assigned by this firewall.

See *DHCPv6 Status*.

DNS Resolver
Shows the contents of the DNS resolver infrastructure cache.

See *DNS Resolver Status*.

Filter Reload
Shows the status of the last filter reload request, including active reload actions. Also provides a means to force a filter reload, and to force an XMLRPC configuration sync when HA is configured.

See *Filter Reload Status*.

Gateways
Shows the status of gateways, and gateway groups for multiple WANs.

See *Gateway Status*.

Interfaces
Shows the hardware status for network interfaces, equivalent to using `ifconfig` on the console.

See *Interface Status*.

IPsec
Shows the status of any configured IPsec tunnels.

See *IPsec Status*.

Monitoring
Shows graphed data for system statistics such as bandwidth used, CPU usage, firewall states, etc.

See *Monitoring Graphs*.

NTP
Shows the status of the Network Time Protocol server daemon.

See *NTP Daemon Status*.

OpenVPN
Shows the status of any configured OpenVPN instances.

See *OpenVPN Server and Client Status*.

Queues
Shows the status of traffic shaping queues.

See *ALTQ Traffic Shaper Queue Monitoring*.

Services
Shows the status of system and package service daemons.

See *Service Status*.

System logs
Shows logs from the system and system services such as the firewall, DHCP, VPNs, etc.

See *System Logs*.

Traffic graph
Displays a dynamic real-time traffic graph for an interface.

See *Traffic Graphs*.

UPnP IGD & PCP
Shows a list of any currently active UPnP IGD & PCP port forwards. This entry is only present when the firewall contains more than one interface.

See *UPnP IGD & PCP Status*.

Wireless
Shows a list of any currently available wireless networks in range, along with signal levels. This menu entry is only present if the firewall has an assigned wireless interface.

See *Wireless Status*.

36.7 Diagnostics

Items under the **Diagnostics** menu perform various diagnostic and administrative tasks.

ARP Table
Displays a list of devices as seen locally by the firewall. The list includes an IP address, MAC address, Hostname, the Interface where the device was seen, and other related information.

See *ARP Table*.

Authentication
Tests authentication to a defined RADIUS or LDAP server.

See *Troubleshooting Authentication*.

Backup & Restore
Backup and restore configuration files.

See *Backup and Recovery*.

Command Prompt
Execute shell commands or PHP code, and upload/download files to/from the firewall.

> **Warning: Use with caution**

See *Command Prompt*.

DNS Lookup
Executes a DNS lookup to resolve hostnames for diagnostic purposes, and to test connectivity to DNS servers.

See *DNS Lookup*.

Edit File
Edit a file on the firewall filesystem.

See *Editing Files on the Firewall*.

Factory defaults
Resets the configuration back to default. Be aware, however, that this does not alter the filesystem or uninstall package files; it only changes configuration settings.

See *Resetting to Factory Defaults*.

GEOM Mirrors
If the firewall contains a GEOM disk mirror, this page shows the status of the mirror and provides controls for managing the mirror.

Halt system
Shuts down the firewall and turns off the power where possible.

See *Halting and Powering Off the Firewall*.

Limiter Info
Shows the status of any Limiters and the traffic flowing inside them.

See *Checking Limiter Usage*.

NDP Table
Shows a list of local IPv6 devices as seen by the firewall. The list includes an IPv6 address, MAC address, hostname (if known to the firewall), and the interface.

See *NDP Table*.

Packet Capture
Perform a packet capture to inspect traffic, and then view or download the results.

See *Packet Capture GUI*.

pfInfo
Displays statistics about the packet filter, including general traffic rates, connection rates, state table info, and various other counters.

See *pfInfo*.

pfTop
Displays a list of the top active connections by a selectable metric such as bytes, rate, age, etc.

See *pfTop*.

Ping
Sends ICMP echo requests to a given IP address, sent via a chosen interface.

See *Ping Host*.

Reboot
Reboots the firewall. This can take several minutes to complete, depending on the hardware and enabled features.

See *Rebooting the Firewall*.

Routes
Shows the contents of the routing table.

See *Route Table Contents*.

SMART Status
Displays diagnostic information about disk drives on compatible hardware. Can also run drive tests.

See *S.M.A.R.T. Hard Disk Status*.

Sockets
Displays a list of processes on the firewall that are bound to network ports, listening for connections or making connections outbound from the firewall itself.

See *Viewing Active Network Sockets*.

States
Shows the currently active firewall states.

See *Viewing Firewall States in the GUI*.

States Summary
Displays information about the state table, to see activity summarized by IP address.

See *Firewall States Summary*.

System Activity
Shows memory usage and a list of active processes and system threads on the firewall, the output is from `top -aSH`.

See *System Activity (Top)*.

Tables
Displays and edits the contents of various firewall tables and aliases.

See *Firewall Table Contents*.

Test Port
: Performs a simple TCP connection test from the firewall to determine if a remote host is accepting connections on a specified port.

  See *Testing a TCP Port*.

Traceroute
: Trace the route taken by packets between this firewall and a remote system.

  See *Traceroute*.

This section is a guide to the standard menu choices available in pfSense® software. This guide will help to quickly identify the purpose of a given menu option, and refer to places in the documentation where those options are discussed in further detail.

Packages can add items to any menu, so check each menu or consult the documentation for a package to locate its menu entries. Typically, packages install entries under the **Services** menu, but there are numerous exceptions.

CHAPTER

THIRTYSEVEN

GLOSSARY OF TERMS

DNS
An acronym for Domain Name System.

ICMP
An acronym for Internet Control Message Protocol.

HTTP
An acronym for Hypertext Transfer Protocol.

IP
An acronym for Internet Protocol.

LAN
An acronym for Local Area Network.

NAT
An acronym for Network Address Translation.

SSH
An acronym for Secure Shell.

TCP
An acronym for Transmission Control Protocol.

UDP
An acronym for User Datagram Protocol.

WAN
An acronym for Wide Area Network.

VM
An acronym for Virtual Machine.

VPN
An acronym for Virtual Private Network.

Note: Though technically abbreviations read letter by letter are initialisms, not acronyms, this document refers to both as acronyms to be more accessible to readers.

CHAPTER

THIRTYEIGHT

DEVELOPMENT

These articles cover advanced topics related to developing on or with pfSense® software.

38.1 General Development Information

38.1.1 Software Release Schedule

pfSense Plus software typically targets 3 releases per year, but the schedule varies based on complexity of development for a given release.

Releases for pfSense® CE software are made when they are ready. A public schedule is not available at this time, but release announcements and progress messages are made on the Netgate Blog.

See also:

Some information on upcoming releases can be found in *Versions of pfSense software and FreeBSD*.

Netgate provides maintenance releases of pfSense CE software as needed, typically a couple per year. These include primarily bug fixes and security updates.

To follow the progress of a release, visit the pfSense Redmine.

38.1.2 Reporting Issues with pfSense Software

This page serves as a guide for providing legitimate bug reports with pfSense® software. Most of the bug reports from outside users are not bugs at all, but incorrect configurations. Often user reports do not contain nearly enough information for a developer to replicate the problem. If a bug report does not contain the appropriate information to verify a legitimate bug, developers have no choice but to reject the report.

Attention: The pfSense Redmine site is not a discussion platform and is never to be used for assistance requests.

The first step is typically to ask about the problem using one of the available resources for assistance, such as the Netgate Forum. An issue report can be opened if a specific bug is found that can be reproduced by developers.

An example of an invalid bug report is `XYZ doesn't work!` without appropriate accompanying detail. An alternative, to make that a legitimate bug report, is:

```
Subject:
  Feature XYZ generates an invalid configuration for option A on
  ``thispage.php``
```

(continues on next page)

(continued from previous page)

```
Description:
  The underlying @xyz.conf@ has @option1=1@ where it should be @option1=5@
  when option A is checked on @thispage.php@ in the web interface.
```

See also:

How to report bugs effectively.

Where to submit

For anything that is not a confirmed, specific, detailed bug report, post to the Netgate Forum first. The forum is a platform where the problem can be discussed and the specifics required to replicate the issue can be identified.

After discussion and confirmation of a specific, legitimate bug report on the Netgate Forum, please open a ticket in the pfSense Redmine, including a link back to the discussion in question.

Attention: The pfSense Redmine site is not a discussion platform and is never to be used for assistance requests.

What to Include

When submitting a bug report to the pfSense Redmine site, fill out the report completely and include enough supporting information to reproduce the issue.

Use the following guidelines when completing the issue report.

File issues with the base system under pfSense and issues with packages under pfSense Packages.

Note: Not all fields will be available to all users.

Tracker
: Use the appropriate issue tracker type, following these guidelines:

  Bug
  : Problems, unexpected behavior, crashes, or when the other categories do not apply.

  Regression
  : A function that worked previously but failed during development or when upgrading to a more recent release.

  Feature
  : Feature requests or changes in (working) behavior.

  Todo
  : Tasks or work that need to be completed that are not specifically bugs or features.

Subject
: A brief but complete and accurate description of the problem. If the problem is specific to one page or file, prefix the subject with that page filename.

  Example:

```
Save and Force Update button does not perform any action on ``services_
↪rfc2136_edit.php``
```

Description
A full description of the problem. Include any relevant detail, supporting evidence such as log entries, and if possible a complete recount of how to reproduce the bug that includes every necessary step.

> **Warning: Information in this issue tracker is public!** Do not include any personal information such as usernames, passwords, private certificates, keys, e-mail addresses, IP addresses, and so on. In the rare case when such information is helpful in diagnosing a problem, it can be transmitted privately.

- Please use appropriate formatting, such as `<pre>` tags around log data or command output. Attach lengthy output in a text file rather than including it inline in the description.
- Attach files with supporting information, such as logs or crash dumps, if relevant. Check for personal data before attaching files and obfuscate/mask info as needed.

Status
Leave this as *New*.

Priority
Leave this as *Normal* or set lower for minor issues.

> **Warning: Higher priorities must only be set by developers when appropriate**.

Assignee
Leave this **empty** unless directed by a developer to assign it to them directly.

Category
Pick the closest relevant category for the issue, if possible, or a generic category such as "Web Interface" for GUI issues and "Operating System" for OS issues.

Target Version/Plus Target Version
Leave this **empty** unless directed by a developer to assign a specific target.

> **Warning:** Developers will assign a target version after evaluating the issue.

Affected Version
Pick the version number of the firewall experiencing the issue.

Affected Architecture
Pick *All* unless the issue is known to only affect a single architecture.

Leave other fields **blank**, such as **Parent Task**, **Start/End Date**, and so on.

38.1.3 Obtaining Panic Information for Developers

Crash dump functionality is built into every kernel on current versions of pfSense® software, but behavior varies based on architecture. Crash dumps are automatically saved on installations with swap space, or printed to the console on other platforms (e.g. ARM).

Viewing and Submitting a Crash Dump

After a panic or crash leading to a reboot, the operating system attempts to recover the contents of the crash dump while booting back up.

If the OS was able to read and process the contents of the crash dump, it converts the crash dump into a crash report. The GUI then displays a prompt on the dashboard to view the contents of the crash report.

Note: For privacy reasons, crash reports cannot be automatically submitted to Netgate for review.

Download the report or copy the contents from the GUI, and review the included data to ensure there is no private or identifiable information inside. If the data is OK, create a post on the Netgate Forum with the contents of the crash report. Attaching the crash report archive file(s) from the GUI is the best practice as that yields the most information in the easiest format to read.

Note: At a minimum, include the backtrace portion of `ddb.txt` from the crash dump archive. The backtrace is located after `db:0:kdb.enter.default> bt` in that file.

Crash Dump Format

The crash dumps are in FreeBSD textdump format and held in `/var/crash` after the OS recovers the data.

Crash dump archives are named `textdump.tar.<n>` where `<n>` starts at `0` and is incremented if older archives are still present in `/var/crash`.

Crash dump archives contain the following files:

`config.txt`
: The kernel configuration file

`ddb.txt`
: The output from the debugger scripts run during the crash dump

`msgbuf.txt`
: The contents of the kernel message buffer (`dmesg`)

`panic.txt`
: The panic message, if available. Typically, a more detailed version of this is found at the end of `msgbuf.txt`.

`version.txt`
: The kernel version string

Serial Console Crash Dump

On hardware with a serial console, connect to the serial console and record the scrollback buffer when a crash happens, or there may not be a way to retrieve the crash dump otherwise.

Install without Swap Space

If the installation does not contain any swap space, it may not be able to take a crash dump or automatically restart, and may stop at a `db>` prompt on the console. Capture the output there, and also the output of the `bt` command at that prompt, then manually reboot.

38.1.4 FreeBSD Issue Policy

The pfSense® team relies on the work provided by the FreeBSD project as the base operating system of pfSense software. On occasion, pfSense software users will run into problems with FreeBSD that are outside the scope with which Netgate can help. This page provides basic guidance on what to do in these scenarios.

Most frequently, these issues are driver or hardware-specific bugs. Netgate developers are most familiar with hardware sold by Netgate, and drivers for that hardware in FreeBSD. Thus, that is the only hardware for which Netgate can offer assistance.

When pfSense software users encounter driver problems on third party hardware, the best practice is to install a stock FreeBSD release on the hardware, replicate the problem, and report it to the appropriate FreeBSD list. This is likely the only way such problems will get resolved in future releases of FreeBSD. Alternatively, use different hardware known to work well with pfSense software, such as hardware from the Netgate Store.

If an installation of pfSense software encounters a kernel panic, Netgate developers may be able to help analyze the debugger backtrace. Post the crash dump contents on the Netgate Forum.

See also:

See: *Obtaining Panic Information for Developers*

The primary exception to this is for issues that affect many pfSense software users and/or Netgate hardware customers. Examples of this could be problems with PF, CARP, IPsec, or any number of other components. In these cases, Netgate will track and work to resolve issues of this nature. However, Netgate cannot coordinate testing and fixes for one-off issues and edge cases encountered by single users.

38.1.5 Requesting New pfSense Features

To request new pfSense® features first submit a feature request on pfSense Redmine.

If the feature request is approved, develop the functionality and submit a pull request to the applicable pfSense repositories.

38.1.6 System Patches Package

The System Patches package manages patches which change the behavior of pfSense® software. These patches may be bundled with the package, fetched from the official code repository, pasted in, or even uploaded from other sources.

This package makes it simple obtain official recommended security patches and bug fixes from Netgate between releases, as well as to test and deploy custom changes.

Installing the package

This package is available in the package repository:

- Navigate to **System > Packages**, **Available Packages** tab.
- Find **System Patches** in the list, or search for it.
- Click  **Install** at the end of its row
- Click  **Confirm** to confirm the action and complete the installation.

Once the installation finishes, Patches may be managed at **System > Patches**.

Recommended System Patches

The System Patches package is bundled with recommended patches for specific versions of pfSense software. This list is below the custom system patches area on **System > Patches**. The list of recommended patches is different for each version and typically includes security patches and important bug fixes. The list only displays patches relevant to the current version of pfSense software running on the device.

Note: Not every release has recommended patches! In most cases a brand-new release would not have any entries in the recommended patches list unless a significant issue was discovered after the release was published.

The list of recommended patches includes a brief description of the patch and what it fixes or changes. There are typically links to relevant Redmine issues, Security Advisories, or similar sources. There may also be special notes or instructions to follow for patch entries.

Managing Recommended Patches

Applying and reverting recommended patches works the same way as it does for manual patch entries, as described in *Managing Patch Entries*. There are a few differences, however. Options not relevant to recommended patches are not available for entries in the recommended patches area. This includes the "Fetch" action as well as the ability to reorder, edit, add, or delete entries.

In most cases administrators will want to click  **Apply All Recommended** which will do as it says and apply all the currently unapplied patches in the recommended patches list. Recommended patches may also be applied or reverted one at a time.

Warning: Occasionally a recommended patch also requires a manual action, such as restarting a service or rebooting the device. Read the description of each recommended patch entry carefully for such notes and follow the suggested steps after applying the patch.

Warning: There is no need to take any action with these patches before upgrading pfSense software. Newer releases will already contain the older patches, and the entries in the list will be different after the release.

Updating Recommended Patches

Recommended patches are bundled within the System Patches package itself, so updating the package will also update the list of available recommended patches.

> **Warning:** Take care when checking for package updates when there is a newer release of pfSense software available. Ensure the device is set to use its current version update branch under **System > Updates** rather than the branch for the new version.

Patch Settings

When creating or editing a Custom System Patch entry, the following settings are available:

Description
Text identifying the patch for reference.

URL/Commit ID
A commit ID for the pfSense CE software repository on GitHub, or the full URL to a patch file.

After saving the patch, use the **Fetch** button to download the patch content to the firewall.

Patch Contents
The contents of the patch in unified diff format.

When using a URL or commit ID, this should be blank when first saving but will contain the patch content after fetching.

Patch File Upload
A button to populate the **Patch Contents** by selecting a file on the client computer.

Path Strip Count
The number of path components to remove from the paths in patch metadata.

GitHub commit IDs and URLs should be count of `2` (default and fixed automatically on save). Patches from other sources will need to be set appropriately.

For example, if a path like `a/src/etc/inc/filter.inc` is in the patch header, the package should strip off the `a/src` so a strip count of `2` is needed. If it's deeper, such as `home/me/patches/etc/inc/filter.inc`, strip however many levels are necessary, which in this example would be `3`.

Base Directory
The package assumes a base directory of `/` for patches by default, but an alternate base may be applied if a patch does not supply a full path to the file it is patching (e.g. `/usr/local/www`).

Ignore Whitespace
Whether the patching process should ignore whitespace differences in the patch data.

Patches from GitHub should work with either whitespace setting, patches from other sources may need the option set to ignore whitespace, especially if they contain DOS line endings rather than UNIX or if the patch content lost tabs when copying and pasting.

Auto Apply
Whether the package will attempt to apply this patch on each boot of the firewall.

For patches which are included in future releases of pfSense software this is unnecessary as the appropriate fixes are included in the new release and need not be applied again. For manual custom changes this may be necessary to ensure these customizations are restored after upgrades.

The patches may be reordered in the list to arrange them so they apply in a specific order automatically, in case one patch depends on a previous patch.

Patch ID
When editing an existing patch, the GUI displays its unique ID in this field.

Managing Patch Entries

Manage patch entries at **System > Patches**.

The **Custom System Patches** list is for patches added manually by firewall administrators. The list has the following functions:

Select/Move
Selects entries to move or delete.

Clicking the  icon moves selected patches to this position, altering the order of patches. This may be relevant with auto-apply if one patch depends upon another.

Description
Text describing the patch, for reference.

Fetch
A button to download the patch content from its source, either a custom URL or a GitHub commit ID.

Apply
Attempt to apply this patch.

Revert
Attempt to revert this patch.

View
View the contents of the patch data.

Debug
Test the patch and interpret the results, this will display information about why a patch may not apply or restore cleanly. The output will include a detailed analysis of the results and can optionally display full detail of patch failures.

Auto Apply
A read only indication of whether this patch entry has the auto-apply option enabled.

Edit

The  icon edits this patch entry.

Delete

The  icon deletes this patch entry.

Add New Patch
Creates a new patch entry.

Delete Patches
Deletes all selected patch entries.

Note: The GUI does not display buttons unless they are relevant.

The lower section contains **Recommended System Patches** for the specific running version of pfSense software as described earlier in this document under *Recommended System Patches*. The controls in this section are limited as there is no need to edit the entries or alter the list.

Warning: There is typically no need to revert Custom or Recommended patch entries before or after upgrading pfSense software. Newer releases may contain the same fix as an older patch, which means the patch may appear to be applied after upgrading. Reverting such patches will **remove** the fix from the new release, bringing back the old bug. As such, the best course of action is to delete outdated Custom System Patch entries without reverting them.

Adding a Custom Patch

- Navigate to **System > Patches**
- **Read the text and warnings!**
- Click  **Add New Patch** under the **Custom System Patches** section
- Enter *Patch Settings* as described previously using one of the following styles:
  - Commit ID (e.g. `4573641589d50718b544b778cea864cfd725078a`) in the **URL/Commit ID** field
  - GitHub commit URL (e.g. `https://github.com/pfsense/pfsense/commit/4573641589d50718b544b778cea864cfd725078a`) in the **URL/Commit ID** field
  - GitHub Pull Request (PR) URL with '.diff' appended, such as `https://github.com/pfsense/pfsense/pull/XXXX.diff` where XXXX is the PR number

    Set Path Strip = `2` if it does not adjust automatically
  - Full URL to a patch from another source (e.g. `https://redmine.pfsense.org/attachments/594/0001-Add-support-for-aliases-in-DNS-Forwarder-fixes-2410.patch`) in the **URL/Commit ID** field
  - Leave **URL/Commit ID** blank and paste the contents of a patch into **Patch Contents** text area or upload a patch file
- Click **Save**

Applying/Reverting a patch

If a URL or commit ID was entered for a patch, the entry in the patch list will have a **Fetch** button.

Click **Fetch** and firewall will retrieve the patch content. This **does not** apply the patch.

To apply the patch, click **Apply**. The package will then apply the patch. The available link for the patch will then change to say **Revert** instead.

To revert, click **Revert**.

See also:

Managing Patch Entries

Troubleshooting

- It is normal for an already-applied patch to show only a revert button. Similarly, if an older patch is present on a newer release which includes the fix, it will appear as already applied.

  Warning: Do not revert these patches after upgrading! – Doing so will undo the fix and cause problems on the new release.

- If one patch relies upon another patch being applied first, their usual actions may not appear unless taken in the proper order. For example, patch 2 may not show an Apply button until after patch 1 is applied. Likewise for reverting.
- Click  **Re-Fetch** for remote patches to make sure the package has a clean copy of the patch content.
- Click **Debug** to run a test and then click **Detail** next to either the apply or revert line to get the full patch output
- If the above test output mentions **No file to patch**, double-check the **Path Strip Count** and/or the **Base Directory**.
- If every part of a patch fails, try toggling **Ignore Whitespace**.

Known issues

See also:

The pfSense bug tracker contains a list of known issues with this package.

Package Assistance

Netgate TAC can provide assistance for this package to customers with an active TAC subscription.

38.1.7 Using the PHP Shell

Using the PHP developer shell on pfSense® software allows manipulation of the firewall configuration directly without using the GUI. Using this mechanism also allows rapid deployment of pfSense software and/or the setup of exotic configurations.

The shell can be started from console menu option `12` or from the CLI by executing `pfSsh.php`.

Example Session

The following shows an example session, with the text coming from the `help` command in the PHP shell.

Follow each line or group of lines to run with `exec;`:

```
*** Welcome to pfSense ***

 WAN (wan)       -> vmx0       -> v4/DHCP4: 198.51.100.3/24
                                  v6/DHCP6: 2001:db8::ffff:22d6/128
 LAN (lan)       -> vmx1       -> v4: 10.3.0.1/24
                                  v6/t6: 2001:db8:1:eee0:20c:29ff:fe45:260/60
```

(continues on next page)

(continued from previous page)

```
 0) Logout (SSH only)                  9) pfTop
 1) Assign Interfaces                 10) Filter Logs
 2) Set interface(s) IP address       11) Restart GUI
 3) Reset admin account and password  12) PHP shell + pfSense tools
 4) Reset to factory defaults         13) Update from console
 5) Reboot system                     14) Disable Secure Shell (sshd)
 6) Halt system                       15) Restore recent configuration
 7) Ping host                         16) Restart PHP-FPM
 8) Shell

Enter an option: 12

Starting the pfSense developer shell....

Welcome to the pfSense developer shell

Type "help" to show common usage scenarios.

Available playback commands:
     changepassword disablecarp disabledhcpd disablerefererchеck enableallowallwan␣
↪enablecarp
     enablesshd externalconfiglocator generateguicert gitsync installpkg listpkg␣
↪removepkgconfig
     removeshaper restartdhcpd restartipsec svc uninstallpkg

pfSense shell: help

     Enter a series of commands and then execute the set with "exec".

     For example:
     echo "foo"; // php command
     echo "foo2"; // php command
     ! echo "heh" # shell command
     exec

     Example commands:

     record <recordingfilename>
     stoprecording
     showrecordings

     parse_config(true);  # reloads the $config array

     $temp = print_r($config, true);
     more($temp);

     /* to output a configuration array */
     print_r($config);

     /* to output the interfaces configuration portion of config.xml */
     print_r($config['interfaces']);
```

(continues on next page)

(continued from previous page)

```
    /* to output the dhcp server configuration */
    print_r($config['dhcpd']);

    /* to exit the  developer shell */
    exit

    /* to output possible wireless modes for an interface */
    print_r(get_wireless_modes(\"ath0\"));

    /* to enable SSH */
    $config['system']['enablesshd'] = true;

    /* change OPTX to the OPT interface name such as BACKHAUL */
    $config['interfaces']['optx']['wireless']['standard'] = "11a";
    $config['interfaces']['optx']['wireless']['mode'] = "hostap";
    $config['interfaces']['optx']['wireless']['channel'] = "6";

    /* to enable dhcp server for an optx interface */
    $config['dhcpd']['optx']['enable'] = true;
    $config['dhcpd']['optx']['range']['from'] = "192.168.31.100";
    $config['dhcpd']['optx']['range']['to'] = "192.168.31.150";

    /* to disable the firewall filter */
    $config['system']['disablefilter'] = true;

    /* to enable an interface and configure it as a DHCP client */
    $config['interfaces']['optx']['disabled'] = false;
    $config['interfaces']['optx']['ipaddr'] = "dhcp";

    /* to enable an interface and set a static IPv4 address */
    $config['interfaces']['wan']['enable'] = true;
    $config['interfaces']['wan']['ipaddr'] = "192.168.100.1";
    $config['interfaces']['wan']['subnet'] = "24";

    /* to save out the new configuration (config.xml) */
    write_config();

    /* to reboot the system after saving */
    system_reboot_sync();
```

Playback Scripts

There are several pre-defined playback scripts for the PHP Shell that automate simple tasks or enable access to the GUI.

These scripts are run from within the PHP shell like so:

```
pfSense shell: playback scriptname
```

Warning: The PHP shell is not designed to run multiple playback scripts in a single session. After executing a playback script, exit the PHP shell and start it again. Otherwise, subsequent scripts may fail to run. Alternately,

run the playback scripts using the CLI method.

The scripts can also run from the command line:

```
# pfSsh.php playback scriptname
```

Warning: This must be run from an actual shell prompt over SSH or at the console. Do not attempt to run this through the GUI.

changepassword

This script changes the password for a user, and also prompts to reset the account properties if it is disabled or expired.

cryptconfig

This script can encrypt or decrypt configuration backups using the same method as doing so in the GUI when creating or restoring backups.

In either case, the script will prompt for the encryption passphrase. When decrypting, this passphrase must exactly match the passphrase used to encrypt the configuration file.

Usage:

```
# pfSsh.php playback cryptconfig <action> <input filename> <output filename>
```

To encrypt a configuration file:

```
# pfSsh.php playback crypt encrypt /conf/config.xml /root/config-backup.xml
```

To decrypt a configuration file:

```
# pfSsh.php playback crypt decrypt /root/config-backup.xml /root/config.xml
```

disablecarp / enablecarp

These scripts disable and enable CARP high availability functions, and will deactivate CARP type Virtual IP addresses.

This action **does not** persist across reboots.

disablecarpmaint / enablecarpmaint

These scripts disable and enable CARP maintenance mode, which leaves CARP active but demotes this unit, so the other node can assume control.

This maintenance mode **will** persist across reboots.

disabledhcpd

This script removes all DHCP configuration from the firewall, effectively disabling the DHCP service and completely removing all of its settings.

disablerefercheck

This script disables the `HTTP_REFERER` check mentioned in *Browser HTTP_REFERER enforcement*. This can help gain access to the GUI if a browser session is triggering this protection.

enableallowallwan

This script adds an allow all rule for IPv4 and IPv6 to the WAN interface.

> **Warning:** Be extremely careful with this option, it is meant to be a **temporary** measure to gain access to services on the WAN interface of the firewall in situations where the LAN is not usable. Once proper access rules are put in place, remove the rules added by this script.

enablesshd

This script enables the SSH daemon, the same as the console menu option or GUI option.

externalconfiglocator

This script looks for a `config.xml` file on an external device, such as a USB thumb drive, and will move it in place for use by the firewall.

gatewaystatus

This script prints the current gateway status and statistics. It also accepts an optional parameter `brief` which prints only the gateway name and status, omitting the addresses and statistical data.

generateguicert

This script creates a new self-signed certificate for the firewall and activates it for use in the GUI. This is useful in cases where the previous certificate is invalid or otherwise not usable. It fills in the certificate details using the firewall hostname and other custom information.

gitsync

This complex script synchronizes the PHP and other script sources with files from the pfSense CE software GitHub repository. It is most useful on CE development snapshots to pick up changes from more recent commits.

> **Warning:** This script can be dangerous to use in other circumstances. Only use this under the direction of a knowledgeable developer or Netgate TAC representative.

> **Warning:** This script is not currently compatible with pfSense Plus software.

If the script is run without any parameters it prints a help message outlining its use. More information can be found at *Using gitsync to Update pfSense® Software Between Snapshots*.

installpkg / listpkg / uninstallpkg

These scripts interface with the package system and operate identically to the GUI package manager functions. These are primarily used for debugging package issues, comparing information in `config.xml` compared to the package database.

pfanchordrill

This script recursively searches through PF anchors and prints any NAT or firewall rules it finds. This can help track down unexpected behavior in areas such as UPnP or Captive Portal which rely on rules in anchors that are not otherwise visible in the GUI.

pftabledrill

This script prints the contents of all PF tables, which contain addresses used in firewall aliases as well as built-in system tables for features such as bogon network blocking, snort, and GUI/SSH lockout. This script is useful for checking if a specific IP address is found in any table, rather than searching individually.

removepkgconfig

This script removes all traces of package configuration data from the running `config.xml`. This can be useful if a package has corrupted settings or has otherwise left the packages in an inconsistent state.

removeshaper

This script removes ALTQ traffic shaper settings, which can be useful if the shaper configuration is preventing rules from loading or is otherwise incorrect and preventing proper operation of the firewall.

resetwebgui

This script resets the GUI settings for widgets, dashboard columns, the theme, and other GUI-related settings. It can return the GUI, particularly the dashboard, to a stable state if it is not functioning properly.

restartallwan

This script disables and re-enables each WAN-type interface, which reapplies the interface configuration.

restartdhcpd

This script stops and restarts the DHCP daemon.

restartipsec

This script rewrites and reloads the IPsec configuration for strongSwan.

svc

This script controls the services running on the firewall, similar to interacting with services at **Status > Services**.

The general form of the command is:

```
# pfSsh.php playback svc <action> <service name> [service-specific options]
```

The **action** can be `stop`, `start`, or `restart`.

The **service name** is the name of the services as found under **Status > Services**. If the name includes a space, enclose the name in quotes.

The **service-specific options** vary depending on the service, they are used to uniquely identify services with multiple instances, such as OpenVPN or Captive Portal entries.

Examples:

Stop the `miniupnpd` service:

```
# pfSsh.php playback svc stop miniupnpd
```

Restart OpenVPN client with ID 2:

```
# pfSsh.php playback svc restart openvpn client 2
```

Start the Captive Portal process for zone “MyZone”:

```
# pfSsh.php playback svc start captiveportal MyZone
```

upgradeconfig

Runs the current configuration, or optionally an arbitrary configuration file, through the configuration upgrade process.

The general form of the command is:

```
# pfSsh.php playback upgradeconfig [<input filename> <output filename>]
```

To force the current running configuration through the upgrade process:

```
# pfSsh.php playback upgradeconfig
```

To run the configuration upgrade process on a different configuration file:

```
# pfSsh.php playback upgradeconfig /root/oldconfig.xml /root/newconfig.xml
```

Warning: Some configuration upgrade steps may contain a directive to write the configuration which may unintentionally replace the current running configuration with the copy being upgraded at that step. After running this script on an arbitrary file, visit the configuration history (*Configuration History*) and roll back to the correct configuration if it was replaced. The upgrade process does not apply settings, so no additional action should be necessary.

Recording and Playback

Aside from the predefined scripts, users may create their own playback scripts by recording sessions.

This can easily automate a number of commands to repeat them later, saving time and effort.

Recording a session

```
# pfSsh.php
[...]

pfSense shell: record resetrrd
Recording of resetrrd started.
pfSense shell: require_once("filter.inc");
pfSense shell: require("shaper.inc");
pfSense shell: require_once("rrd.inc");
pfSense shell: ! rm /var/db/rrd/*.rrd
pfSense shell: enable_rrd_graphing();
pfSense shell: setup_gateways_monitor();
pfSense shell: stoprecording
Recording stopped.
pfSense shell: exit
```

Playing back a session

```
# pfSsh.php
[...]

pfSense shell: playback resetrrd

Playback of file resetrrd started.

pfSense shell: exit
```

Warning: The PHP shell is not designed to run multiple playback scripts in a single session. After executing a playback script, exit the PHP shell and start it again. Otherwise, subsequent scripts may fail to run. Alternately, run the playback scripts using the CLI method.

Sessions can be played back directly from the command line as well:

```
# pfSsh.php playback resetrrd
```

38.2 pfSense Software Development

38.2.1 Getting Started with pfSense® Software Development

There is no single specific starting point for joining the pfSense® software development effort, but the following items are helpful in getting started:

- Review the *Developer Style Guide*.
- Become familiar with the pfSense CE software git repositories and GitHub in general.

  Current repositories used for developing pfSense software and its dependencies include:

  pfSense
  : The main source repository for pfSense CE software, containing the GUI code, builder code, and related scripts.

  FreeBSD-src
  : OS source code used to build pfSense CE software.

  FreeBSD-ports
  : Build information for software used in pfSense CE software, and code for custom programs, daemons, modules, and packages.

- Review the list of open bug reports and other issues.
- Submit changes as *pull requests on GitHub*.

Netgate developers will review submissions, offer feedback, and merge the changes if they are acceptable.

38.2.2 PHP 8.x and Later Development

PHP 7.x is at its end of life, and so development is moving to PHP 8.x on future releases of pfSense software, such as pfSense Plus software version 23.01 and pfSense CE software version 2.7.x.

Migrating to PHP 8.1 introduces some backward incompatible changes from 7.x in how some expressions are evaluated.

- *Configuration Access*
  - *Reading Configuration Values*
  - *Testing if Settings are Enabled*
  - *Writing Configuration Values*
  - *Deleting Configuration Values*
- *Array Access Functions*
- *String to Number Comparison*
- *Practical Examples*
  - *Iterating over Arrays of Items in the Configuration*
  - *Replacing isset() to Determine if an Item is Enabled*
  - *Accessing Items from Variable Paths*
  - *Default Values*

Configuration Access

One of the more pervasive idioms in the pfSense code base and its packages that is now problematic is the method by which code historically traverses the configuration tree. The problematic method indexes the configuration data into multiple levels of nested associative arrays, e.g.:

```
$foo = $config['section']['subsection']['item'];
```

The configuration array tree is not fully populated with all possible keys, therefore any of the keys in the expression above may not exist in the array being indexed. The evaluation of an expression indexing an associative array with a key that it does not contain results in an empty string. Expressions indexing a string are allowed in PHP 7.4, and in the case that the key is a string the expression evaluates to an empty string again. However, in PHP 8.1 these expressions now raise errors and terminate the evaluation of PHP code. To correct these errors and to progress toward decoupling the implementation of the runtime configuration store from the rest of the code base, the development tree now includes utility functions to access the `$config` variable instead of accessing it directly. These functions are defined in `/etc/inc/config.lib.inc`.

Reading Configuration Values

The following function should be used to read values from the configuration:

```
config_get_path(string $path, $default = null)
```

Given a path with the separator /, look for an item by indexing into the configuration tree by the identifiers in the path.

Returns `$default` if the item or any intermediary in the path cannot be found.

Examples

Basic example accessing one value nested inside a section:

PHP 7.x Style:

```
$foo = $config['section']['item'];
```

PHP 8.x Style:

```
$foo = config_get_path('section/item');
```

More complicated example accessing a value inside a subsection multiple levels deep:

PHP 7.x Style:

```
init_config_arr('section', 'subsection');
if ($config['section']['subsection']['item']) {
    $foo = $config['section']['subsection']['item'];
} else {
    $foo = "Not Found";
}
```

PHP 8.x Style:

```
$foo = config_get_path('section/subsection/item', "Not Found");
```

Testing if Settings are Enabled

A common way to mark options and features as enabled or disabled in the configuration is via "presence" style variables. Where if the value is present in the configuration, the option is enabled. If the value is missing, it's disabled.

There is a new shortcut for testing these options:

```
config_path_enabled($path, $enable_key = "enable")
```

This function determines if `$enable_key` is a key into the array at path `$path`, and the value is non-null. This is used to obtain the same semantics as checking if the value is present in the configuration, but in a safe manner.

Note: Some sections of the configuration use different semantics for indicating if a feature is enabled which `config_path_enabled()` does not interpret at this time. Only use `config_path_enabled()` to replace `isset()` style expressions.

Examples

This first example is for an option which is enabled or disabled based on the presence of a configuration element named default "enable":

PHP 7.x Style:

```
$bar_enabled = isset($config['foo']['bar']['enable']);
```

PHP 8.x Style:

```
$bar_enabled = config_path_enabled('foo/bar');
```

This next example tests whether the element "baz" is present in the configuration at the same level of the previous example.

PHP 7.x Style:

```
$baz_enabled = isset($config['foo']['bar']['baz']);
```

PHP 8.x Style:

```
$baz_enabled = config_path_enabled('foo/bar', 'baz');
```

Writing Configuration Values

This function sets an item at a given the configuration path:

```
config_set_path(string $path, $value, $default = null)
```

If any intermediate section in the path cannot be found, or is an empty value, the function creates an array for it. If any intermediary in the path is unexpectedly a scalar value, the function returns the `$default` value to indicate an error.

Examples

This basic example sets a value in the configuration and then writes the changes:

PHP 7.x Style:

```
$config['foo']['bar']['item'] = 'newvalue';
write_config('Update settings');
```

PHP 8.x Style:

```
config_set_path('foo/bar/item', "newvalue");
write_config('Update settings');
```

This slightly more complex example reads a value from the configuration, updates the option with a new value, then stores the results.

PHP 7.x Style:

```
$bar = &$config['foo']['bar'];
/* ... */
$bar['item'] = 'newvalue';
/* ... */
write_config('Update settings');
```

PHP 8.x Style:

```
$bar = config_get_path('foo/bar');
/* ... */
$bar['item'] = 'value';
/* ... */
config_set_path('foo/bar', $bar);
write_config('Update settings');
```

Deleting Configuration Values

To unset (delete) a value from the configuration, use this function:

```
config_del_path(string $path)
```

This function completely removes a path from the configuration tree by running `unset()` on the array element, if it exists.

This obtains the same result as `unset()`, but in a safe manner.

Examples

PHP 7.x Style:

```
unset($config['foo']['bar']['item']);
```

PHP 8.x Style:

```
config_del_path('foo/bar/item');
```

Array Access Functions

The configuration-based functions utilize more general array functions which can be used directly for similar tasks on any array.

```
array_get_path(array &$arr, string $path, $default = null)
array_set_path(array &$arr, string $path, $value, $default = null)
array_path_enabled(array &$arr, string $path, $enable_key = “enable”)
array_del_path(array &$arr, string $path)
```

String to Number Comparison

In the past, PHP would return `true` when comparing `0` to an empty string (`''`) but now it returns `false`. If the code in question must test against an empty string to know if a value is usable, do not cast the value to `int` or take other equivalent actions. Along the same lines, if the value must be compared numerically, cast it to `int` after first checking if it's an empty or otherwise usable value.

Note: While the behavior of comparing `0` to an empty string has changed, the result of `empty(0)` and `empty('0')` are still `true` on PHP 8.x.

Before:

```
$varusersamountoftime = ($users['varusersamountoftime'] ?: '');
$varusersamountoftime = (int) $varusersamountoftime * 60;
/* ... */
if ($varusersamountoftime != '') {
    /* ... */
}
```

In this example the value of the variable is always cast to `int` which means that an empty value is changed to `0`. On PHP 7.x, the later `if` would evaluate to `false` if the value was `0` but on PHP 8.x, the test will evaluate to `true`.

After:

```
if ($users['varusersamountoftime']) {
    $varusersamountoftime = (int) $users['varusersamountoftime'] * 60;
} else {
    $varusersamountoftime = '';
}
/* ... */
if ($varusersamountoftime != '') {
    /* ... */
}
```

This code will evaluate the same on both PHP 7.x and 8.x because it only changes the value to `int` when the variable is set and has a non-zero value, and it is an empty string otherwise.

Practical Examples

Iterating over Arrays of Items in the Configuration

Previously, one or more checks would have to be made to see if the array in the configuration is present, checking the value with `isset()` or `is_array()`. With the config interface, a suitable default value of an empty array can be specified in a call to `config_get_path()`, and the result can be directly iterated. The resulting code is more concise.

This example iterates over aliases in the configuration:

PHP 7.x Style:

```
if (isset($config['aliases']) &&
    is_array($config['aliases']) &&
    isset($config['aliases']['alias']) &&
    is_array($config['aliases']['alias'])) {
```

(continues on next page)

(continued from previous page)

```
    foreach ($config['aliases']['alias'] as $aliased) {
        if ($aliased['name'] == $alias_name) {
            return filter_generate_nested_alias($aliased['name']);
        }
    }
}
```

The above example needs multiple safety checks to ensure that each level exists and is an array before it attempts to iterate over the array. Some of this can be bypassed by running a `init_config_arr()` or similar, but it still requires wrapping it in a test before iterating.

Contrast that with this much simpler example using the new utility function:

PHP 8.x Style:

```
foreach (config_get_path('aliases/alias', []) as $aliased) {
    if ($aliased['name'] == $alias_name) {
        return filter_generate_nested_alias($aliased['name']);
    }
}
```

This next example demonstrates a common method used to check all OpenVPN server and client instances:

PHP 7.x Style:

```
init_config_arr(array('openvpn'));
foreach (array('server', 'client') as $mode) {
    if (is_array($config['openvpn']["openvpn-{$mode}"])) {
        foreach ($config['openvpn']["openvpn-{$mode}"] as $id => $setting) {
            /* ... */
        }
    }
}
```

This becomes much simpler and does not require the extra initialization or tests:

PHP 8.x Style:

```
foreach (array('server', 'client') as $mode) {
    foreach (config_get_path("openvpn/openvpn-{$mode}", []) as $id => $setting) {
        /* ... */
    }
}
```

This example iterates over all installed packages. This requires extra checks because the user may not have any packages installed, or may even not have a section in the configuration with package settings.

PHP 7.x Style:

```
if (is_array($config['installedpackages']) &&
    is_array($config['installedpackages']['package'])) {
    foreach ($config['installedpackages']['package'] as $pkg) {
        if ($pkg['name'] == $package_name) {
            return $pkg['descr'];
        }
```

(continues on next page)

(continued from previous page)

```
    }
}
```

PHP 8.x Style:

```
foreach (config_get_path('installedpackages/package', []) as $pkg) {
   if ($pkg['name'] == $package_name) {
      return $pkg['descr'];
   }
}
```

Replacing isset() to Determine if an Item is Enabled

As with other configuration accesses, using `config_path_enabled()` is safe if any part of the path does not exist, and will return `false` if that is the case. If the array key for the element that indicates something is enabled is not the default `enable`, the name can be passed as another parameter.

PHP 7.x Style:

```
$assignedif = convert_real_interface_to_friendly_interface_name($vlanif);
if ($assignedif) {
    if (isset($config['interfaces'][$assignedif]['enable'])) {
        interface_configure($assignedif, true);
    }
}
```

PHP 8.x Style:

```
$assignedif = convert_real_interface_to_friendly_interface_name($vlanif);
if ($assignedif) {
    if (config_path_enabled("interfaces/{$assignedif}")) {
        interface_configure($assignedif, true);
    }
}
```

This example checks if an option is enabled the same way, but using a specific key name:

PHP 7.x Style:

```
if (!isset($config['system']['ipv6allow'])) {
    /* ... */
}
```

PHP 8.x Style:

```
if (!config_path_enabled('system','ipv6allow')) {
    /* ... */
}
```

Accessing Items from Variable Paths

Before, extra care needed to be taken to initialize multi-level arrays and to check before accessing certain areas, for example:

```
init_config_arr(array('captiveportal'));
$a_cp = &$config['captiveportal'];
/* ... */
if ($a_cp[$cpzone]) {
    /* ... */
    $pconfig['certref'] = $a_cp[$cpzone]['certref'];
    /* ... */
}
```

Now this can be done safely without the extra steps:

```
$pconfig['certref'] = config_get_path("captiveportal/{$cpzone}/certref");
```

Note: Note the use of double quotes in the path to allow variable substitution.

Default Values

One of the primary benefits of this new style is the ability to easily accommodate default values without a lot of extra logic. There are examples of this in previous sections above where a default array is returned (`[]`) to ensure that a returned value is always an array.

In this example, a value is populated either from the configuration or using a globally defined default:

PHP 7.x Style:

```
init_config_arr('syslog');
$syslogcfg = $config['syslog'];
$log_size = isset($syslogcfg['logfilesize']) ? $syslogcfg['logfilesize'] : $g['default_
↪log_size'];
```

PHP 8.x Style:

```
$log_size = config_get_path('syslog/logfilesize', $g['default_log_size']);
```

38.2.3 Referencing Tickets in Commit Messages

By placing a special keyword in a commit message the issue tracking system (pfSense Redmine) can associate a commit with a specific ticket automatically, creating a link in the ticket to the relevant commits.

When using these keywords immediately follow them by a `#` and then the ticket number, such as `Ticket #1234`. They are not case-sensitive.

The following keywords will reference a ticket but take no action on the ticket status:

- `refs`, `references`, `IssueID`, `ticket`, `bug`, `feature`, `todo`, `redmine`

The following keywords will not only reference the ticket, but automatically move the ticket to a feedback state:

- `fix`, `fixes`, `fixed`, `close`, `closes`, `closed`, `resolve`, `resolves`, `resolved`, `implement`, `implements`, `implemented`, `finish`, `finishes`, `finished`

Keep this in mind when submitting changes in GitHub pull requests for existing issues.

38.2.4 Submitting a Pull Request via GitHub

Submitting a Pull Request (PR) via GitHub is the fastest and best way to contribute source code changes to the pfSense® CE project.

Using a PR allows developers to easily review and comment on changes, allows easy testing via patches from the *System Patches* Package, and allows the changes to be easily merged into the project.

Creating a PR is a relatively straightforward process, but there are a few guidelines and suggestions to follow when submitting contributions:

- Ensure the bug or feature has an entry on the pfSense Redmine

  Create a new entry if one does not exist.

  The only exceptions to this are for very minor typo/wording fixes.

- Read through and follow the *Developer Style Guide*
- Read through the GitHub documentation on pull requests
- Create a fork of the correct repository (e.g. pfSense/pfSense for the base system, or pfSense/FreeBSD-ports for packages)
- Always submit pull/merge requests to the `master` branch of the `pfSense/pfSense` repository or the `devel` branch of `pfSense/FreeBSD-ports`.

  The exception to this is when making PRs for multiple branches due syntax or other differences which require variations that prevent a regular cherry-pick or similar action from copying the commit to other branches

- When making commits to the fork, *reference the relevant Redmine entry*
- When ready to submit the changes, create a pull request from the fork to the appropriate branch
  - The PR title should be a short summary of the changes and include a reference to the relevant Redmine issue number(s)
  - The PR description should include a longer explanation of the proposed changes and link to the relevant Redmine issue number(s)
- After submitting the pull request, add a link to the PR on the related Redmine issue to cross-reference the entries.

38.2.5 Checking the Current FreeBSD Version

Versions of pfSense software and FreeBSD contains a table of the versions used in various releases of pfSense® software and FreeBSD.

Alternately, inspect the version on a running firewall manually using one of the following methods.

Dashboard

From the Dashboard in the GUI, look at the **System Information** widget. The widget displays the FreeBSD version under the **Version** section.

Command Line

In the *SSH console* or *Execute Shell Command* field in the GUI, run the following command:

```
uname -mrs
```

38.2.6 Creating Dashboard Widgets

Getting Started

Creating widgets is simple. First, create the HTML code for the widget, save it to a file named `widget_name.widget.php`, and put it into the `/usr/local/www/widgets/widgets` directory on the firewall.

Do not include any `<body>`, `<html>`, or `$pgtitle` definitions, etc. Only include basic HTML code for what needs to be displayed. Examine the current widget source code for examples.

The file must be named in the `my_name.widget.php` format. No spaces are allowed in the filename. The name the GUI displays is the name of the file. For example the Traffic Graphs widget file is named `traffic_graphs.widget.php`.

And that's it! The dashboard handles the rest of the widget behavior (e.g. buttons, border, dragging, sequence, etc.) automatically.

To include custom code into the widget upon rendering:

PHP
: Create a file named `widget_name.inc` and place the `.inc` file in the `/usr/local/www/widgets/include` directory.

JavaScript
: Create a file named `widget_name.js` and place the `.js` file in the `/usr/local/www/widgets/javascript` directory.

Saving Data

Saving configuration data for a widget requires more work and is beyond the scope of this document. Look at the source code for existing widgets which have settings and follow their example, including:

- The form for the settings and all related fields, save button, etc.
- Input validation
- Storing settings by writing the configuration
- Redirect the user back to the dashboard

Customizing the Title and Linking to page

By default, the name of the widget file is what is shown on the Widget in the dashboard. This title can be changed and also have a link to another page inserted.

To configure the widget to use a certain name other than the name of the file, create a `.inc` file with the same base name as the widget and place that file into the `/usr/local/www/widgets/include` directory.

For example if a widget is named `abc.widget.php`, the include file must be `abc.inc`.

In this `.inc` file use the following code:

```
<?php //set variable for custom title
$abc_title = "A B C custom";
$abc_title_link = "abc.php";
?>
```

Use the widget name in the variable names as shown in the example.

An example of these values can be taken from the `interfaces.inc` file for the Interfaces widget:

```
<?php [...]
$interfaces_title = gettext("Interfaces");
$interfaces_title_link = "status_interfaces.php";
?>
```

38.2.7 Enabling Additional PHP Modules

In certain cases, such as packages or customized Captive Portal code, additional PHP extensions may be required that are not enabled.

There are a handful of PHP extensions available in the package repository which are not included in the base distribution of pfSense® software.

The extensions included and activated vary by pfSense software version, look in `/usr/local/lib/php/` to see which extensions are present, and check the output of `php -m` to see what is enabled on a firewall already.

Examples of additional extensions available to install may include:

- `mysqli`
- `pdo_mysql`
- `pgsql`

To activate one of these, install the appropriate version-specific package from the command line, for example:

```
# pkg search php
<locate the desired package>
# pkg install -y php74-mysqli
# /etc/rc.php-fpm_restart; /etc/rc.restart_webgui
# php -m | grep mysqli
mysqli
```

A reboot would also fully activate the module, but should not be necessary.

38.2.8 Executing Commands at Boot

There are three primary options for executing custom commands at boot time: `shellcmd`, `earlyshellcmd`, and shell scripts.

The Shellcmd package can manage the `shellcmd` and `earlyshellcmd` tags in the GUI, so `config.xml` values need not be edited by hand.

At boot time the firewall executes the `earlyshellcmd` entries first and the `shellcmd` entries much later in the boot process. Shell scripts are executed at the very end of the boot process when initializing packages.

The `shellcmd` and `earlyshellcmd` options are preferable as they are contained within in the configuration file. As such they do not typically require additional modifications should the storage medium be replaced and reinstalled, or if the configuration is restored to a different piece of hardware.

shellcmd option

The hidden `config.xml` option <shellcmd> makes the firewall run a command towards the end of the boot process.

To add a `shellcmd` to a configuration, either use the `shellcmd` package or edit `config.xml` by hand (*XML Configuration File*).

To edit the `config.xml`:

- Back it up via **Diagnostics > Backup/restore**
- Open the XML backup file in a text editor that properly handles UNIX line endings.
- Add a new line above the </system> line such as the following:

```
<shellcmd>mycommand -a -b -c 123</shellcmd>
```

  Where `mycommand -a -b -c 123` is the command to run.

- Save the changes to the configuration
- Restore the modified configuration

Multiple lines may be added to execute multiple commands.

earlyshellcmd option

The hidden `config.xml` option <earlyshellcmd> makes the firewall run a command at the beginning of the boot process.

Note: In most cases <shellcmd> is more appropriate, though this may be necessary in some circumstances.

The process to add an <earlyshellcmd> tag is the same as <shellcmd>. Either use the `shellcmd` package or edit it in by hand.

This should result in a tag such as the following in the configuration:

```
<earlyshellcmd>mycommand -a -b -c 123</earlyshellcmd>
```

Multiple <earlyshellcmd> lines can be present to execute multiple commands.

Shell script option

Any shell script can be placed in the `/usr/local/etc/rc.d/` directory.

The filename must end in `.sh`, and it must be marked as executable (`chmod +x myscript.sh`).

The firewall will execute every shell script ending in `.sh` in this directory at boot time and also during certain system events (e.g. interface link changes, IP address changes, and gateway events).

38.2.9 Using a Debug Kernel

pfSense® software has debug symbols for the kernel and modules that can be added to aid in debugging. To install the debug symbols, use the following command:

```
pkg install pfSense-kernel-debug
```

38.2.10 Using gitsync to Update pfSense® Software Between Snapshots

Most often upgrading to a new development snapshot is the best way to get updated code when tracking a development version of pfSense® software (alpha, beta, RC, etc.).

However, since new snapshots are only built once per day in most cases, one can often get by with pulling new code from the `pfSense` CE Git repository instead of reloading a whole new snapshot. This is useful for testing code changes committed since the last snapshot. This process is known as a `gitsync`.

Warning: This process is not compatible with pfSense Plus software.

Warning: A `gitsync` only synchronizes PHP changes without any binary changes. At times, PHP changes require associated binary changes that only come from an upgrade using a snapshot or other upgrade image.

Unless development is followed closely and the ramifications of all changes are understood, or unless breakage is not a concern, do not use this!

For most users, this action should only be taken if directed to do so by a developer.

This only works with code or files that are not compiled. For example: PHP Code, configuration files, GUI pages, etc.

There are two ways to perform a `gitsync`, both perform the same function:

Method 1:

From the console menu, press option `12` to start a developer shell, then type:

```
> playback gitsync master
```

Method 2:

From a normal shell (console menu option 8), type the following command:

```
# pfSsh.php playback gitsync master
```

Warning: This must be run from an actual shell prompt over SSH or at the console. Do not attempt to run this through the GUI.

The `master` part of the command tells the gitsync process to grab the code for the `master` branch, a.k.a. HEAD, main, etc. That can be replaced with a version-specific branch such as `RELENG_x_y_z`.

For example, on 2.8.0-RELEASE, to sync post-release code changes, use:

```
# pfSsh.php playback gitsync RELENG_2_8_0
```

Troubleshooting

The `gitsync` script attempts to install `git` automatically. However, if an error occurs the `git` package may need to be added manually. This can be done from a shell prompt:

```
# pkg install git
```

Git Protocol changed/URL Moved

Occasionally something on GitHub changes and renders an old clone of the repository broken. The most recent change was a change to disable certain protocols in favor of only using HTTPS. In this case it's necessary to delete the old clone so that `gitsync` can create a fresh copy.

```
# rm -rf /root/pfsense/
# pfSsh.php playback gitsync master
```

38.2.11 Development Branch Names

During the development cycle of pfSense® software work happens on specific branches for specific purposes.

HEAD (master)

HEAD, also known as -HEAD, `master`, `plus-master`, `devel`, `plus-devel`, or `main`, refers to the latest bleeding edge development version of pfSense software, where all new features are first added. The specific branch name varies depending on the repository.

When a release nears, this branch of a given repository is branched to create a RELENG (Release Engineering) branch. This follows the FreeBSD project's development model.

RELENG_x_y_z

A branch with a complete version number, such as `RELENG_x_y_z` (CE) or `plus-RELENG_yy_mm` (Plus) is a maintenance branch for a specific release, `x.y.z` (CE) or `yy.mm` (Plus). This may be updated after a release and eventually become the next patch release, if necessary.

Others

For links to other branches, past and present, see *Versions of pfSense software and FreeBSD*.

See also:

- *Developer Style Guide*

38.3 pfSense Package Development

38.3.1 Developing Packages

Developers familiar with FreeBSD, `pkg`, and PHP, will find it fairly straightforward to create packages for pfSense® software. End users and organizations can benefit from developing a package that does not exist.

To submit a new package, create a *pull request* on GitHub for the pfSense software FreeBSD-Ports Repository so Netgate can evaluate the work for inclusion into the package repository for access by all users.

When developing packages, always target the latest development version of pfSense software first.

pfSense Software Package System

On pfSense software, every pfSense software package is also a FreeBSD port. These are installed and managed via `pkg`, even when using the GUI to add or remove packages. Binary packages from FreeBSD are added as dependencies of the pfSense software package, so they are installed automatically as well, along with any of their own required dependencies.

The basic idea is to make packages for pfSense software similar to FreeBSD packages, but with customization. One way this is achieved is by adding metadata about packages to the firewall configuration when it is installed, and also creating the configuration screen of an application using XML. pfSense software provides an optional framework to create the web interface and to store it in the XML configuration file of the firewall. The package writer is expected to convert the data from XML to the native format of the application.

Package System

pfSense software packages are typically composed of:

- A Manifest File
- Package configuration file(s)
- Supporting files (`.inc` files, additional `.php` web interface files, etc.)

See also:

See *Package Port Directory Structure* for a more in-depth list of files and their locations in the package structure.

Manifest File

The manifest file is located inside the package's port directory:

```
<category>/pfSense-pkg-<package name>/files/usr/local/share/pfSense-pkg-<package name>/
↪info.xml
```

For example:

```
sysutils/pfSense-pkg-Cron/files/usr/local/share/pfSense-pkg-Cron/info.xml
```

This file contains basic information about the package. The format of the manifest XML file is as follows:

```
<pfsensepkgs>
    <package>
        <name>someprogram</name>
        <descr><![CDATA[Some cool program.]]></descr>
        <version>%%PKGVERSION%%</version>
        <configurationfile>someprogram.xml</configurationfile>
    </package>
</pfsensepkgs>
```

Note: The `version` entry is updated automatically, use the template as shown.

Package Configuration Files

The manifest specifies a Package Configuration File for the package using the `config_file` tag. The convention is to keep this file inside the `files/usr/local/pkg/` directory inside the port structure for the package.

The easiest way to get a feel for the format is to look at existing packages and how they use these configuration files, how their fields look, how the code behaves, and so on.

The format is:

```
<?xml version="1.0" encoding="utf-8" ?>
<packagegui>
  <copyright></copyright>
  <name></name>
  <title></title>
  <include_file></include_file>
  <aftersaveredirect></aftersaveredirect>
  <menu>
    <name></name>
    <section></section>
    <configfile></configfile>
    <tooltiptext></tooltiptext>
    <url>/pkg.php?xml=package.xml</url>
  </menu>
  <tabs>
    <tab>
      <text></text>
      <url></url>
```

(continues on next page)

(continued from previous page)

```
      <active/>
      <tab_level/>
    </tab>
  </tabs>
  <service>
    <name></name>
    <rcfile></rcfile>
    <executable></executable>
    <description></description>
  </service>
  <plugins>
    <item>
      <type>plugin_name</type>
    <item>
  </plugins>
  <adddeleteeditpagefields>
    <columnitem>
      <fielddescr></fielddescr>
      <fieldname></fieldname>
    </columnitem>
  </adddeleteeditpagefields>
  <fields>
    <field>
      <fielddescr></fielddescr>
      <fieldname></fieldname>
      <description></description>
      <size></size>
      <type></type>
    </field>
  </fields>
  <custom_php_global_functions><!-- PHP function call --></custom_php_global_functions>
  <custom_php_install_command><!-- PHP function call --></custom_php_install_command>
  <custom_php_pre_deinstall_command><!-- PHP function call --></custom_php_pre_deinstall_
↪command>
  <custom_php_deinstall_command><!-- PHP function call --></custom_php_deinstall_command>
  <custom_add_php_command><!-- PHP function call --></custom_add_php_command>
  <custom_add_php_command_late><!-- PHP function call --></custom_add_php_command_late>
  <custom_delete_php_command><!-- PHP function call --></custom_delete_php_command>
  <custom_php_resync_config_command><!-- PHP function call --></custom_php_resync_config_
↪command>
  <start_command><!-- PHP function call --></start_command>
  <custom_php_service_status_command><!-- PHP function call --></custom_php_service_
↪status_command>
  <custom_php_validation_command><!-- PHP function call --></custom_php_validation_
↪command>
  <custom_php_after_head_command><!-- PHP function call --></custom_php_after_head_
↪command>
  <custom_php_command_before_form><!-- PHP function call --></custom_php_command_before_
↪form>
  <custom_php_after_form_command><!-- PHP function call --></custom_php_after_form_
↪command>
</packagegui>
```

Field types

interfaces_selection
Combo/list box with interfaces list:

```
<field>
  <fielddescr>Interface Selection</fielddescr>
  <fieldname>interfaces</fieldname>
  <type>interfaces_selection</type>
  <description>Select interfaces to listen on</description>
  <multiple/><!-- (optional) -->
  <size>10</size><!-- (optional) -->
  <hideinterfaceregex>(wan|loopback)</hideinterfaceregex><!-- (optional) -->
  <showvirtualips/><!-- (optional) -->
  <showips/><!-- (optional) -->
  <showlistenall/><!-- (optional) -->
</field>
```

checkbox
Field with text description and an enable/disable checkbox:

```
<field>
  <fielddescr>Enable</fielddescr>
  <fieldname>enable_package</fieldname>
  <type>checkbox</type>
  <description>Select this option to enable this config</description>
</field>
```

input
Single line text edit element

```
<field>
  <fielddescr>username</fielddescr>
  <fieldname>username</fieldname>
  <type>input</type>
  <description>Enter package username</description>
</field>
```

password
Special input element for passwords, all input will be masked with * symbol on GUI but clear text on XML config file:

```
<field>
  <fielddescr>password</fielddescr>
  <fieldname>password</fieldname>
  <type>password</type>
  <description>Enter password</description>
</field>
```

textarea
Multi-line text edit element:

```
<field>
  <fielddescr>Custom options</fielddescr>
```

(continues on next page)

(continued from previous page)

```
  <fieldname>custom_options</fieldname>
  <type>textarea</type>
  <description>Paste custom config here</description>
  <encoding>base64</encoding><!-- (optional) -->
</field>
```

select
Combo Box with drop-down list items:

```
<field>
  <fielddescr>Some Choice</fielddescr>
  <fieldname>some_choice</fieldname>
  <description><![CDATA[Select a choice]]></description>
  <type>select</type>
    <options>
    <option><name>Choice A</name><value>a</value></option>
    <option><name>Choice B</name><value>b</value></option>
    </options>
  <multiple/><!-- (optional) -->
  <size>10</size><!-- (optional) -->
</field>
```

info
Information text without any options to select:

```
<field>
  <fielddescr>Additional info</fielddescr>
  <fieldname>just_info</fieldname>
  <type>info</type>
  <description>show info text on package GUI</description>
</field>
```

button
Additional buttons to take additional actions on packages:

```
<field>
 <fielddescr>Reload config</fielddescr>
 <fieldname>reload</fieldname>
 <type>button</type>
 <description>click to force a config reload</description>
 <placeonbottom/><!-- Use this option to place the button besides save default␣
↪button -->
</field>
```

In package `.inc` file, to check what button was selected, use:

```
if (($_POST['Submit'] == 'Save') {
   /* Do save stuff */
}
if (($_POST['Submit'] == 'Reload') ||
    !isset($_POST['Submit'])) {
   /* Do reload stuff */
}
```

Field groups

combinefields

Several options can be combined into a single row as an option group using <combinefields>

For example this block groups three options onto a single row:

```
<field>
        <fielddescr>Option 1</fielddescr>
        <fieldname>option1</fieldname>
        <description>Option 1</description>
        <type>input</type>
        <combinefields>begin</combinefields>
</field>
<field>
        <fielddescr>Option 2</fielddescr>
        <fieldname>option1</fieldname>
        <description>Option 2</description>
        <type>input</type>
<field>
        <fielddescr>Option 3</fielddescr>
        <fieldname>option3</fieldname>
        <description>Option 3</description>
        <type>input</type>
        <combinefields>end</combinefields>
</field>
```

rowhelper

Used in pkg_edit.php to add multiple config lines like a table on package GUI. Inside <rowhelper>, add any field type described above:

```
<field>
<fielddescr><![CDATA[Lists]]></fielddescr>
<fieldname>none</fieldname>
<description><![CDATA['Format' - Choose the file format that url will retrieve or␣
↪local file format.]]></description>
<type>rowhelper</type>
  <rowhelper>
    <rowhelperfield>
    <fielddescr>Format</fielddescr>
    <fieldname>format</fieldname>
    <type>select</type>
      <options>
      <option><name>gz</name><value>gz</value></option>
      <option><name>txt</name><value>txt</value></option>
     </options>
    </rowhelperfield>
    <rowhelperfield>
      <fielddescr>URL or local file</fielddescr>
      <fieldname>url</fieldname>
      <type>input</type>
      <size>75</size>
    </rowhelperfield>
  </rowhelper>
```

(continues on next page)

(continued from previous page)

```
</field>
```

adddeleteeditpagefields
: Used with `pkg.php` to have multiple configuration items the same XML page. Useful to access lists, users lists, multi daemon configurations, etc.:

```
<adddeleteeditpagefields>
  <columnitem>
    <fielddescr>Alias</fielddescr>
    <fieldname>aliasname</fieldname>
  </columnitem>
  <columnitem>
    <fielddescr>Description</fielddescr>
    <fieldname>description</fieldname>
  </columnitem>
  <columnitem>
    <fielddescr>Action</fielddescr>
    <fieldname>action</fieldname>
  </columnitem>
  <columnitem>
    <fielddescr>Update Frequency</fielddescr>
    <fieldname>cron</fieldname>
  </columnitem>
</adddeleteeditpagefields>
```

Binaries from FreeBSD

The actual binaries are normal FreeBSD package binaries for that particular program. Once listed as a dependency for a pfSense package in its *Makefile*, the builders compile the dependencies automatically and copy them to the pfSense software package servers. There is no need to specify these in XML.

Updating Packages

When updating a package is it important to bump the version in its *Makefile* otherwise the package will not be rebuilt and made available to others.

Repository Branches

When submitting changes, they are typically submitted to the `devel` branch of the FreeBSD-Ports Repository. In order to show to all users, the changes must be placed in the current release branch as well, such as RELENG_2_8_0.

Ideally, the changes should be submitted to the development branches and tested on systems pulling packages from the development repository. Once the changes have been tested, they can be placed into the release branch for deployment to a wider audience.

Testing/Building Individual Packages

Archive files from pkg may be copied to the firewall and added with pkg directly. The good thing about using pkg is that the GUI packages and CLI packages are all the same. Files for pfSense software packages are all kept together inside the archive, dependencies such as FreeBSD packages are in separate archives.

The package may be compiled on a local FreeBSD 15.0-CURRENT@bf06074106cf builder, then `pkg delete` the old version and then `pkg add` the new one or use any other pkg operations needed.

For example, a basic package like Cron is `pfSense-pkg-Cron-0.3.3`, so if a new copy is built and put on the firewall:

```
# pkg add /path/to/file/pfSense-pkg-Cron-0.3.3.txz
```

It will also work with `pkg add` and a URL to an HTTP or HTTPS web server.

The process for making a package is:

- Check out the FreeBSD-Ports Repository.
- *Locate the port directory*
- Make changes
- Run `make package`:

```
$ git clone git@github.com:pfsense/FreeBSD-ports.git pfSense-ports
$ cd pfSense-ports/blah/pfSense-pkg-foo/
[hack, hack, hack]
$ make package  (might need sudo)
$ scp work/pkg/pfSense-pkg-foo* root@myfirewall:.
```

And then on the firewall:

```
# pkg add pfSense-pkg-foo-<version>.txz
```

Poudriere could also be setup for a custom repository but in most cases that will be overkill.

There are additional considerations when adding files, like updating the *plist*, and crafting a new pfSense package from scratch may be tricky if there is no prior knowledge of how the FreeBSD ports tree works.

38.3.2 Package Port Directory Structure

The directory structure of a package for pfSense® software is similar to that of a traditional FreeBSD port.

For more information on working with FreeBSD Ports, see bsd.port.mk.

This page uses the simple Cron package as an example, many other packages are similar. See *FreeBSD Ports Used for Packages* for links to existing packages to copy/clone from.

Category
: First is the category, which roughly lines up with the category for the package with the caveat that if a pfSense package is based on a FreeBSD port, it should be in the same location (e.g. `haproxy` is under `net/haproxy`)

  In the case of Cron this is `sysutils`:

```
FreeBSD-ports/sysutils/
```

Main package directory
: Inside the category directory, it is always prefixed with `pfSense-pkg-`:

```
FreeBSD-ports/sysutils/pfSense-pkg-Cron/
```

Note: From here on, the `FreeBSD-ports` prefix will be omitted for brevity.

`Makefile`

Includes version information, information about binaries, dependencies, install procedures, where to copy files, and so on. Copy an existing one for a similar package and adjust as needed (but do so carefully):

```
sysutils/pfSense-pkg-Cron/Makefile
```

pfSense software standard package `install/deinstall` scripts

Note: These files are identical for all packages. Copy the contents from another existing package.

```
sysutils/pfSense-pkg-Cron/files/pkg-deinstall.in
sysutils/pfSense-pkg-Cron/files/pkg-install.in
```

Description

A brief text description of the package:

```
sysutils/pfSense-pkg-Cron/pkg-descr
```

Packing List

A list of files installed by the package, for specifics on the format, see the links above:

```
sysutils/pfSense-pkg-Cron/pkg-plist
```

Files Directory

The directory where custom files are placed which will be copied to the firewall:

```
sysutils/pfSense-pkg-Cron/files/
```

The structure under the `files/` directory should follow the same conventions as files on the pfSense installation, which typically follows hier(7) from FreeBSD. For example, executable scripts would go under `files/usr/local/bin/`

Privileges

An optional file containing privilege information:

```
sysutils/pfSense-pkg-Cron/files/etc/inc/priv/cron.priv.inc
```

Package Code and Configuration

The include files and XML files for the package:

```
sysutils/pfSense-pkg-Cron/files/usr/local/pkg/cron.inc
sysutils/pfSense-pkg-Cron/files/usr/local/pkg/cron.xml
```

XML Metadata

```
sysutils/pfSense-pkg-Cron/files/usr/local/share/pfSense-pkg-Cron/info.xml
```

GUI-Accessible Files

Files that go into the main directory of the web server:

```
sysutils/pfSense-pkg-Cron/files/usr/local/www/
```

Note: The Cron package uses a subdirectory under `/usr/local/www` for its files, which is optional.

```
sysutils/pfSense-pkg-Cron/files/usr/local/www/packages/cron/cron.php
sysutils/pfSense-pkg-Cron/files/usr/local/www/packages/cron/cron_edit.php
sysutils/pfSense-pkg-Cron/files/usr/local/www/packages/cron/index.php
```

38.3.3 FreeBSD Ports Used for Packages

The list of packages in the pfSense® software copy of FreeBSD-ports includes:

- ACME (security/pfSense-pkg-acme)
- apcupsd (sysutils/pfSense-pkg-apcupsd)
- arping (net/pfSense-pkg-arping)
- arpwatch (net-mgmt/pfSense-pkg-arpwatch)
- Avahi (net/pfSense-pkg-Avahi)
- Backup (sysutils/pfSense-pkg-Backup)
- bandwidthd (net-mgmt/pfSense-pkg-bandwidthd)
- bind (dns/pfSense-pkg-bind)
- Cellular (net/pfSense-pkg-cellular)
- Cron (sysutils/pfSense-pkg-Cron)
- Darkstat (net-mgmt/pfSense-pkg-darkstat)
- filer (sysutils/pfSense-pkg-filer)
- FreeRADIUS3 (net/pfSense-pkg-freeradius3)
- FRR (net/pfSense-pkg-frr)
- FTP Client Proxy (ftp/pfSense-pkg-FTP_Client_Proxy)
- HAProxy (net/pfSense-pkg-haproxy)
- HAProxy-devel (net/pfSense-pkg-haproxy-devel)
- iperf (benchmarks/pfSense-pkg-iperf)
- LADVD (net/pfSense-pkg-LADVD)
- LCDproc (sysutils/pfSense-pkg-LCDproc)
- Lightsquid (www/pfSense-pkg-Lightsquid)
- lldpd (net-mgmt/pfSense-pkg-lldpd)
- Mail Reports (mail/pfSense-pkg-mailreport)
- MTR (net/pfSense-pkg-mtr-nox11)
- NET-SNMP (net-mgmt/pfSense-pkg-net-snmp)
- Netgate Firmware Upgrade (sysutils/pfSense-pkg-Netgate_Firmware_Upgrade)

- nmap (security/pfSense-pkg-nmap)
- Node Exporter (sysutils/pfSense-pkg-node_exporter)
- Notes (sysutils/pfSense-pkg-Notes)
- nrpe (net-mgmt/pfSense-pkg-nrpe)
- ntopng (net/pfSense-pkg-ntopng)
- NUT (sysutils/pfSense-pkg-nut)
- Open-VM-Tools (emulators/pfSense-pkg-Open-VM-Tools)
- OpenVPN Client Export (security/pfSense-pkg-openvpn-client-export)
- pfBlockerNG (net/pfSense-pkg-pfBlockerNG)
- PIMD (net/pfSense-pkg-pimd)
- RRD Summary (sysutils/pfSense-pkg-RRD_Summary)
- Service Watchdog (sysutils/pfSense-pkg-Service_Watchdog)
- shellcmd (sysutils/pfSense-pkg-Shellcmd)
- siproxd (net/pfSense-pkg-siproxd)
- snmptt (net-mgmt/pfSense-pkg-snmptt)
- Snort (security/pfSense-pkg-snort)
- softflowd (net-mgmt/pfSense-pkg-softflowd)
- Squid (www/pfSense-pkg-squid)
- SquidGuard (www/pfSense-pkg-squidGuard)
- Status_Traffic_Totals (net/pfSense-pkg-Status_Traffic_Totals)
- stunnel (security/pfSense-pkg-stunnel)
- Sudo (security/pfSense-pkg-sudo)
- Suricata (security/pfSense-pkg-suricata)
- syslog-ng (sysutils/pfSense-pkg-syslog-ng)
- System Patches (sysutils/pfSense-pkg-System_Patches)
- Telegraf (net-mgmt/pfSense-pkg-Telegraf)
- tftpd (ftp/pfSense-pkg-tftpd)
- tinc (security/pfSense-pkg-tinc)
- WireGuard (net/pfSense-pkg-WireGuard)
- Zabbix-Agent (net-mgmt/pfSense-pkg-zabbix-agent)
- Zabbix-Proxy (net-mgmt/pfSense-pkg-zabbix-proxy)
- Zeek (security/pfSense-pkg-zeek)

Note that each package is under a category and prefixed with `pfSense-pkg-`.

See *Package Port Directory Structure* for details about the structure of files inside each package directory.

38.3.4 Compiling Software on the Firewall

pfSense® software intentionally does not include a full environment for compiling software (make, headers/includes, sources, etc.) on the installed firewall. Those tools are left out for security reasons.

A virtual machine or separate system can be setup to compile software, and then the compiled binaries/packages/software can be moved over to the firewall.

When doing this, install a version of FreeBSD that matches up with the version of pfSense software currently in use. A list can be found here: *Versions of pfSense software and FreeBSD*

Alternately, install pre-compiled FreeBSD packages as described here: *Installing FreeBSD Packages*

CHAPTER

THIRTYNINE

REFERENCES

These documents are for reference purposes and may be about general topics that do not fit into other categories or topics about the documentation itself.

39.1 Typographic Conventions

Throughout this documentation the authors use conventions to denote certain concepts, information, or actions. The following list gives examples of how to format these items in the documentation.

Menu Selections
: **Firewall > Rules**

GUI Item Labels/Names
: **Destination**

Buttons
:  **Add**

Prompt for input
: **Proceed?**

Input from the user (text)
: `Rule Description`

Input from the user (selection)
: *WAN*

File Names
: `/boot/loader.conf`

Names of commands or programs
: `gzip`

Commands Typed at a shell prompt

```
# ls -l
```

Long shell command-line examples may be split using the backslash (\) for shell line continuation.

Command Output

```
-rw-r--r--  2 root  wheel   887 Apr 12 09:49 .cshrc
```

Special Notes

Note: Consider this …

Warnings

Warning: Watch out!

Tips

Tip: The best practice is …

References

See also:

For more information …

The list above also servers as an example of a "definition list" used to define sets of terms or options and their meanings.

39.2 Documentation Quality Guidelines

39.2.1 Overview

This document aims to give an idea of what types of information should be in articles to keep them as useful as possible for other users of the documentation.

For information about style and formatting, see *Style Guide*.

39.2.2 What to include in articles

- Concise, detailed information – the more accurate the information, the more useful it is.
- Good spelling and grammar
- Good formatting (*Style Guide*)
- Proper categories to classify the article. For example, place the new article in the correct area of the documentation site (e.g. articles about OpenVPN go under `source/vpn/openvpn/`).

39.2.3 What not to include in articles

- Poor grammar
- Inconsistent or poor backing data
- Disparaging comments
- Ambiguity – try to make articles clear, so there is no confusion
- Avoid creating a new directory or category for a single article if possible. If more will be added, then it is OK.
- Redundant information already covered in other documents. Especially in recipes.

- Do not start a recipe with installing pfSense software or other basic tasks. Make assumptions and link to other articles for topics already covered elsewhere. There is no need to reinvent the wheel and have multiple overlapping copies of the same information. This helps keep recipes as concise as possible!

39.2.4 Where to see why an article was flagged

Articles flagged for cleanup should have corresponding issues on GitHub discussing the required changes.

39.3 Style Guide

To make this documentation easier for users, the style of articles should be consistent and clear. The following guidelines are best practices and strong suggestions. Text found to not be following the Language Style/Grammar guidelines may be edited and corrected at any time.

See also:

See *Documentation Quality Guidelines* for information about how entries in the documentation should be written and what they should contain.

39.3.1 Trademarks

- The first use of `pfSense®` **must** have the ® symbol on each and every page.
- Trademarks, such as "pfSense®", are proper adjectives and **not** verbs or nouns.

  In most cases this means that a reference must be phrased as "pfSense® software" instead of "pfSense®" on its own.
- Include "Plus" or "CE" when referring to specific builds of pfSense software. If a topic applies to any build, then this extra reference is not necessary.

  Example: "pfSense® Plus software", "pfSense® CE software"
- Be respectful of Trademarks, including those of third parties.
- Where feasible, try to include relevant marks for third parties or use their preferred nomenclature, phrasing, formatting, and so on.

  For example, use "Proxmox® VE" and not "Proxmox" when referring to the Proxmox Virtualization Environment as that is the way Proxmox Server Solutions GmbH prefers others to reference the product.

Examples:

Firewalls

- Good: Firewall running pfSense® software
- Bad: pfSense® firewall, pfSense® box

Versions

- Good: pfSense® Plus software version 23.01
- Bad: pfSense® 23.01

Installations

- Good: Installation of pfSense® software
- Bad: pfSense® install

39.3.2 Referring to items involving pfSense software

Refer to a firewall running pfSense as a "firewall" or "node". Avoid other similar terms ("router", "system", "box", etc.) for clarity and consistency.

39.3.3 Capitalization

Capitalize terms correctly! Especially **pfSense**!

No other capitalization of "pfSense" may be used except in a URL which is acceptable as lowercase (e.g. https://www.pfsense.org). If a sentence begins with "pfSense" the first letter must remain lowercase.

Note: If any other usage of "pfSense" or a misspelling of same is present in a document ("PFsense", "PFSense", "pfSence", etc), fix it immediately.

Other special notes for capitalization:

- WebGUI, IPsec, OpenVPN, Internet, Ethernet, VPN, DNS, PPPoE, IPv4, IPv6, NPt, strongSwan, pfsync, pftop, JavaScript, WireGuard.

39.3.4 What to Avoid

Avoid addressing the user directly ("you", "your", etc.)
Rewrite sentences to avoid addressing the user when found. Exceptions may be made for quoted/cited text or other unavoidable circumstances.

Avoid references to the writer ("I", "we")
Except when making specific recommendations, which is OK to avoid using passive voice. "We recommend" is better than "It is recommended", though the ideal phrasing would be "the best practice is".

Avoid the use of words such as "should", "could", "might"
Words that do not commit to a specific action/result are undesirable. For example "This should happen" or "That might appear". Some instances are expected/required when making recommendations, but reword where feasible.

Avoid the use of Weasel words
See Weasel Words for reference.

Avoid redundant phrases
This especially includes acronym references that duplicate words: "WAN Network", "LAN Network", "DUID Identifier", "6RD Rapid Deployment". Remove the redundant word(s) and/or use alternate phrasing ("WAN Subnet" or "Network on the WAN interface" rather than "WAN Network")

Avoid unnecessary abbreviations and shortening of words
This creates ambiguity, for example:

- Avoid using "IP" or "IPs" to refer to IP addresses. Use the full form "IP address" instead to remove ambiguity.
- Avoid using "config" when "config.xml" or "configuration" is more clear.
- Avoid using "ovpn" to mean "OpenVPN" except in cases when the OS-level interfaces are being referenced (`ovpncX` is OK, "Use OVPN instead" is not.)

Avoid using "here" for links
Do not make links for "here", "click here", or similar phrasing. They provide no context for the link, cause redundancy in phrasing, and cause problems for users that require accessibility functions such as screen readers.

See also:

See recommendations from W3C Tips and uxmovement.

Avoid awkward possessive references
For example: "firewall's", "pfSense's".

Avoid gender-specific pronouns
Example: "his", "hers"

Avoid leaving out necessary hyphens
Example: "howto" should be "how-to"

Avoid "Britishisms" or other non-en_US style spellings
Use "Flavor", not "Flavour". Use "Specialize", not "Specialise"

Avoid confusing bandwidth specifications

- Avoid confusing bits and bytes. Use Big **B** for bytes, little **b** for bits.
- Avoid ambiguity when specifying bandwidth measurements. Bandwidth should always have *time* component, such as *per second*. "5 Mbit/s" is much clearer than "5 Mbps", "5 Mb", or "5MB".

Avoid passive voice when necessary
When it makes more sense to do so, try to make it clear who is taking an action in a sentence, and use active voice instead of passive voice where possible in these cases.

- Bad: "When a packet is routed"
- Okay: "When a packet is routed by the firewall"
- Better: "When the firewall routes a packet"

There are times when the person or item taking an action isn't as important as the receiver of the action. The performer of the action may be unknown, assumed, or unimportant. In these cases the passive voice is acceptable. For example when describing GUI interactions, "When X is clicked" may be less awkward than repeating "When the user clicks X".

39.3.5 Crafting Instructions

When forming lists of instructions, start each item with an action word when possible: Click x, Select x, Enter x, Navigate to x.

When instructing a user to reach a specific page or place in the GUI, reference this action as "Navigate to".

39.3.6 Example Text

When offering examples, keep the following in mind:

- Domain names should use `example.com` or another reserved name from RFC 2606.
- IP address examples should be taken from subnets reserved for documentation in RFC 6890: `192.0.2.0/24`, `198.51.100.0/24`, `203.0.113.0/24` or the traditional RFC 1918 networks `192.168.0.0/16`, `172.16.0.0/12`, or `10.0.0.0/8` if the documentation subnets are not sufficient.
  - In some cases where additional unique examples are needed, use the benchmarking subnet `198.18.0.0/15`.

39.3.7 High Availability / CARP References

- Refer to the cluster as a "High Availability Cluster" or "HA Cluster" and **not** as a "CARP Cluster".
- Use the term "node" as in "cluster node" for referencing an individual unit.
- Use the term "primary" for the primary node, never "master" as this can be confused with the CARP VIP status.
- Use the term "secondary" for the secondary node, never "backup" or "slave" as this is outdated terminology and can be confused with the CARP VIP status.
- Use the terms "Sync interface" or "Interconnect interface" when referring to the dedicated interface between HA Cluster nodes. **Never refer to that interface as a "CARP interface"**.

39.4 Formatting Guide

The pfSense® software documentation is built using Sphinx/reStructuredText (reST). The formatting is similar in some ways to Markdown, but has significant differences. To get a feel for the formatting, and read through this document.

Tip: Test out how different markup is rendered using the Online reStructuredText editor. Additional information can be found at A primer on reStructuredText and reST/Sphinx cheat sheet.

39.4.1 Filenames

When adding new pages or images use all lowercase letters and hyphens instead of spaces for the filename. This is commonly referred to as a slug, or a "slugified" version of the text. For example, this file is named `formatting-guide.rst`. When possible, try to find a meaningful single word for the filename instead of using lengthy names.

Tip: Use an online slug generator.

39.4.2 Text

In general try to keep text in logical paragraphs wrapped at 80 characters. This ensures the source is easy for everyone to read no matter where it is being edited. For long pages with several sections that may only be relevant to some users, split the page into several smaller documents.

Basic Inline Formatting

Add basic inline formatting to the text as follows:

- One asterisk: `*text*` for *emphasis* (italics),
- Two asterisks: `**text**` for **strong emphasis** (boldface), and
- Double backticks: ``` ``text`` ``` for `code samples`.
- To start a block of literal formatted text such as code, console input/output, log data, etc., use `.. code-block::` followed by a blank line. Prefix each line inside the block with three spaces:

```
.. code-block::

   code
```

These can be applied to text in various ways within the documentation:

- Menu references use bold text and ">" with spaces in between to separate menus from menu items: **System > General**
  - Navigation that also refers to a tab name should be formatted like so: **System > Advanced**, **Miscellaneous** tab.
- GUI text references and option names use bold text: **Description**
- Text to be entered or replaced by the user uses backticks: Enter `192.168.1.1` for the **IPv4 Address**.
- Options selected by the user from a list or drop-down are italic: Select *WAN* for the **Interface**.
- File names and paths use backticks: `/root`
- Commands names inline with other text use backticks, "The `sudo` command ...".
- Shell commands being demonstrated or directed use code blocks. Start a code block as usual and then use # followed by a space to simulate a command prompt:

```
# ls -l /root
```

- Program output also uses code blocks. Blank lines in output can either be blank so that a line does not contain only whitespace:

```
# someprogram
Output:

Foo
```

- Do not leave any trailing whitespace on lines or any lines containing only whitespace.
- Lines should be indented with spaces, not tabs. Tabs may be used inside literal blocks.
- Use UTF-8 encoding.

39.4.3 Headings

Headings consist of text and a line of characters underneath ("underline") the **same length** as the text. The specific characters must be consistent to denote sections of the same depth in a single file. Parts and chapters also use a similar row of characters above the text ("overline").

- `####` with overline, for parts
- `****` with overline, for chapters
- `====`, for sections
- `----`, for subsections
- `~~~~`, for subsubsections
- ```` ```` ````, for paragraphs

Example:

```
This is a Section Heading
=========================
```

The specific characters are not as important as being consistent. The parser considers **each file separately**, so there is no need to track and be consistent across multiple files even if they are in the same chapter/section/etc.

Note: The headings are also how the "On This Page" section is generated. When possible, use headings that create an outline of the content, making it easy for the reader to scan.

39.4.4 Lists

Sphinx offers several types of item lists suitable for various purposes.

Unordered Lists

Place an asterisk followed by a space at the start of a paragraph and indent two additional spaces for any lines that wrap.

```
* This is a bulleted list.
* It has two items, the second
  item uses two lines.
```

Which renders as:

- This is a bulleted list.
- It has two items, the second item uses two lines.

Ordered lists

Ordered lists are auto-numbered by prefixing a line with `#.` followed by a space. Indent lines that wrap with three additional spaces:

```
#. This is a numbered list.
#. It has two items too.
   This second item has two lines, too.
```

Which renders as:

1. This is a numbered list.
2. It has two items too. This second item has two lines, too.

Nested lists

Nested lists are possible, but be aware that they must be separated from the parent list items by blank lines:

```
* this is
* a list

  * with a nested list
  * and some sub-items

* and here the parent list continues
```

Which renders as:

- this is
- a list
  - with a nested list
  - and some sub-items
- and here the parent list continues

Definition lists

Definition lists are created as follows:

```
Term (up to a line of text)
  Definition of the term, which must be indented

  and can even consist of multiple paragraphs

Next Term
  Description.
```

Note: The term itself cannot have more than one line of text.

Which renders as:

Term (up to a line of text)
Definition of the term, which must be indented

and can even consist of multiple paragraphs

Next Term
Description.

> **Warning:** Ordered, unordered, definition, and field lists cannot be nested more than **three** levels deep, or the PDF build will fail.

Field Lists

Field lists are perfect for lists of options. They start with the field name starting and ending with a colon, followed by text describing the field. The content should start on a new line under the field name and should be indented two space. Though starting on a new line is optional in the parser, starting on a blank line makes the formatting cleaner and easier to follow.

```
:Option Name:
  What it does.
:Option 2:
  Another option. This is a long description that wraps to the next line,
  with two spaces indentation.
:Third Option:
  Something else.
```

Which renders as:

Option Name
What it does.

Option 2
Another option. This is a long description that wraps to the next line, with two spaces indentation.

Third Option
Something else.

> **Warning:** Ordered, unordered, definition, and field lists cannot be nested more than **three** levels deep, or the PDF build will fail.

39.4.5 Links

External Link

Separate the link and the target definition, like this:

```
This is a paragraph that contains `a link`_.
```

```
.. _a link: http://example.com/
```

Place the target definition at the bottom of the page in alphabetical order.

Note: If the link text will contain a colon, escape it in both the link text and the definition, for example:

```
See `Link\: Stuff`_.
```

```
.. _Link\: Stuff: http://example.com/stuff
```

Cross-Reference to Section of Document

To make a cross-reference to another document, first create a label immediately before the section title:

```
.. _label-some-section:

Some Section
------------
```

And then in the other document, reference it using `:ref:` and the given label:

```
See :ref:`label-some-section` for more information
```

Cross-Reference to Entire Document

If a cross-reference will instead reference an entire document rather than a specific section, use the `:doc:` method instead.

For example, to reference this entire document, `/references/style-guide.rst`, use the following text, omitting the file extension:

```
:doc:`/references/style-guide`
```

39.4.6 Images

Images

Place images in the `source/_static` directory in the same folder structure as the page that the image is going to be posted on. For example, an image within `source/references/fomatting-guide.rst` would go in `source/_static/references/image.png`.

```
.. image::
   /_static/filename.*
   :align:
     center
   :alt:
     Alternative text that describes the image
   :target:
     /_static/filename.png
```

Note: `:target:` is optional and only necessary if it is a large image.

Figures

Place figures in the `source/_static` directory in the same folder structure as the page that the image is going to be posted on. For example, an image within `source/references/fomatting-guide.rst` would go in `source/_static/references/image.png`.

Figures are similar to images, but need a unique label and a caption for proper in-text references.

```
.. _figure-my-stuff:
.. figure::
   /_static/stuff.*
   :figclass:
     align-center
   :target:
     /_static/stuff.png

   This is the caption
```

Note: `:target:` is optional and only necessary if it is a large image.

Which can be referred to using the following:

```
An example is shown in Figure :ref:`figure-my-stuff`.
```

Warning: The indention is important! The caption *must* be aligned properly with the other attributes!

Inline Images

Inline images (no breaks above or below, aka inline with the text) require using a substitution. Since inline images are typically inserted on many pages, inline image files can be placed in the root of `source/_static` or an appropriate subdirectory.

Many common icon substitutions are available in a common substitutions file which is automatically included in all source files and usable as follows:

```
To add a blah, click |fa-plus|.
```

To do this in a one-off fashion, use a substitution within the same file:

```
Click |image_icon_edit| to edit the entry
<rest of page>
.. |image_icon_edit| image:: _static/icons/fa-pencil.*
```

39.4.7 Tables

Grid Tables

The grid must be "painted", they look like this example:

```
+------------------------+------------+----------+----------+
| Header row, column 1   | Header 2   | Header 3 | Header 4 |
| (header rows optional) |            |          |          |
+========================+============+==========+==========+
| body row 1, column 1   | column 2   | column 3 | column 4 |
+------------------------+------------+----------+----------+
| body row 2             | ...        | ...      |          |
+------------------------+------------+----------+----------+
```

Simple Tables

These are easier to write, but are limited: they must contain more than one row, and the first column cells cannot contain multiple lines. They look like this:

```
=====  =====  =======
A      B      A and B
=====  =====  =======
False  False  False
True   False  False
False  True   False
True   True   True
=====  =====  =======
```

39.4.8 Table of Contents

Every file has to a part of a `toctree` or **Table of Contents** tree, as this is how the side navigation is built.

Reference RST files by their filenames without their `.rst` extension. It is also possible to link to external resources if necessary, as shown with the YouTube link:

```
.. toctree::
   :maxdepth:
     2

   filename1
   filename2
   subdirectory/index
   Example YouTube Video <https://youtu.be/Cwz7vWu_KO0>
```

Local Table of Contents

Sometimes it is useful to add the table of contents of the current page:

```
.. contents::
   :depth:
     2
```

39.4.9 Admonitions (Colored Boxes)

Admonitions are text, distinguished in formatted boxes, that bring attention to important items.

Admonitions are similar to other directives, starting with two periods and a space followed by the admonition name and then two colons, e.g. `.. note::`.

Though the parser allows the text to start after the directive, the best practice is to start the text content on a line **under** the admonition which makes the formatting easier to follow and wrap consistently.

The most common example is a "Note" box:

```
.. note::
   This is a note, it will be surrounded by a note box when it is built.
```

Which renders as:

Note: This is a note, it will be surrounded by a note box when it is built.

Admonitions are available for a wide variety of types, including:

- `note`
- `tip`
- `warning`
- `attention`
- `caution`
- `danger`
- `error`
- `hint`
- `important`
- `seealso`

39.4.10 Substitutions

reST substitutions are pieces of text and/or markup which the parser will replace with another (usually longer) string or element. Substitutions are referred to in the text by `|name|`.

Substitutions are useful for inline images, links or references to be used in many documents, or even commonly repeated strings which may need to change over time. Rather than finding and replacing every reference so such items manually, they can be managed from a shared source.

Substitutions are defined like footnotes with explicit markup blocks, for example:

```
.. |name| replace:: replacement *text*
```

or this:

```
.. |caution| image:: warning.png
              :alt:
                Warning!
```

To use substitutions for multiple documents, put them into a separate file and include it into all documents where they will be used, using the `include` directive. Give the include file a file name extension differing from that of other source files, such as `.rsti`, to avoid Sphinx finding it as a standalone document.

Note: The common substitutions file is automatically included in all documents. Check that file before adding more substitutions in other files. Substitutions which will be widely used in many documents should be placed there.

39.4.11 Literal (code) Blocks

Briefly described earlier, literal or "code" blocks allow for pre-formatted text, most commonly used for source code, shell commands, command output, and so on.

A code block for general use, such as for log data, shell input/output, or command examples should be started with `.. code-block::` followed by a blank line. Prefix each non-blank line inside the block with three spaces.

```
.. code-block::

   code
```

The lines inside the code block must be indented to the same level, usually three spaces.

For a blank line inside the code block, use a completely empty line, not a line containing only whitespace.

For more complex examples, syntax highlighting can be used for source code using the `code-block` directive followed by the type of content inside the block:

```
.. code-block::
   html
   :linenos:

   <b>some HTML</b>
```

Which renders as:

```
<b>some HTML</b>
```

39.5 Updating Documentation

The primary avenue to make Netgate aware of problems in the documentation is to open an Issue on Redmine. In addition to requesting corrections, this may also be used to request new content or to propose other changes.

When creating an issue, please be specific and reference documents by filename where appropriate.

39.6 Developer Style Guide

This page covers rules and styles to be used when submitting code for inclusion in pfSense® software.

39.6.1 Modularity and Organization

Historically pfSense software, and before it M0n0wall, have placed all the functionality required to obtain configuration data, display it, edit it, validate it, format it, and save it in a single PHP file (web page). Much use was made of global variables to allow functions access to the required data. These are both undesirable strategies.

The "everything in one file" approach leads to large, difficult to manage files. Functions declared in PHP web pages cannot be accessed from outside that page leading to duplication, bloat and maintenance headaches. Global variables are fragile and often mysterious, particularly when the variable is declared in another file. What does `global $p_interface_remote;` mean? What type is it, who/what sets it, where is it declared, is it multi-user safe, etc.?

pfSense software functionality, particularly in PHP web pages, should adhere to MVC methodologies. The PHP web page file should contain only the code to display and edit information. The code to retrieve it, validate it, update it and save it should be provided in discrete functions included ("required") from a separate file (preferably in `/usr/local/pfSense/include/www`).

All functions should be organized in a modular fashion with clearly defined inputs and return values documented in the comments. The use of global variables should be minimized. Global variables such as `$g`, `$config`, `$input_errors`, etc. are unavoidable and so acceptable, but where possible values should be passed in either directly or by reference.

References:

- https://en.wikipedia.org/wiki/Modular_programming
- https://en.wikipedia.org/wiki/Model%E2%80%93view%E2%80%93controller

39.6.2 Developer Rules

- Never commit untested code.
- If a developer breaks the code, they fix it, even if their code breaks another subsystem. This is not glamorous, but it is the right thing to do (or back the code out until a proper fix can be obtained).
- If a developer commits a kernel change that requires a userland configuration change of any type, then that developer must also make sure there is PHP code to control the userland change as needed (different sysctl tunables, different means of configuring something, upgrade code etc.). This could be done by the same person or coordinated with other developers so all the changes go in at the same time.
- Internal developers (employees, etc.) should use branch development and merge requests, do not commit directly to master branch or equivalent where possible.
- Avoid changing the XML configuration structure whenever possible. Where that is infeasible, do not commit code that changes the XML configuration structure without adding configuration upgrade code at the same time.

- Never push a commit to any branch that knowingly causes a regression.
- Any major or high risk changes must first be done in a git clone (community contributors), branch (employees/internal developers) and reviewed by another committer before merging into the `master` branch.

  Using a *pull request* or merge request for this workflow is acceptable, and is how every change from those outside the development team is handled.
- Mention relevant ticket numbers in commit messages so that Redmine can associate the changes. See *Referencing Tickets in Commit Messages*.
- When possible, compose and format commit messages similar to entries in the change logs. For example:
  - Brief first sentence that describes the commit
  - Start with an action word describing the nature of the change ("Adds…", "Changes…", "Improves…", "Corrects…")
  - Reference the ticket number in the first sentence (see previous bullet point)
  - Examples:
    - "Change X to Y which fixes #1234"
    - "Correct check for XYZ in some_page.php to prevent badthing. Fixes #2345"
    - "Add coolnewthing to some_page.php. Implements #3456"
  - Add a blank line after the first sentence.
  - Longer, more detailed explanations can be placed in the body of the commit message.
- Always use full paths when calling an executable (e.g. `/usr/bin/grep` NOT `grep`)

39.6.3 HTML Specific Rules

Note: Incorrect HTML code is treated as broken code. Breaking the code is not allowed. A C compiler for example would complain in most cases if a developer breaks the code syntactically. Web browsers may ignore invalid code, but this does not mean that the code is not broken. The broken code must be fixed by the person who committed the invalid code.

---

pfSense software uses the XHTML DOCTYPE in the GUI code. The DOCTYPE enforces code against the following ruleset:

- Use lower case tag names and not a mix of uppercase and lowercase tag names
- Breaks must be closed (`<br />`)
- Image tags must be closed (`<img />`)
- An image tag *always* has an alt attribute, though it may contain no value
- Horizontal rule tags must be closed (`<hr />`)
- HTML form fields must be closed (`<input />`)
- Ampersands in a URL (e.g. within a `href` attribute) must be coded as an HTML entity (e.g. `&`)
- Special characters (e.g. umlauts) must be coded as HTML Entities:
- A `<table />` tag does not have a name attribute
- A `<div />` tag does not have a name attribute

- A tag does not have a name attribute
- A tag does not have a name attribute
- Checkbox checked attributes must be coded as checked="checked"
- HTML field disabled attributes must be coded as disabled="disabled"
- HTML field readonly attributes must be coded as readonly="readonly"
- Any HTML <input /> field has a type attribute (e.g. type="text")
- Opening , <b> tags must have a matching closing tag (e.g.)
- <table /> tags do not contain a <form /> tag
- The type attribute of a <form /> tag must contain a lower case value (e.g. type="post" or type="get")
- The language=JavaScript attribute for <script /> tags is deprecated. Use the type="text/javascript" attribute instead.
- Always use lowercase attribute names for calling JavaScript events (e.g. onclick="foobar();")
- The <embed /> tag is deprecated, use <object /> instead
- If a style attribute is assigned to an HTML element it must be enclosed by quotes, for example: element.style.borderTop = "2px solid #990000";

It is possible to syntax check code in Firefox with the HTML validator plugin, or use the W3C validator. The latter is possible in Opera even for RFC 1918 networks.

39.6.4 PHP Specific Rules

Rule #1: Any time a rule is broken, there must be proper justification and documentation.

General rules

- All PHP files must start with a header block in English
- Use descriptive variable names in English
- Use lowercase variable names ($my_very_long_var_name) or Camel Case names ($myVeryLongVarName)
- When referencing variables inline in double-quoted strings, use braces around the variable names:

```
$foo = "bar{$bar}bar";
```

- Avoid shell execution if at all possible. If it is unavoidable, ensure all variables passed to the shell are escaped using escapeshellarg() or similar
- Do not print user input back to the user without encoding it in some way (e.g. htmlspecialchars())
- Add comments in English whenever necessary or helpful
- Use // or /* */ style syntax for single line comments, do not use #
- Use /* */ style syntax for multi-line comments
- Use elseif and not else if when given a choice. The else if variant only works with braced syntax and not colon syntax (e.g. if: ... elseif: ... endif;).
- For testing the same variable against multiple strings or values directly, use a switch statement rather than a long chain of if/elseif/elseif/elseif/[...]/else statements.

- Add `TODO:` comments, when there is something to be done
- Add `FIXME:` comments, when something is broken
- add `NOTE:` comments, when there is something important other people should know beyond a traditional comment, for example a warning about not changing code in certain ways.
- Try to code in a readable way:

```
$header = "<head>{$foo}</head>";
$message = "SOME{$bar}TEXT";
```

  Is easier to read than:

```
$header="<head>".$foo."</head>";
$message = "SOME" . $bar . "TEXT";
```

- Try to simplify code for better readability:

```
if ($bool1)
   if ($bool2)
       if ($bool3)
           do_it();
whatever();
```

- Should be written as:

```
if ($bool1 && $bool2 && $bool3) {
   do_it();
}
whatever();
```

- Do not set unnecessary or single-use variables:

```
$is_set = isset($var);
if ($is_set) ...
```

- Loop variables are `$i`, `$j`, `$k`, …
  - Do **NOT** use `$g` for a loop variable, as it conflicts with the global `$g` used by pfSense software
- All `switch` statements must have a `default`
- In classes, use `private`, `protected` and `public`, not `var` for attribute declaration
- Do not to use deprecated or obsolete syntax or functions
  - Keep an eye on future versions of PHP to avoid using functions that will be deprecated in the future as well
- If a PHP-internal function is an alias for another function, use the original (i.e. use `exit()` instead of `die()`)

Indent style

- Use K&R, BSD KNF variant style:

```
if ($x == $y) {
   something();
   ...
} else {
   somethingelse();
   ...
}
finalthing();
```

- When creating `if`, `for`, `foreach`, and other similar block style structures, even if there is only one statement inside, the use of braces is required.

  For example, good:

```
if ($foo) {
   something();
}
```

  Not good:

```
if($foo)
   something();
```

- If a conditional statement must span multiple lines, indent using four spaces to align with the start of the conditional above it:

```
   if ($foo1 && $foo2 && $foo3 && $foo4 && $foo5 && $foo6 &&
       $foo7 && $foo8 && $foo9) {
           something();
   }
```

- Do not put be a space between a function name and its argument list:

```
isset($myvar);
```

  - Conditional/control statements such as `if`, `foreach`, and `switch` are exceptions to this. Those must have a space before the parenthesis.

- … but **do** separate function arguments with a single space:

```
do_something($foo, 27, false);
```

- Use tabs for indentation – NOT spaces or a mixture of both
- … but spaces are OK in the middle of a line and for conditional alignment
- Use a tab stop of 8, rather than 4, in an editor.
- Ensure there is NO trailing whitespace at the end of a line, for example spaces or tabs when there is no more text afterward
- Ensure there is NO whitespace on empty lines. For example, a line must not contain only spaces or only tabs

Configuration Manipulation

- Boolean values which are false should be unset:

```
$config['system']['enablesshd'] = "no";
```

  should be:

```
unset($config['system']['enablesshd']);
```

39.6.5 JavaScript Specific Rules

- pfSense software does not try to be compatible with outdated browsers, so do not take special measures to use code required by old/obsolete browsers or rendering engines
- pfSense software includes, among other JavaScript resources, Bootstrap and jQuery. While native JavaScript is best for simple tasks, if a developer can accomplish a goal easily using an included library, they can use it instead
- pfSense software does not currently utilize `transpiler` or similar utilities
- Take special care with user input or statements/variables that can be populated with user input to avoid creating a vulnerability vector such as XSS. User fields must be encoded or otherwise sanitized
  - For example, be extremely cautions of values inserted into JavaScript via PHP variables. `json_encode()` can help avoid a situation where a user-supplied string could include text such as quotes or semicolons that leads to execution of arbitrary JavaScript

39.6.6 Shell Script Specific Rules

- Use braces in **all** variable references for proper parameter expansion:

```
${SOMETHING}
```

39.6.7 Ports/Packages Specific Rules

When working with the `pkg` system and FreeBSD ports structure, adhere to the FreeBSD guidelines for code in these files.

Useful resources for working with `pkg` and ports include:

- The FreeBSD Porter's Handbook
- The FreeBSD Ports bsd.port.mk file
- Use portlint to check the syntax of the Makefile and other supporting files
  - Install `portlint` on a FreeBSD system and run the following command inside the root directory of the port:

```
portlint -CN
```

- Run the following command to make sure the contents of `pkg-plist` are correct:

```
make -DNO_DEPENDS check-plist
```

Other Guidelines:

- A port version or revision must increase for the port to be rebuilt, otherwise changes will not propagate to the `pkg` servers to be picked up by clients
  - For very minor changes, add or increase the `PORTREVISION` line immediately beneath `PORTVERSION` in the `Makefile`, starting at `1`, for example: A second revision would be `PORTREVISION=2`
  - For more significant changes, increase `PORTVERSION`
    - When increasing `PORTVERSION`, completely remove any `PORTREVISION` line, do not comment it out
  - Do not add or change `PORTEPOCH` except under direction of a committer

39.6.8 External Code

Code that has been imported from an external source does not need to be changed to fit these guidelines.

39.6.9 Editor Configuration

The pfSense software project uses a similar coding style to FreeBSD, which has editor configurations for Emacs and Vim. The FreeBSD man page style(9) contains additional relevant material.

CHAPTER

FORTY

LICENSING

pfSense® software uses a combination of Open Source software subject to several different licenses.

The following list shows each Open Source component along with its license.

Package Name	License
beep	*BSD4CLAUSE*
bind-tools	*MPL20*
bsnmp-regex	*BSD3CLAUSE*
bsnmp-ucd	*BSD2CLAUSE*
ca_root_nss	*MPL20*
ccid	*LGPL21*
choparp	*BSD3CLAUSE*
cpdup	*BSD2CLAUSE*
curl	*MIT*
cyrus-sasl	*BSD4CLAUSE*
dbus	*GPLv2*
devcpu-data	*BSD2CLAUSE*
devcpu-data-amd	*EULA*
devcpu-data-intel	*EULA*
dhcp6	*BSD3CLAUSE*
dhcpleases	*APACHE20*
dhcpleases6	*APACHE20*
dmidecode	*GPLv2*
dnsmasq	*GPLv2*
dpinger	*BSD2CLAUSE*
expat	*MIT*
gettext-runtime	*LGPL21+* and *GPLv3+*
glib	*LGPL20*
hostapd	*BSD3CLAUSE*
icu	*ICU*
iftop	*GPLv2*
igmpproxy	*GPLv2+*
indexinfo	*BSD2CLAUSE*
ipmitool	*BSD3CLAUSE*
isc-dhcp44-client	*MPL20*
isc-dhcp44-relay	*MPL20*
isc-dhcp44-server	*MPL20*
json-c	*MIT*
ldns	*BSD3CLAUSE*
libargon2	*CC0-1.0*

continues on next page

Table 1 – continued from previous page

Package Name	License
libedit	*BSD2CLAUSE*
libevent	*BSD3CLAUSE*
libffi	*MIT*
libgcrypt	*LGPL21+* and *GPLv2+*
libgpg-error	*LGPL21* and *GPLv2*
libiconv	*GPLv3*
libidn2	*GPLv3*
libinotify	*MIT*
libltdl	*LGPL21*
liblz4	*BSD2CLAUSE* and *GPLv2*
libmcrypt	*LGPL21+*
libnghttp2	*MIT*
libssh2	*BSD3CLAUSE*
libucl	*BSD2CLAUSE*
libunistring	*GPLv2* and *GFDL* and *LGPL3+*
libuv	*NODE*
libxml2	*MIT* and *TRIO*
libxslt	*MIT*
links	*GPLv2*
lua-resty-core	*BSD2CLAUSE*
lua-resty-lrucache	*BSD2CLAUSE*
luajit-openresty	*MIT* and *PD*
lzo2	*GPLv2*
minicron	*BSD2CLAUSE*
miniupnpd	*BSD3CLAUSE*
mobile-broadband-provider-info	*PD*
mpd5	*BSD3CLAUSE*
mpdecimal	*BSD2CLAUSE*
nginx	*BSD2CLAUSE*
nss_ldap	*GPLv2*
ntp	*BSD2CLAUSE*
oniguruma	*BSD2CLAUSE*
openldap24-client	*OPENLDAP*
opensc	*LGPL21*
openvpn	*GPLv2*
openvpn-auth-script	*APACHE20*
pam_ldap	*GPLv2+* and *LGPL20+*
pam_mkhomedir	*BSD4CLAUSE*
pcre	*BSD3CLAUSE*
pcre2	*BSD3CLAUSE*
pcsc-lite	*BSD3CLAUSE* and *GPLv3+*
perl5	*GPLv1+* or *ART10*
pftop	*BSD2CLAUSE*
php74	*PHP301*
php74-bcmath	*PHP301*
php74-bz2	*PHP301*
php74-ctype	*PHP301*
php74-curl	*PHP301*
php74-dom	*PHP301*
php74-filter	*PHP301*

continues on next page

Table 1 – continued from previous page

Package Name	License
php74-gettext	*PHP301*
php74-intl	*PHP301*
php74-json	*PHP301*
php74-ldap	*PHP301*
php74-mbstring	*PHP301*
php74-opcache	*PHP301*
php74-openssl	*PHP301*
php74-openssl_x509_crl	*MIT*
php74-pcntl	*PHP301*
php74-pdo	*PHP301*
php74-pdo_sqlite	*PHP301*
php74-pear	*PHP301*
php74-pear-Auth_RADIUS	*BSD3CLAUSE*
php74-pear-Cache_Lite	*LGPL21*
php74-pear-Crypt_CHAP	*BSD3CLAUSE*
php74-pear-HTTP_Request2	*BSD3CLAUSE*
php74-pear-Mail	*BSD3CLAUSE*
php74-pear-Net_IPv6	*BSD2CLAUSE*
php74-pear-Net_SMTP	*BSD2CLAUSE*
php74-pear-Net_Socket	*BSD2CLAUSE*
php74-pear-Net_URL2	*BSD3CLAUSE*
php74-pear-XML_RPC2	*PHP301*
php74-pecl-mcrypt	*PHP301*
php74-pecl-radius	*BSD3CLAUSE*
php74-pecl-rrd	*PHP301*
php74-phpseclib	*MIT*
php74-posix	*PHP301*
php74-readline	*PHP301*
php74-session	*PHP301*
php74-shmop	*PHP301*
php74-simplepie	*BSD3CLAUSE*
php74-simplexml	*PHP301*
php74-sockets	*PHP301*
php74-sqlite3	*PHP301*
php74-sysvmsg	*PHP301*
php74-sysvsem	*PHP301*
php74-sysvshm	*PHP301*
php74-tokenizer	*PHP301*
php74-xml	*PHP301*
php74-xmlreader	*PHP301*
php74-xmlwriter	*PHP301*
php74-zlib	*PHP301*
pkg	*BSD2CLAUSE*
py38-ply	*BSD3CLAUSE*
py38-setuptools	*MIT*
python38	*PSFL*
radvd	*RADVD*
rate	*GPLv2*
readline	*GPLv3*
rrdtool	*GPLv2*

continues on next page

Table 1 – continued from previous page

Package Name	License
scponly	*BSD2CLAUSE*
smartmontools	*GPLv2+*
sqlite3	*PD*
sshguard	*BSD2CLAUSE*
strongswan	*GPLv2*
uclcmd	*BSD2CLAUSE*
unbound	*BSD3CLAUSE*
vstr	*LGPL21+*
wol	*GPLv2+*
wpa_supplicant	*BSD3CLAUSE*
xinetd	*XINETD*

40.1 BSD2CLAUSE License Text for multiple packages

The following license text applies to multiple packages:

Listing 1: Download: BSD2CLAUSE for multiple packages

```
Copyright <YEAR> <COPYRIGHT HOLDER>

Redistribution and use in source and binary forms, with or without modification,
are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this
   list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice,
   this list of conditions and the following disclaimer in the documentation
   and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR
ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON
ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.2 BSD3CLAUSE License Text for multiple packages

The following license text applies to multiple packages:

Listing 2: Download: BSD3CLAUSE for multiple packages

```
Copyright <YEAR> <COPYRIGHT HOLDER>

Redistribution and use in source and binary forms, with or without modification,
are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this
   list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice,
   this list of conditions and the following disclaimer in the documentation
   and/or other materials provided with the distribution.

3. Neither the name of the copyright holder nor the names of its contributors
   may be used to endorse or promote products derived from this software without
   specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR
ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON
ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.3 BSD4CLAUSE License Text for multiple packages

The following license text applies to multiple packages:

Listing 3: Download: BSD4CLAUSE for multiple packages

```
Copyright (c) <year> <owner>. All rights reserved.

Redistribution and use in source and binary forms, with or without modification,
are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this
   list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice,
   this list of conditions and the following disclaimer in the documentation
   and/or other materials provided with the distribution.
```

(continues on next page)

(continued from previous page)

```
3. All advertising materials mentioning features or use of this software must
   display the following acknowledgement:

   This product includes software developed by <the organization>.

4. Neither the name of <the copyright holder> nor the names of its contributors
   may be used to endorse or promote products derived from this software without
   specific prior written permission.

THIS SOFTWARE IS PROVIDED BY <COPYRIGHT HOLDER> "AS IS" AND ANY EXPRESS OR
IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT
SHALL <COPYRIGHT HOLDER> BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR
BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING
IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY
OF SUCH DAMAGE.
```

40.4 GPLv2+ License Text for multiple packages

The following license text applies to multiple packages:

Listing 4: Download: GPLv2+ for multiple packages

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
```

(continues on next page)

(continued from previous page)

```
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
```

(continues on next page)

(continued from previous page)

```
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.
```

(continues on next page)

(continued from previous page)

```
In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
```

(continues on next page)

(continued from previous page)

```
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
```

(continues on next page)

(continued from previous page)

```
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
```

(continues on next page)

(continued from previous page)

```
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
```

(continues on next page)

(continued from previous page)

```
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.5 LGPL21+ License Text for multiple packages

The following license text applies to multiple packages:

Listing 5: Download: LGPL21+ for multiple packages

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
```

(continues on next page)

(continued from previous page)

```
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be

introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it becomes
a de-facto standard.  To achieve this, non-free programs must be
```

(continues on next page)

(continued from previous page)

```
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must

be combined with the library in order to run.

                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
```

(continues on next page)

(continued from previous page)

```
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
```

(continues on next page)

(continued from previous page)

```
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.
```

(continues on next page)

(continued from previous page)

```
  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
```

(continues on next page)

(continued from previous page)

```
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you

distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
```

(continues on next page)

(continued from previous page)

```
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
```

(continues on next page)

(continued from previous page)

```
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded.  In such case, this License incorporates the limitation as if
written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
```

(continues on next page)

(continued from previous page)

```
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.


                     END OF TERMS AND CONDITIONS

           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms of the
ordinary General Public License).

  To apply these terms, attach the following notices to the library.  It is
safest to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James Random Hacker.

  <signature of Ty Coon>, 1 April 1990
```

(continues on next page)

(continued from previous page)

```
  Ty Coon, President of Vice

That's all there is to it!
```

40.6 MIT License Text for multiple packages

The following license text applies to multiple packages:

Listing 6: Download: MIT for multiple packages

```
Copyright <YEAR> <COPYRIGHT HOLDER>

Permission is hereby granted, free of charge, to any person obtaining a copy of
this software and associated documentation files (the "Software"), to deal in
the Software without restriction, including without limitation the rights to
use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of
the Software, and to permit persons to whom the Software is furnished to do so,
subject to the following conditions:

The above copyright notice and this permission notice shall be included in all
copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS
FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR
COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER
IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN
CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

40.7 PD License Text for multiple packages

The following license text applies to multiple packages:

Listing 7: Download: PD for multiple packages

```
This software is in the public domain and does not require a license.

https://en.wikipedia.org/wiki/Public_domain

The public domain consists of all the creative work to which no exclusive
intellectual property rights apply. Those rights may have expired, been
forfeited, expressly waived, or may be inapplicable.
```

40.8 PSFL License Text for multiple packages

The following license text applies to multiple packages:

Listing 8: Download: PSFL for multiple packages

```
A. HISTORY OF THE SOFTWARE
==========================

Python was created in the early 1990s by Guido van Rossum at Stichting
Mathematisch Centrum (CWI, see http://www.cwi.nl) in the Netherlands
as a successor of a language called ABC.  Guido remains Python's
principal author, although it includes many contributions from others.

In 1995, Guido continued his work on Python at the Corporation for
National Research Initiatives (CNRI, see http://www.cnri.reston.va.us)
in Reston, Virginia where he released several versions of the
software.

In May 2000, Guido and the Python core development team moved to
BeOpen.com to form the BeOpen PythonLabs team.  In October of the same
year, the PythonLabs team moved to Digital Creations, which became
Zope Corporation.  In 2001, the Python Software Foundation (PSF, see
https://www.python.org/psf/) was formed, a non-profit organization
created specifically to own Python-related Intellectual Property.
Zope Corporation was a sponsoring member of the PSF.

All Python releases are Open Source (see http://www.opensource.org for
the Open Source Definition).  Historically, most, but not all, Python
releases have also been GPL-compatible; the table below summarizes
the various releases.

    Release         Derived     Year        Owner       GPL-
                    from                                compatible? (1)

    0.9.0 thru 1.2              1991-1995   CWI         yes
    1.3 thru 1.5.2  1.2         1995-1999   CNRI        yes
    1.6             1.5.2       2000        CNRI        no
    2.0             1.6         2000        BeOpen.com  no
    1.6.1           1.6         2001        CNRI        yes (2)
    2.1             2.0+1.6.1   2001        PSF         no
```

(continues on next page)

(continued from previous page)

```
    2.0.1           2.0+1.6.1   2001        PSF         yes
    2.1.1           2.1+2.0.1   2001        PSF         yes
    2.1.2           2.1.1       2002        PSF         yes
    2.1.3           2.1.2       2002        PSF         yes
    2.2 and above   2.1.1       2001-now    PSF         yes

Footnotes:

(1) GPL-compatible doesn't mean that we're distributing Python under
    the GPL.  All Python licenses, unlike the GPL, let you distribute
    a modified version without making your changes open source.  The
    GPL-compatible licenses make it possible to combine Python with
    other software that is released under the GPL; the others don't.

(2) According to Richard Stallman, 1.6.1 is not GPL-compatible,
    because its license has a choice of law clause.  According to
    CNRI, however, Stallman's lawyer has told CNRI's lawyer that 1.6.1
    is "not incompatible" with the GPL.

Thanks to the many outside volunteers who have worked under Guido's
direction to make these releases possible.


B. TERMS AND CONDITIONS FOR ACCESSING OR OTHERWISE USING PYTHON
===============================================================

Python software and documentation are licensed under the
Python Software Foundation License Version 2.

Starting with Python 3.8.6, examples, recipes, and other code in
the documentation are dual licensed under the PSF License Version 2
and the Zero-Clause BSD license.

Some software incorporated into Python is under different licenses.
The licenses are listed with code falling under that license.


PYTHON SOFTWARE FOUNDATION LICENSE VERSION 2
--------------------------------------------

1. This LICENSE AGREEMENT is between the Python Software Foundation
("PSF"), and the Individual or Organization ("Licensee") accessing and
otherwise using this software ("Python") in source or binary form and
its associated documentation.

2. Subject to the terms and conditions of this License Agreement, PSF hereby
grants Licensee a nonexclusive, royalty-free, world-wide license to reproduce,
analyze, test, perform and/or display publicly, prepare derivative works,
distribute, and otherwise use Python alone or in any derivative version,
provided, however, that PSF's License Agreement and PSF's notice of copyright,
i.e., "Copyright (c) 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010,
2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022 Python Software␣
```

(continues on next page)

(continued from previous page)

```
↪Foundation;
All Rights Reserved" are retained in Python alone or in any derivative version
prepared by Licensee.

3. In the event Licensee prepares a derivative work that is based on
or incorporates Python or any part thereof, and wants to make
the derivative work available to others as provided herein, then
Licensee hereby agrees to include in any such work a brief summary of
the changes made to Python.

4. PSF is making Python available to Licensee on an "AS IS"
basis.  PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR
IMPLIED.  BY WAY OF EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND
DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS
FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON WILL NOT
INFRINGE ANY THIRD PARTY RIGHTS.

5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON
FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS
A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON,
OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

6. This License Agreement will automatically terminate upon a material
breach of its terms and conditions.

7. Nothing in this License Agreement shall be deemed to create any
relationship of agency, partnership, or joint venture between PSF and
Licensee.  This License Agreement does not grant permission to use PSF
trademarks or trade name in a trademark sense to endorse or promote
products or services of Licensee, or any third party.

8. By copying, installing or otherwise using Python, Licensee
agrees to be bound by the terms and conditions of this License
Agreement.


BEOPEN.COM LICENSE AGREEMENT FOR PYTHON 2.0
-------------------------------------------

BEOPEN PYTHON OPEN SOURCE LICENSE AGREEMENT VERSION 1

1. This LICENSE AGREEMENT is between BeOpen.com ("BeOpen"), having an
office at 160 Saratoga Avenue, Santa Clara, CA 95051, and the
Individual or Organization ("Licensee") accessing and otherwise using
this software in source or binary form and its associated
documentation ("the Software").

2. Subject to the terms and conditions of this BeOpen Python License
Agreement, BeOpen hereby grants Licensee a non-exclusive,
royalty-free, world-wide license to reproduce, analyze, test, perform
and/or display publicly, prepare derivative works, distribute, and
otherwise use the Software alone or in any derivative version,
```

(continues on next page)

(continued from previous page)

```
provided, however, that the BeOpen Python License is retained in the
Software, alone or in any derivative version prepared by Licensee.

3. BeOpen is making the Software available to Licensee on an "AS IS"
basis.  BEOPEN MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR
IMPLIED.  BY WAY OF EXAMPLE, BUT NOT LIMITATION, BEOPEN MAKES NO AND
DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS
FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE WILL NOT
INFRINGE ANY THIRD PARTY RIGHTS.

4. BEOPEN SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF THE
SOFTWARE FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS
AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THE SOFTWARE, OR ANY
DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

5. This License Agreement will automatically terminate upon a material
breach of its terms and conditions.

6. This License Agreement shall be governed by and interpreted in all
respects by the law of the State of California, excluding conflict of
law provisions.  Nothing in this License Agreement shall be deemed to
create any relationship of agency, partnership, or joint venture
between BeOpen and Licensee.  This License Agreement does not grant
permission to use BeOpen trademarks or trade names in a trademark
sense to endorse or promote products or services of Licensee, or any
third party.  As an exception, the "BeOpen Python" logos available at
http://www.pythonlabs.com/logos.html may be used according to the
permissions granted on that web page.

7. By copying, installing or otherwise using the software, Licensee
agrees to be bound by the terms and conditions of this License
Agreement.


CNRI LICENSE AGREEMENT FOR PYTHON 1.6.1
---------------------------------------

1. This LICENSE AGREEMENT is between the Corporation for National
Research Initiatives, having an office at 1895 Preston White Drive,
Reston, VA 20191 ("CNRI"), and the Individual or Organization
("Licensee") accessing and otherwise using Python 1.6.1 software in
source or binary form and its associated documentation.

2. Subject to the terms and conditions of this License Agreement, CNRI
hereby grants Licensee a nonexclusive, royalty-free, world-wide
license to reproduce, analyze, test, perform and/or display publicly,
prepare derivative works, distribute, and otherwise use Python 1.6.1
alone or in any derivative version, provided, however, that CNRI's
License Agreement and CNRI's notice of copyright, i.e., "Copyright (c)
1995-2001 Corporation for National Research Initiatives; All Rights
Reserved" are retained in Python 1.6.1 alone or in any derivative
version prepared by Licensee.  Alternately, in lieu of CNRI's License
```

(continues on next page)

(continued from previous page)

```
Agreement, Licensee may substitute the following text (omitting the
quotes): "Python 1.6.1 is made available subject to the terms and
conditions in CNRI's License Agreement.  This Agreement together with
Python 1.6.1 may be located on the internet using the following
unique, persistent identifier (known as a handle): 1895.22/1013.  This
Agreement may also be obtained from a proxy server on the internet
using the following URL: http://hdl.handle.net/1895.22/1013".

3. In the event Licensee prepares a derivative work that is based on
or incorporates Python 1.6.1 or any part thereof, and wants to make
the derivative work available to others as provided herein, then
Licensee hereby agrees to include in any such work a brief summary of
the changes made to Python 1.6.1.

4. CNRI is making Python 1.6.1 available to Licensee on an "AS IS"
basis.  CNRI MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR
IMPLIED.  BY WAY OF EXAMPLE, BUT NOT LIMITATION, CNRI MAKES NO AND
DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS
FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON 1.6.1 WILL NOT
INFRINGE ANY THIRD PARTY RIGHTS.

5. CNRI SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON
1.6.1 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS
A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON 1.6.1,
OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

6. This License Agreement will automatically terminate upon a material
breach of its terms and conditions.

7. This License Agreement shall be governed by the federal
intellectual property law of the United States, including without
limitation the federal copyright law, and, to the extent such
U.S. federal law does not apply, by the law of the Commonwealth of
Virginia, excluding Virginia's conflict of law provisions.
Notwithstanding the foregoing, with regard to derivative works based
on Python 1.6.1 that incorporate non-separable material that was
previously distributed under the GNU General Public License (GPL), the
law of the Commonwealth of Virginia shall govern this License
Agreement only as to issues arising under or with respect to
Paragraphs 4, 5, and 7 of this License Agreement.  Nothing in this
License Agreement shall be deemed to create any relationship of
agency, partnership, or joint venture between CNRI and Licensee.  This
License Agreement does not grant permission to use CNRI trademarks or
trade name in a trademark sense to endorse or promote products or
services of Licensee, or any third party.

8. By clicking on the "ACCEPT" button where indicated, or by copying,
installing or otherwise using Python 1.6.1, Licensee agrees to be
bound by the terms and conditions of this License Agreement.

        ACCEPT
```

(continues on next page)

(continued from previous page)

```
CWI LICENSE AGREEMENT FOR PYTHON 0.9.0 THROUGH 1.2
--------------------------------------------------

Copyright (c) 1991 - 1995, Stichting Mathematisch Centrum Amsterdam,
The Netherlands.  All rights reserved.

Permission to use, copy, modify, and distribute this software and its
documentation for any purpose and without fee is hereby granted,
provided that the above copyright notice appear in all copies and that
both that copyright notice and this permission notice appear in
supporting documentation, and that the name of Stichting Mathematisch
Centrum or CWI not be used in advertising or publicity pertaining to
distribution of the software without specific, written prior
permission.

STICHTING MATHEMATISCH CENTRUM DISCLAIMS ALL WARRANTIES WITH REGARD TO
THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND
FITNESS, IN NO EVENT SHALL STICHTING MATHEMATISCH CENTRUM BE LIABLE
FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT
OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

ZERO-CLAUSE BSD LICENSE FOR CODE IN THE PYTHON DOCUMENTATION
----------------------------------------------------------------------

Permission to use, copy, modify, and/or distribute this software for any
purpose with or without fee is hereby granted.

THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH
REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY
AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY SPECIAL, DIRECT,
INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM
LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR
OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR
PERFORMANCE OF THIS SOFTWARE.
```

40.9 MPL20 License Text for bind-tools

The following license text applies to bind-tools:

Listing 9: Download: `MPL20` for bind-tools

```
Copyright (C) 1996-2021  Internet Systems Consortium, Inc. ("ISC")

This Source Code Form is subject to the terms of the Mozilla Public
License, v. 2.0. If a copy of the MPL was not distributed with this
file, you can obtain one at https://mozilla.org/MPL/2.0/.
```

(continues on next page)

(continued from previous page)

```
 --------------------------------------------------------------------------------

        Portions of this code release fall under one or more of the
        following Copyright notices.  Please see individual source
        files for details.

        For binary releases also see: OpenSSL-LICENSE.

Copyright (C) 1996-2001  Nominum, Inc.

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINUM DISCLAIMS ALL WARRANTIES
WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NOMINUM BE LIABLE FOR
ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT
OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

 --------------------------------------------------------------------------------

Copyright (C) 1995-2000 by Network Associates, Inc.

Permission to use, copy, modify, and/or distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC AND NETWORK ASSOCIATES DISCLAIMS
ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS.  IN NO EVENT SHALL ISC BE LIABLE
FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR
IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

 --------------------------------------------------------------------------------

Copyright (C) 2002 Stichting NLnet, Netherlands, stichting@nlnet.nl.

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the
above copyright notice and this permission notice appear in all
copies.

THE SOFTWARE IS PROVIDED "AS IS" AND STICHTING NLNET
DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL
STICHTING NLNET BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR
CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS
```

(continues on next page)

(continued from previous page)

```
OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE
OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE
USE OR PERFORMANCE OF THIS SOFTWARE.

The development of Dynamically Loadable Zones (DLZ) for Bind 9 was
conceived and contributed by Rob Butler.

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the
above copyright notice and this permission notice appear in all
copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ROB BUTLER
DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL
ROB BUTLER BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR
CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS
OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE
OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE
USE OR PERFORMANCE OF THIS SOFTWARE.

 ---------------------------------------------------------------------------

Copyright (c) 1987, 1990, 1993, 1994
     The Regents of the University of California.  All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors
   may be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.

 ---------------------------------------------------------------------------
```

(continues on next page)

(continued from previous page)

```
Copyright (C) The Internet Society 2005.  This version of
this module is part of RFC 4178; see the RFC itself for
full legal notices.

(The above copyright notice is per RFC 3978 5.6 (a), q.v.)

 ------------------------------------------------------------------------------

Copyright (c) 2004 Masarykova universita
(Masaryk University, Brno, Czech Republic)
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice,
   this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.

3. Neither the name of the University nor the names of its contributors may
   be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.

 ------------------------------------------------------------------------------

Copyright (c) 1997 - 2003 Kungliga Tekniska Hgskolan
(Royal Institute of Technology, Stockholm, Sweden).
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright
```

(continues on next page)

(continued from previous page)

```
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.

3. Neither the name of the Institute nor the names of its contributors
   may be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INSTITUTE AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE INSTITUTE OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.

 --------------------------------------------------------------------------------

Copyright (c) 1993 by Digital Equipment Corporation.

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies, and that
the name of Digital Equipment Corporation not be used in advertising or
publicity pertaining to distribution of the document or software without
specific, written prior permission.

THE SOFTWARE IS PROVIDED "AS IS" AND DIGITAL EQUIPMENT CORP. DISCLAIMS ALL
WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS.   IN NO EVENT SHALL DIGITAL EQUIPMENT
CORPORATION BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL
DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR
PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS
ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS
SOFTWARE.

 --------------------------------------------------------------------------------

Copyright (C) 1995, 1996, 1997, and 1998 WIDE Project.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
```

(continues on next page)

(continued from previous page)

```
3. Neither the name of the project nor the names of its contributors
   may be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE PROJECT OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.

 --------------------------------------------------------------------------------

Copyright (c) 1999-2000 by Nortel Networks Corporation

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NORTEL NETWORKS DISCLAIMS
ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NORTEL NETWORKS
BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES
OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS,
WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION,
ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS
SOFTWARE.

 --------------------------------------------------------------------------------

Copyright (C) 2004  Nominet, Ltd.

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINET DISCLAIMS ALL WARRANTIES WITH
REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY
AND FITNESS.  IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT,
INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM
LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE
OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR
PERFORMANCE OF THIS SOFTWARE.

 --------------------------------------------------------------------------------

Copyright (c) 1996, David Mazieres <dm@uun.org>
```

(continues on next page)

(continued from previous page)

```
Copyright (c) 2008, Damien Miller <djm@openbsd.org>

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES
WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR
ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF
OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

--------------------------------------------------------------------------------

Copyright (c) 1995, 1997, 1998 The NetBSD Foundation, Inc.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS
``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS
BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------------------

Copyright (C) 2008-2011  Red Hat, Inc.

Permission to use, copy, modify, and/or distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND Red Hat DISCLAIMS ALL WARRANTIES WITH
REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY
AND FITNESS.  IN NO EVENT SHALL Red Hat BE LIABLE FOR ANY SPECIAL, DIRECT,
INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM
```

(continues on next page)

(continued from previous page)

```
LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE
OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR
PERFORMANCE OF THIS SOFTWARE.

--------------------------------------------------------------------------------

Copyright (c) 2013-2014, Farsight Security, Inc.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1. Redistributions of source code must retain the above copyright
notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in the
documentation and/or other materials provided with the distribution.

3. Neither the name of the copyright holder nor the names of its
contributors may be used to endorse or promote products derived from
this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR
CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;
OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF
ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------------------

Copyright (c) 2014 by Farsight Security, Inc.

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

   http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```

40.10 MPL20 License Text for ca_root_nss

The following license text applies to ca_root_nss:

Listing 10: Download: MPL20 for ca_root_nss

```
NSS is available under the Mozilla Public License, version 2, a copy of which
is below.

Note on GPL Compatibility
-------------------------

The MPL 2, section 3.3, permits you to combine NSS with code under the GNU
General Public License (GPL) version 2, or any later version of that
license, to make a Larger Work, and distribute the result under the GPL.
The only condition is that you must also make NSS, and any changes you
have made to it, available to recipients under the terms of the MPL 2 also.

Anyone who receives the combined code from you does not have to continue
to dual licence in this way, and may, if they wish, distribute under the
terms of either of the two licences - either the MPL alone or the GPL
alone. However, we discourage people from distributing copies of NSS under
the GPL alone, because it means that any improvements they make cannot be
reincorporated into the main version of NSS. There is never a need to do
this for license compatibility reasons.

Note on LGPL Compatibility
--------------------------

The above also applies to combining MPLed code in a single library with
code under the GNU Lesser General Public License (LGPL) version 2.1, or
any later version of that license. If the LGPLed code and the MPLed code
are not in the same library, then the copyleft coverage of the two
licences does not overlap, so no issues arise.


Mozilla Public License Version 2.0
==================================

1. Definitions
--------------

1.1. "Contributor"
    means each individual or legal entity that creates, contributes to
    the creation of, or owns Covered Software.

1.2. "Contributor Version"
    means the combination of the Contributions of others (if any) used
    by a Contributor and that particular Contributor's Contribution.

1.3. "Contribution"
    means Covered Software of a particular Contributor.
```

(continues on next page)

(continued from previous page)

```
1.4. "Covered Software"
    means Source Code Form to which the initial Contributor has attached
    the notice in Exhibit A, the Executable Form of such Source Code
    Form, and Modifications of such Source Code Form, in each case
    including portions thereof.

1.5. "Incompatible With Secondary Licenses"
    means

    (a) that the initial Contributor has attached the notice described
        in Exhibit B to the Covered Software; or

    (b) that the Covered Software was made available under the terms of
        version 1.1 or earlier of the License, but not also under the
        terms of a Secondary License.

1.6. "Executable Form"
    means any form of the work other than Source Code Form.

1.7. "Larger Work"
    means a work that combines Covered Software with other material, in
    a separate file or files, that is not Covered Software.

1.8. "License"
    means this document.

1.9. "Licensable"
    means having the right to grant, to the maximum extent possible,
    whether at the time of the initial grant or subsequently, any and
    all of the rights conveyed by this License.

1.10. "Modifications"
    means any of the following:

    (a) any file in Source Code Form that results from an addition to,
        deletion from, or modification of the contents of Covered
        Software; or

    (b) any new file in Source Code Form that contains any Covered
        Software.

1.11. "Patent Claims" of a Contributor
    means any patent claim(s), including without limitation, method,
    process, and apparatus claims, in any patent Licensable by such
    Contributor that would be infringed, but for the grant of the
    License, by the making, using, selling, offering for sale, having
    made, import, or transfer of either its Contributions or its
    Contributor Version.

1.12. "Secondary License"
    means either the GNU General Public License, Version 2.0, the GNU
    Lesser General Public License, Version 2.1, the GNU Affero General
```

(continues on next page)

(continued from previous page)

```
    Public License, Version 3.0, or any later versions of those
    licenses.

1.13. "Source Code Form"
    means the form of the work preferred for making modifications.

1.14. "You" (or "Your")
    means an individual or a legal entity exercising rights under this
    License. For legal entities, "You" includes any entity that
    controls, is controlled by, or is under common control with You. For
    purposes of this definition, "control" means (a) the power, direct
    or indirect, to cause the direction or management of such entity,
    whether by contract or otherwise, or (b) ownership of more than
    fifty percent (50%) of the outstanding shares or beneficial
    ownership of such entity.

2. License Grants and Conditions
--------------------------------

2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free,
non-exclusive license:

(a) under intellectual property rights (other than patent or trademark)
    Licensable by such Contributor to use, reproduce, make available,
    modify, display, perform, distribute, and otherwise exploit its
    Contributions, either on an unmodified basis, with Modifications, or
    as part of a Larger Work; and

(b) under Patent Claims of such Contributor to make, use, sell, offer
    for sale, have made, import, and otherwise transfer either its
    Contributions or its Contributor Version.

2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution
become effective for each Contribution on the date the Contributor first
distributes such Contribution.

2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under
this License. No additional rights or licenses will be implied from the
distribution or licensing of Covered Software under this License.
Notwithstanding Section 2.1(b) above, no patent license is granted by a
Contributor:

(a) for any code that a Contributor has removed from Covered Software;
    or

(b) for infringements caused by: (i) Your and any other third party's
```

(continues on next page)

(continued from previous page)

```
    modifications of Covered Software, or (ii) the combination of its
    Contributions with other software (except as part of its Contributor
    Version); or

(c) under Patent Claims infringed by Covered Software in the absence of
    its Contributions.

This License does not grant any rights in the trademarks, service marks,
or logos of any Contributor (except as may be necessary to comply with
the notice requirements in Section 3.4).

2.4. Subsequent Licenses

No Contributor makes additional grants as a result of Your choice to
distribute the Covered Software under a subsequent version of this
License (see Section 10.2) or under the terms of a Secondary License (if
permitted under the terms of Section 3.3).

2.5. Representation

Each Contributor represents that the Contributor believes its
Contributions are its original creation(s) or it has sufficient rights
to grant the rights to its Contributions conveyed by this License.

2.6. Fair Use

This License is not intended to limit any rights You have under
applicable copyright doctrines of fair use, fair dealing, or other
equivalents.

2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted
in Section 2.1.

3. Responsibilities
-------------------

3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any
Modifications that You create or to which You contribute, must be under
the terms of this License. You must inform recipients that the Source
Code Form of the Covered Software is governed by the terms of this
License, and how they can obtain a copy of this License. You may not
attempt to alter or restrict the recipients' rights in the Source Code
Form.

3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:
```

(continues on next page)

(continued from previous page)

```
(a) such Covered Software must also be made available in Source Code
    Form, as described in Section 3.1, and You must inform recipients of
    the Executable Form how they can obtain a copy of such Source Code
    Form by reasonable means in a timely manner, at a charge no more
    than the cost of distribution to the recipient; and

(b) You may distribute such Executable Form under the terms of this
    License, or sublicense it under different terms, provided that the
    license for the Executable Form does not attempt to limit or alter
    the recipients' rights in the Source Code Form under this License.

3.3. Distribution of a Larger Work

You may create and distribute a Larger Work under terms of Your choice,
provided that You also comply with the requirements of this License for
the Covered Software. If the Larger Work is a combination of Covered
Software with a work governed by one or more Secondary Licenses, and the
Covered Software is not Incompatible With Secondary Licenses, this
License permits You to additionally distribute such Covered Software
under the terms of such Secondary License(s), so that the recipient of
the Larger Work may, at their option, further distribute the Covered
Software under the terms of either this License or such Secondary
License(s).

3.4. Notices

You may not remove or alter the substance of any license notices
(including copyright notices, patent notices, disclaimers of warranty,
or limitations of liability) contained within the Source Code Form of
the Covered Software, except that You may alter any license notices to
the extent required to remedy known factual inaccuracies.

3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support,
indemnity or liability obligations to one or more recipients of Covered
Software. However, You may do so only on Your own behalf, and not on
behalf of any Contributor. You must make it absolutely clear that any
such warranty, support, indemnity, or liability obligation is offered by
You alone, and You hereby agree to indemnify every Contributor for any
liability incurred by such Contributor as a result of warranty, support,
indemnity or liability terms You offer. You may include additional
disclaimers of warranty and limitations of liability specific to any
jurisdiction.

4. Inability to Comply Due to Statute or Regulation
---------------------------------------------------

If it is impossible for You to comply with any of the terms of this
License with respect to some or all of the Covered Software due to
statute, judicial order, or regulation then You must: (a) comply with
the terms of this License to the maximum extent possible; and (b)
```

(continues on next page)

(continued from previous page)

```
describe the limitations and the code they affect. Such description must
be placed in a text file included with all distributions of the Covered
Software under this License. Except to the extent prohibited by statute
or regulation, such description must be sufficiently detailed for a
recipient of ordinary skill to be able to understand it.

5. Termination
--------------

5.1. The rights granted under this License will terminate automatically
if You fail to comply with any of its terms. However, if You become
compliant, then the rights granted under this License from a particular
Contributor are reinstated (a) provisionally, unless and until such
Contributor explicitly and finally terminates Your grants, and (b) on an
ongoing basis, if such Contributor fails to notify You of the
non-compliance by some reasonable means prior to 60 days after You have
come back into compliance. Moreover, Your grants from a particular
Contributor are reinstated on an ongoing basis if such Contributor
notifies You of the non-compliance by some reasonable means, this is the
first time You have received notice of non-compliance with this License
from such Contributor, and You become compliant prior to 30 days after
Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent
infringement claim (excluding declaratory judgment actions,
counter-claims, and cross-claims) alleging that a Contributor Version
directly or indirectly infringes any patent, then the rights granted to
You by any and all Contributors for the Covered Software under Section
2.1 of this License shall terminate.

5.3. In the event of termination under Sections 5.1 or 5.2 above, all
end user license agreements (excluding distributors and resellers) which
have been validly granted by You or Your distributors under this License
prior to termination shall survive termination.

************************************************************************
*                                                                      *
*  6. Disclaimer of Warranty                                           *
*  -------------------------                                           *
*                                                                      *
*  Covered Software is provided under this License on an "as is"       *
*  basis, without warranty of any kind, either expressed, implied, or  *
*  statutory, including, without limitation, warranties that the       *
*  Covered Software is free of defects, merchantable, fit for a        *
*  particular purpose or non-infringing. The entire risk as to the     *
*  quality and performance of the Covered Software is with You.        *
*  Should any Covered Software prove defective in any respect, You     *
*  (not any Contributor) assume the cost of any necessary servicing,   *
*  repair, or correction. This disclaimer of warranty constitutes an   *
*  essential part of this License. No use of any Covered Software is   *
*  authorized under this License except under this disclaimer.         *
*                                                                      *
```

(continues on next page)

(continued from previous page)

```
****************************************************************************

****************************************************************************
*                                                                          *
*  7. Limitation of Liability                                              *
*  --------------------------                                              *
*                                                                          *
*  Under no circumstances and under no legal theory, whether tort          *
*  (including negligence), contract, or otherwise, shall any               *
*  Contributor, or anyone who distributes Covered Software as              *
*  permitted above, be liable to You for any direct, indirect,             *
*  special, incidental, or consequential damages of any character          *
*  including, without limitation, damages for lost profits, loss of        *
*  goodwill, work stoppage, computer failure or malfunction, or any        *
*  and all other commercial damages or losses, even if such party          *
*  shall have been informed of the possibility of such damages. This       *
*  limitation of liability shall not apply to liability for death or       *
*  personal injury resulting from such party's negligence to the           *
*  extent applicable law prohibits such limitation. Some                   *
*  jurisdictions do not allow the exclusion or limitation of               *
*  incidental or consequential damages, so this exclusion and              *
*  limitation may not apply to You.                                        *
*                                                                          *
****************************************************************************

8. Litigation
-------------

Any litigation relating to this License may be brought only in the
courts of a jurisdiction where the defendant maintains its principal
place of business and such litigation shall be governed by laws of that
jurisdiction, without reference to its conflict-of-law provisions.
Nothing in this Section shall prevent a party's ability to bring
cross-claims or counter-claims.

9. Miscellaneous
----------------

This License represents the complete agreement concerning the subject
matter hereof. If any provision of this License is held to be
unenforceable, such provision shall be reformed only to the extent
necessary to make it enforceable. Any law or regulation which provides
that the language of a contract shall be construed against the drafter
shall not be used to construe this License against a Contributor.

10. Versions of the License
---------------------------

10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section
10.3, no one other than the license steward has the right to modify or
```

(continues on next page)

(continued from previous page)

```
publish new versions of this License. Each version will be given a
distinguishing version number.

10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version
of the License under which You originally received the Covered Software,
or under the terms of any subsequent version published by the license
steward.

10.3. Modified Versions

If you create software not governed by this License, and you want to
create a new license for such software, you may create and use a
modified version of this License if you rename the license and remove
any references to the name of the license steward (except to note that
such modified license differs from this License).

10.4. Distributing Source Code Form that is Incompatible With Secondary
Licenses

If You choose to distribute Source Code Form that is Incompatible With
Secondary Licenses under the terms of this version of the License, the
notice described in Exhibit B of this License must be attached.

Exhibit A - Source Code Form License Notice
-------------------------------------------

  This Source Code Form is subject to the terms of the Mozilla Public
  License, v. 2.0. If a copy of the MPL was not distributed with this
  file, You can obtain one at http://mozilla.org/MPL/2.0/.

If it is not possible or desirable to put the notice in a particular
file, then You may include the notice in a location (such as a LICENSE
file in a relevant directory) where a recipient would be likely to look
for such a notice.

You may add additional accurate notices of copyright ownership.

Exhibit B - "Incompatible With Secondary Licenses" Notice
---------------------------------------------------------

  This Source Code Form is "Incompatible With Secondary Licenses", as
  defined by the Mozilla Public License, v. 2.0.
```

40.11 LGPL21 License Text for ccid

The following license text applies to ccid:

Listing 11: Download: LGPL21 for ccid

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
        51 Franklin St, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations
below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.
```

(continues on next page)

(continued from previous page)

```
  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be

introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it
becomes a de-facto standard.  To achieve this, non-free programs must
be allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
```

(continues on next page)

(continued from previous page)

```
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must

be combined with the library in order to run.

                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control
compilation and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.
```

(continues on next page)

(continued from previous page)

```
  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.
```

(continues on next page)

(continued from previous page)

```
Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.
```

(continues on next page)

(continued from previous page)

```
  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
```

(continues on next page)

(continued from previous page)

```
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at least
    three years, to give the same user the materials specified in
    Subsection 6a, above, for a charge no more than the cost of
    performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you

distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
```

(continues on next page)

(continued from previous page)

```
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply, and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
```

(continues on next page)

(continued from previous page)

```
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License
may add an explicit geographical distribution limitation excluding those
countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
```

(continues on next page)

(continued from previous page)

```
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.


                     END OF TERMS AND CONDITIONS

           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms
of the ordinary General Public License).

  To apply these terms, attach the following notices to the library.
It is safest to attach them to the start of each source file to most
effectively convey the exclusion of warranty; and each file should
have at least the "copyright" line and a pointer to where the full
notice is found.


    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA  02110-1301  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or
your school, if any, to sign a "copyright disclaimer" for the library,
if necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James
  Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice
```

(continues on next page)

(continued from previous page)

```
That's all there is to it!
```

40.12 MIT License Text for curl

The following license text applies to curl:

Listing 12: Download: MIT for curl

```
COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996 - 2021, Daniel Stenberg, <daniel@haxx.se>, and many
contributors, see the THANKS file.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose
with or without fee is hereby granted, provided that the above copyright
notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN
NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM,
DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR
OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE
OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not
be used in advertising or otherwise to promote the sale, use or other dealings
in this Software without prior written authorization of the copyright holder.
```

40.13 BSD4CLAUSE License Text for cyrus-sasl

The following license text applies to cyrus-sasl:

Listing 13: Download: BSD4CLAUSE for cyrus-sasl

```
/* CMU libsasl
 * Tim Martin
 * Rob Earhart
 * Rob Siemborski
 */
/*
 * Copyright (c) 1998-2003 Carnegie Mellon University.  All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
```

(continues on next page)

(continued from previous page)

```
*
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in
*    the documentation and/or other materials provided with the
*    distribution.
*
* 3. The name "Carnegie Mellon University" must not be used to
*    endorse or promote products derived from this software without
*    prior written permission. For permission or any other legal
*    details, please contact
*      Office of Technology Transfer
*      Carnegie Mellon University
*      5000 Forbes Avenue
*      Pittsburgh, PA  15213-3890
*      (412) 268-4387, fax: (412) 268-7395
*      tech-transfer@andrew.cmu.edu
*
* 4. Redistributions of any form whatsoever must retain the following
*    acknowledgment:
*    "This product includes software developed by Computing Services
*     at Carnegie Mellon University (http://www.cmu.edu/computing/)."
*
* CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO
* THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY
* AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE
* FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
* WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN
* AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING
* OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.
*/
```

40.14 GPLv2 License Text for dbus

The following license text applies to dbus:

Listing 14: Download: GPLv2 for dbus

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

                            Preamble
```

(continues on next page)

(continued from previous page)

```
  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
```

(continues on next page)

(continued from previous page)

```
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
```

(continues on next page)

(continued from previous page)

```
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
```

(continues on next page)

(continued from previous page)

```
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
```

(continues on next page)

(continued from previous page)

```
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
```

(continues on next page)

(continued from previous page)

```
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
```

(continues on next page)

(continued from previous page)

```
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.15 EULA License Text for devcpu-data-amd

The following license text applies to devcpu-data-amd:

Listing 15: Download: EULA for devcpu-data-amd

```
Copyright (C) 2010-2018 Advanced Micro Devices, Inc., All rights reserved.

Permission is hereby granted by Advanced Micro Devices, Inc. ("AMD"),
free of any license fees, to any person obtaining a copy of this
microcode in binary form (the "Software") ("You"), to install,
reproduce, copy and distribute copies of the Software and to permit
persons to whom the Software is provided to do the same, subject to
the following terms and conditions.  Your use of any portion of the
Software shall constitute Your acceptance of the following terms and
conditions. If You do not agree to the following terms and conditions,
do not use, retain or redistribute any portion of the Software.

If You redistribute this Software, You must reproduce the above
copyright notice and this license with the Software.
Without specific, prior, written permission from AMD, You may not
reference AMD or AMD products in the promotion of any product derived
from or incorporating this Software in any manner that implies that
AMD endorses or has certified such product derived from or
incorporating this Software.

You may not reverse engineer, decompile, or disassemble this Software
or any portion thereof.

THE SOFTWARE IS PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED
WARRANTY OF ANY KIND, INCLUDING BUT NOT LIMITED TO WARRANTIES OF
MERCHANTABILITY, NONINFRINGEMENT, TITLE, FITNESS FOR ANY PARTICULAR
```

(continues on next page)

(continued from previous page)

```
PURPOSE, OR WARRANTIES ARISING FROM CONDUCT, COURSE OF DEALING, OR
USAGE OF TRADE. IN NO EVENT SHALL AMD OR ITS LICENSORS BE LIABLE FOR
ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR
LOSS OF PROFITS, BUSINESS INTERRUPTION, OR LOSS OF DATA OR
INFORMATION) ARISING OUT OF AMD'S NEGLIGENCE, GROSS NEGLIGENCE, THE
USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF AMD HAS BEEN ADVISED
OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME JURISDICTIONS
PROHIBIT THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR
INCIDENTAL DAMAGES OR THE EXCLUSION OF IMPLIED WARRANTIES, THE ABOVE
LIMITATION MAY NOT APPLY TO YOU.

Without limiting the foregoing, the Software may implement third party
technologies for which You must obtain licenses from parties other
than AMD. You agree that AMD has not obtained or conveyed to You, and
that You shall be responsible for obtaining the rights to use and/or
distribute the applicable underlying intellectual property rights
related to the third party technologies. These third party
technologies are not licensed hereunder.

If You use the Software (in whole or in part), You shall adhere to all
applicable U.S., European, and other export laws, including but not
limited to the U.S. Export Administration Regulations ("EAR"), (15
C.F.R. Sections 730 through 774), and E.U. Council Regulation (EC) No
1334/2000 of 22 June 2000. Further, pursuant to Section 740.6  of the
EAR, You hereby certify that, except pursuant to a license granted by
the United States Department of Commerce Bureau of Industry and
Security or as otherwise permitted pursuant to a License Exception
under the U.S. Export Administration Regulations ("EAR"), You will not
(1) export, re-export or release to a national of a country in Country
Groups D:1, E:1 or E:2 any restricted technology, software, or source
code You receive hereunder, or (2) export to Country Groups D:1, E:1
or E:2 the direct product of such technology or software, if such
foreign produced direct product is subject to national security
controls as identified on the Commerce Control List (currently found
in Supplement 1 to Part 774 of EAR). For the most current Country
Group listings, or for additional information about the EAR or Your
obligations under those regulations, please refer to the U.S. Bureau
of Industry and Security?s website at ttp://www.bis.doc.gov/.
```

40.16 EULA License Text for devcpu-data-intel

The following license text applies to devcpu-data-intel:

Listing 16: Download: EULA for devcpu-data-intel

```
The terms of the software license agreement included with any software you download will␣
↪control your use of the software.

INTEL SOFTWARE LICENSE AGREEMENT
```

(continues on next page)

(continued from previous page)

```
IMPORTANT - READ BEFORE COPYING, INSTALLING OR USING.
Do not use or load this software and any associated materials (collectively,
the "Software") until you have carefully read the following terms and
conditions. By loading or using the Software, you agree to the terms of this
Agreement. If you do not wish to so agree, do not install or use the Software.

LICENSES: Please Note:
- If you are a network administrator, the "Site License" below shall
apply to you.
- If you are an end user, the "Single User License" shall apply to you.
- If you are an original equipment manufacturer (OEM), the "OEM License"
shall apply to you.

SITE LICENSE. You may copy the Software onto your organization's computers
for your organization's use, and you may make a reasonable number of
back-up copies of the Software, subject to these conditions:

1. This Software is licensed for use only in conjunction with Intel
component products. Use of the Software in conjunction with non-Intel
component products is not licensed hereunder.
2. You may not copy, modify, rent, sell, distribute or transfer any part
of the Software except as provided in this Agreement, and you agree to
prevent unauthorized copying of the Software.
3. You may not reverse engineer, decompile, or disassemble the Software.
4. You may not sublicense or permit simultaneous use of the Software by
more than one user.
5. The Software may include portions offered on terms in addition to those
set out here, as set out in a license accompanying those portions.

SINGLE USER LICENSE. You may copy the Software onto a single computer for
your personal, noncommercial use, and you may make one back-up copy of the
Software, subject to these conditions:

1. This Software is licensed for use only in conjunction with Intel
component products. Use of the Software in conjunction with non-Intel
component products is not licensed hereunder.
2. You may not copy, modify, rent, sell, distribute or transfer any part
of the Software except as provided in this Agreement, and you agree to
prevent unauthorized copying of the Software.
3. You may not reverse engineer, decompile, or disassemble the Software.
4. You may not sublicense or permit simultaneous use of the Software by
more than one user.
5. The Software may include portions offered on terms in addition to those
set out here, as set out in a license accompanying those portions.

OEM LICENSE: You may reproduce and distribute the Software only as an
integral part of or incorporated in Your product or as a standalone
Software maintenance update for existing end users of Your products,
excluding any other standalone products, subject to these conditions:

1. This Software is licensed for use only in conjunction with Intel
component products. Use of the Software in conjunction with non-Intel
```

(continues on next page)

(continued from previous page)

```
component products is not licensed hereunder.
2. You may not copy, modify, rent, sell, distribute or transfer any part
of the Software except as provided in this Agreement, and you agree to
prevent unauthorized copying of the Software.
3. You may not reverse engineer, decompile, or disassemble the Software.
4. You may only distribute the Software to your customers pursuant to a
written license agreement. Such license agreement may be a "break-the-
seal" license agreement. At a minimum such license shall safeguard
Intel's ownership rights to the Software.
5. The Software may include portions offered on terms in addition to those
set out here, as set out in a license accompanying those portions.

NO OTHER RIGHTS. No rights or licenses are granted by Intel to You, expressly
or by implication, with respect to any proprietary information or patent,
copyright, mask work, trademark, trade secret, or other intellectual property
right owned or controlled by Intel, except as expressly provided in this
Agreement.

OWNERSHIP OF SOFTWARE AND COPYRIGHTS. Title to all copies of the Software
remains with Intel or its suppliers. The Software is copyrighted and
protected by the laws of the United States and other countries, and
international treaty provisions. You may not remove any copyright notices
from the Software. Intel may make changes to the Software, or to items
referenced therein, at any time without notice, but is not obligated to
support or update the Software. Except as otherwise expressly provided, Intel
grants no express or implied right under Intel patents, copyrights,
trademarks, or other intellectual property rights. You may transfer the
Software only if the recipient agrees to be fully bound by these terms and if
you retain no copies of the Software.

LIMITED MEDIA WARRANTY. If the Software has been delivered by Intel on
physical media, Intel warrants the media to be free from material physical
defects for a period of ninety days after delivery by Intel. If such a defect
is found, return the media to Intel for replacement or alternate delivery of
the Software as Intel may select.

EXCLUSION OF OTHER WARRANTIES. EXCEPT AS PROVIDED ABOVE, THE SOFTWARE IS
PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND
INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A
PARTICULAR PURPOSE. Intel does not warrant or assume responsibility for the
accuracy or completeness of any information, text, graphics, links or other
items contained within the Software.

LIMITATION OF LIABILITY. IN NO EVENT SHALL INTEL OR ITS SUPPLIERS BE LIABLE
FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, LOST PROFITS,
BUSINESS INTERRUPTION, OR LOST INFORMATION) ARISING OUT OF THE USE OF OR
INABILITY TO USE THE SOFTWARE, EVEN IF INTEL HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS PROHIBIT EXCLUSION OR
LIMITATION OF LIABILITY FOR IMPLIED WARRANTIES OR CONSEQUENTIAL OR INCIDENTAL
DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. YOU MAY ALSO HAVE
OTHER LEGAL RIGHTS THAT VARY FROM JURISDICTION TO JURISDICTION.
```

(continues on next page)

(continued from previous page)

```
TERMINATION OF THIS AGREEMENT. Intel may terminate this Agreement at any time
if you violate its terms. Upon termination, you will immediately destroy the
Software or return all copies of the Software to Intel.

APPLICABLE LAWS. Claims arising under this Agreement shall be governed by the
laws of California, excluding its principles of conflict of laws and the
United Nations Convention on Contracts for the Sale of Goods. You may not
export the Software in violation of applicable export laws and regulations.
Intel is not obligated under any other agreements unless they are in writing
and signed by an authorized representative of Intel.

GOVERNMENT RESTRICTED RIGHTS. The Software is provided with "RESTRICTED
RIGHTS." Use, duplication, or disclosure by the Government is subject to
restrictions as set forth in FAR52.227-14 and DFAR252.227-7013 et seq. or its
successor. Use of the Software by the Government constitutes acknowledgment
of Intel's proprietary rights therein. Contractor or Manufacturer is Intel
2200 Mission College Blvd., Santa Clara, CA 95052.
```

40.17 BSD3CLAUSE License Text for dhcp6

The following license text applies to dhcp6:

Listing 17: Download: BSD3CLAUSE for dhcp6

```
$KAME: COPYRIGHT,v 1.2 2004/07/29 19:02:18 jinmei Exp $

Copyright (C) 1998-2004 WIDE Project.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. Neither the name of the project nor the names of its contributors
   may be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE PROJECT OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
```

(continues on next page)

(continued from previous page)

```
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.
```

40.18 APACHE20 License Text for dhcpleases6

The following license text applies to dhcpleases6:

Listing 18: Download: APACHE20 for dhcpleases6

```
                                 Apache License
                           Version 2.0, January 2004
                        http://www.apache.org/licenses/

   TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

   1. Definitions.

      "License" shall mean the terms and conditions for use, reproduction,
      and distribution as defined by Sections 1 through 9 of this document.

      "Licensor" shall mean the copyright owner or entity authorized by
      the copyright owner that is granting the License.

      "Legal Entity" shall mean the union of the acting entity and all
      other entities that control, are controlled by, or are under common
      control with that entity. For the purposes of this definition,
      "control" means (i) the power, direct or indirect, to cause the
      direction or management of such entity, whether by contract or
      otherwise, or (ii) ownership of fifty percent (50%) or more of the
      outstanding shares, or (iii) beneficial ownership of such entity.

      "You" (or "Your") shall mean an individual or Legal Entity
      exercising permissions granted by this License.

      "Source" form shall mean the preferred form for making modifications,
      including but not limited to software source code, documentation
      source, and configuration files.

      "Object" form shall mean any form resulting from mechanical
      transformation or translation of a Source form, including but
      not limited to compiled object code, generated documentation,
      and conversions to other media types.

      "Work" shall mean the work of authorship, whether in Source or
      Object form, made available under the License, as indicated by a
      copyright notice that is included in or attached to the work
      (an example is provided in the Appendix below).

      "Derivative Works" shall mean any work, whether in Source or Object
      form, that is based on (or derived from) the Work and for which the
```

(continues on next page)

(continued from previous page)

```
   editorial revisions, annotations, elaborations, or other modifications
   represent, as a whole, an original work of authorship. For the purposes
   of this License, Derivative Works shall not include works that remain
   separable from, or merely link (or bind by name) to the interfaces of,
   the Work and Derivative Works thereof.

   "Contribution" shall mean any work of authorship, including
   the original version of the Work and any modifications or additions
   to that Work or Derivative Works thereof, that is intentionally
   submitted to Licensor for inclusion in the Work by the copyright owner
   or by an individual or Legal Entity authorized to submit on behalf of
   the copyright owner. For the purposes of this definition, "submitted"
   means any form of electronic, verbal, or written communication sent
   to the Licensor or its representatives, including but not limited to
   communication on electronic mailing lists, source code control systems,
   and issue tracking systems that are managed by, or on behalf of, the
   Licensor for the purpose of discussing and improving the Work, but
   excluding communication that is conspicuously marked or otherwise
   designated in writing by the copyright owner as "Not a Contribution."

   "Contributor" shall mean Licensor and any individual or Legal Entity
   on behalf of whom a Contribution has been received by Licensor and
   subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   copyright license to reproduce, prepare Derivative Works of,
   publicly display, publicly perform, sublicense, and distribute the
   Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   (except as stated in this section) patent license to make, have made,
   use, offer to sell, sell, import, and otherwise transfer the Work,
   where such license applies only to those patent claims licensable
   by such Contributor that are necessarily infringed by their
   Contribution(s) alone or by combination of their Contribution(s)
   with the Work to which such Contribution(s) was submitted. If You
   institute patent litigation against any entity (including a
   cross-claim or counterclaim in a lawsuit) alleging that the Work
   or a Contribution incorporated within the Work constitutes direct
   or contributory patent infringement, then any patent licenses
   granted to You under this License for that Work shall terminate
   as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the
   Work or Derivative Works thereof in any medium, with or without
   modifications, and in Source or Object form, provided that You
   meet the following conditions:
```

(continues on next page)

(continued from previous page)

```
    (a) You must give any other recipients of the Work or
        Derivative Works a copy of this License; and

    (b) You must cause any modified files to carry prominent notices
        stating that You changed the files; and

    (c) You must retain, in the Source form of any Derivative Works
        that You distribute, all copyright, patent, trademark, and
        attribution notices from the Source form of the Work,
        excluding those notices that do not pertain to any part of
        the Derivative Works; and

    (d) If the Work includes a "NOTICE" text file as part of its
        distribution, then any Derivative Works that You distribute must
        include a readable copy of the attribution notices contained
        within such NOTICE file, excluding those notices that do not
        pertain to any part of the Derivative Works, in at least one
        of the following places: within a NOTICE text file distributed
        as part of the Derivative Works; within the Source form or
        documentation, if provided along with the Derivative Works; or,
        within a display generated by the Derivative Works, if and
        wherever such third-party notices normally appear. The contents
        of the NOTICE file are for informational purposes only and
        do not modify the License. You may add Your own attribution
        notices within Derivative Works that You distribute, alongside
        or as an addendum to the NOTICE text from the Work, provided
        that such additional attribution notices cannot be construed
        as modifying the License.

    You may add Your own copyright statement to Your modifications and
    may provide additional or different license terms and conditions
    for use, reproduction, or distribution of Your modifications, or
    for any such Derivative Works as a whole, provided Your use,
    reproduction, and distribution of the Work otherwise complies with
    the conditions stated in this License.

  5. Submission of Contributions. Unless You explicitly state otherwise,
     any Contribution intentionally submitted for inclusion in the Work
     by You to the Licensor shall be under the terms and conditions of
     this License, without any additional terms or conditions.
     Notwithstanding the above, nothing herein shall supersede or modify
     the terms of any separate license agreement you may have executed
     with Licensor regarding such Contributions.

  6. Trademarks. This License does not grant permission to use the trade
     names, trademarks, service marks, or product names of the Licensor,
     except as required for reasonable and customary use in describing the
     origin of the Work and reproducing the content of the NOTICE file.

  7. Disclaimer of Warranty. Unless required by applicable law or
     agreed to in writing, Licensor provides the Work (and each
     Contributor provides its Contributions) on an "AS IS" BASIS,
```

(continues on next page)

(continued from previous page)

```
   WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
   implied, including, without limitation, any warranties or conditions
   of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A
   PARTICULAR PURPOSE. You are solely responsible for determining the
   appropriateness of using or redistributing the Work and assume any
   risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory,
   whether in tort (including negligence), contract, or otherwise,
   unless required by applicable law (such as deliberate and grossly
   negligent acts) or agreed to in writing, shall any Contributor be
   liable to You for damages, including any direct, indirect, special,
   incidental, or consequential damages of any character arising as a
   result of this License or out of the use or inability to use the
   Work (including but not limited to damages for loss of goodwill,
   work stoppage, computer failure or malfunction, or any and all
   other commercial damages or losses), even if such Contributor
   has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing
   the Work or Derivative Works thereof, You may choose to offer,
   and charge a fee for, acceptance of support, warranty, indemnity,
   or other liability obligations and/or rights consistent with this
   License. However, in accepting such obligations, You may act only
   on Your own behalf and on Your sole responsibility, not on behalf
   of any other Contributor, and only if You agree to indemnify,
   defend, and hold each Contributor harmless for any liability
   incurred by, or claims asserted against, such Contributor by reason
   of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS
```

40.19 APACHE20 License Text for dhcpleases

The following license text applies to dhcpleases:

Listing 19: Download: APACHE20 for dhcpleases

```
                              Apache License
                        Version 2.0, January 2004
                     http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

   "License" shall mean the terms and conditions for use, reproduction,
   and distribution as defined by Sections 1 through 9 of this document.

   "Licensor" shall mean the copyright owner or entity authorized by
```

(continues on next page)

(continued from previous page)

```
      the copyright owner that is granting the License.

      "Legal Entity" shall mean the union of the acting entity and all
      other entities that control, are controlled by, or are under common
      control with that entity. For the purposes of this definition,
      "control" means (i) the power, direct or indirect, to cause the
      direction or management of such entity, whether by contract or
      otherwise, or (ii) ownership of fifty percent (50%) or more of the
      outstanding shares, or (iii) beneficial ownership of such entity.

      "You" (or "Your") shall mean an individual or Legal Entity
      exercising permissions granted by this License.

      "Source" form shall mean the preferred form for making modifications,
      including but not limited to software source code, documentation
      source, and configuration files.

      "Object" form shall mean any form resulting from mechanical
      transformation or translation of a Source form, including but
      not limited to compiled object code, generated documentation,
      and conversions to other media types.

      "Work" shall mean the work of authorship, whether in Source or
      Object form, made available under the License, as indicated by a
      copyright notice that is included in or attached to the work
      (an example is provided in the Appendix below).

      "Derivative Works" shall mean any work, whether in Source or Object
      form, that is based on (or derived from) the Work and for which the
      editorial revisions, annotations, elaborations, or other modifications
      represent, as a whole, an original work of authorship. For the purposes
      of this License, Derivative Works shall not include works that remain
      separable from, or merely link (or bind by name) to the interfaces of,
      the Work and Derivative Works thereof.

      "Contribution" shall mean any work of authorship, including
      the original version of the Work and any modifications or additions
      to that Work or Derivative Works thereof, that is intentionally
      submitted to Licensor for inclusion in the Work by the copyright owner
      or by an individual or Legal Entity authorized to submit on behalf of
      the copyright owner. For the purposes of this definition, "submitted"
      means any form of electronic, verbal, or written communication sent
      to the Licensor or its representatives, including but not limited to
      communication on electronic mailing lists, source code control systems,
      and issue tracking systems that are managed by, or on behalf of, the
      Licensor for the purpose of discussing and improving the Work, but
      excluding communication that is conspicuously marked or otherwise
      designated in writing by the copyright owner as "Not a Contribution."

      "Contributor" shall mean Licensor and any individual or Legal Entity
      on behalf of whom a Contribution has been received by Licensor and
      subsequently incorporated within the Work.
```

(continues on next page)

(continued from previous page)

```
2. Grant of Copyright License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   copyright license to reproduce, prepare Derivative Works of,
   publicly display, publicly perform, sublicense, and distribute the
   Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   (except as stated in this section) patent license to make, have made,
   use, offer to sell, sell, import, and otherwise transfer the Work,
   where such license applies only to those patent claims licensable
   by such Contributor that are necessarily infringed by their
   Contribution(s) alone or by combination of their Contribution(s)
   with the Work to which such Contribution(s) was submitted. If You
   institute patent litigation against any entity (including a
   cross-claim or counterclaim in a lawsuit) alleging that the Work
   or a Contribution incorporated within the Work constitutes direct
   or contributory patent infringement, then any patent licenses
   granted to You under this License for that Work shall terminate
   as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the
   Work or Derivative Works thereof in any medium, with or without
   modifications, and in Source or Object form, provided that You
   meet the following conditions:

   (a) You must give any other recipients of the Work or
       Derivative Works a copy of this License; and

   (b) You must cause any modified files to carry prominent notices
       stating that You changed the files; and

   (c) You must retain, in the Source form of any Derivative Works
       that You distribute, all copyright, patent, trademark, and
       attribution notices from the Source form of the Work,
       excluding those notices that do not pertain to any part of
       the Derivative Works; and

   (d) If the Work includes a "NOTICE" text file as part of its
       distribution, then any Derivative Works that You distribute must
       include a readable copy of the attribution notices contained
       within such NOTICE file, excluding those notices that do not
       pertain to any part of the Derivative Works, in at least one
       of the following places: within a NOTICE text file distributed
       as part of the Derivative Works; within the Source form or
       documentation, if provided along with the Derivative Works; or,
       within a display generated by the Derivative Works, if and
       wherever such third-party notices normally appear. The contents
       of the NOTICE file are for informational purposes only and
```

(continues on next page)

(continued from previous page)

```
        do not modify the License. You may add Your own attribution
        notices within Derivative Works that You distribute, alongside
        or as an addendum to the NOTICE text from the Work, provided
        that such additional attribution notices cannot be construed
        as modifying the License.

    You may add Your own copyright statement to Your modifications and
    may provide additional or different license terms and conditions
    for use, reproduction, or distribution of Your modifications, or
    for any such Derivative Works as a whole, provided Your use,
    reproduction, and distribution of the Work otherwise complies with
    the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise,
    any Contribution intentionally submitted for inclusion in the Work
    by You to the Licensor shall be under the terms and conditions of
    this License, without any additional terms or conditions.
    Notwithstanding the above, nothing herein shall supersede or modify
    the terms of any separate license agreement you may have executed
    with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade
    names, trademarks, service marks, or product names of the Licensor,
    except as required for reasonable and customary use in describing the
    origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or
    agreed to in writing, Licensor provides the Work (and each
    Contributor provides its Contributions) on an "AS IS" BASIS,
    WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
    implied, including, without limitation, any warranties or conditions
    of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A
    PARTICULAR PURPOSE. You are solely responsible for determining the
    appropriateness of using or redistributing the Work and assume any
    risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory,
    whether in tort (including negligence), contract, or otherwise,
    unless required by applicable law (such as deliberate and grossly
    negligent acts) or agreed to in writing, shall any Contributor be
    liable to You for damages, including any direct, indirect, special,
    incidental, or consequential damages of any character arising as a
    result of this License or out of the use or inability to use the
    Work (including but not limited to damages for loss of goodwill,
    work stoppage, computer failure or malfunction, or any and all
    other commercial damages or losses), even if such Contributor
    has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing
    the Work or Derivative Works thereof, You may choose to offer,
    and charge a fee for, acceptance of support, warranty, indemnity,
    or other liability obligations and/or rights consistent with this
```

(continues on next page)

(continued from previous page)

```
     License. However, in accepting such obligations, You may act only
     on Your own behalf and on Your sole responsibility, not on behalf
     of any other Contributor, and only if You agree to indemnify,
     defend, and hold each Contributor harmless for any liability
     incurred by, or claims asserted against, such Contributor by reason
     of your accepting any such warranty or additional liability.

   END OF TERMS AND CONDITIONS
```

40.20 GPLv2 License Text for dmidecode

The following license text applies to dmidecode:

Listing 20: Download: GPLv2 for dmidecode

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
```

(continues on next page)

(continued from previous page)

```
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
```

(continues on next page)

(continued from previous page)

```
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
```

(continues on next page)

(continued from previous page)

```
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
```

(continues on next page)

(continued from previous page)

```
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
```

(continues on next page)

(continued from previous page)

```
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>
```

(continues on next page)

(continued from previous page)

```
    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.21 GPLv2 License Text for dnsmasq

The following license text applies to dnsmasq:

Listing 21: Download: GPLv2 for dnsmasq

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
```

(continues on next page)

(continued from previous page)

```
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
```

(continues on next page)

(continued from previous page)

```
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
```

(continues on next page)

(continued from previous page)

```
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
```

(continues on next page)

(continued from previous page)

```
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
```

(continues on next page)

(continued from previous page)

```
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
```

(continues on next page)

(continued from previous page)

```
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.22 MIT License Text for expat

The following license text applies to expat:

Listing 22: Download: MIT for expat

```
Copyright (c) 1998-2000 Thai Open Source Software Center Ltd and Clark Cooper
Copyright (c) 2001-2019 Expat maintainers

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
"Software"), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject to
the following conditions:
```

(continues on next page)

(continued from previous page)

```
The above copyright notice and this permission notice shall be included
in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.
IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY
CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT,
TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE
SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

40.23 GPLv3+ License Text for gettext-runtime

The following license text applies to gettext-runtime:

Listing 23: Download: GPLv3+ for gettext-runtime

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 3, 29 June 2007

 Copyright (C) 2007 Free Software Foundation, Inc. <https://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The GNU General Public License is a free, copyleft license for
software and other kinds of works.

  The licenses for most software and other practical works are designed
to take away your freedom to share and change the works.  By contrast,
the GNU General Public License is intended to guarantee your freedom to
share and change all versions of a program--to make sure it remains free
software for all its users.  We, the Free Software Foundation, use the
GNU General Public License for most of our software; it applies also to
any other work released this way by its authors.  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
them if you wish), that you receive source code or can get it if you
want it, that you can change the software or use pieces of it in new
free programs, and that you know you can do these things.

  To protect your rights, we need to prevent others from denying you
these rights or asking you to surrender the rights.  Therefore, you have
certain responsibilities if you distribute copies of the software, or if
you modify it: responsibilities to respect the freedom of others.
```

(continues on next page)

(continued from previous page)

```
  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must pass on to the recipients the same
freedoms that you received.  You must make sure that they, too, receive
or can get the source code.  And you must show them these terms so they
know their rights.

  Developers that use the GNU GPL protect your rights with two steps:
(1) assert copyright on the software, and (2) offer you this License
giving you legal permission to copy, distribute and/or modify it.

  For the developers' and authors' protection, the GPL clearly explains
that there is no warranty for this free software.  For both users' and
authors' sake, the GPL requires that modified versions be marked as
changed, so that their problems will not be attributed erroneously to
authors of previous versions.

  Some devices are designed to deny users access to install or run
modified versions of the software inside them, although the manufacturer
can do so.  This is fundamentally incompatible with the aim of
protecting users' freedom to change the software.  The systematic
pattern of such abuse occurs in the area of products for individuals to
use, which is precisely where it is most unacceptable.  Therefore, we
have designed this version of the GPL to prohibit the practice for those
products.  If such problems arise substantially in other domains, we
stand ready to extend this provision to those domains in future versions
of the GPL, as needed to protect the freedom of users.

  Finally, every program is threatened constantly by software patents.
States should not allow patents to restrict development and use of
software on general-purpose computers, but in those that do, we wish to
avoid the special danger that patents applied to a free program could
make it effectively proprietary.  To prevent this, the GPL assures that
patents cannot be used to render the program non-free.

  The precise terms and conditions for copying, distribution and
modification follow.

                       TERMS AND CONDITIONS

  0. Definitions.

  "This License" refers to version 3 of the GNU General Public License.

  "Copyright" also means copyright-like laws that apply to other kinds of
works, such as semiconductor masks.

  "The Program" refers to any copyrightable work licensed under this
License.  Each licensee is addressed as "you".  "Licensees" and
"recipients" may be individuals or organizations.

  To "modify" a work means to copy from or adapt all or part of the work
in a fashion requiring copyright permission, other than the making of an
```

(continues on next page)

(continued from previous page)

```
exact copy.  The resulting work is called a "modified version" of the
earlier work or a work "based on" the earlier work.

  A "covered work" means either the unmodified Program or a work based
on the Program.

  To "propagate" a work means to do anything with it that, without
permission, would make you directly or secondarily liable for
infringement under applicable copyright law, except executing it on a
computer or modifying a private copy.  Propagation includes copying,
distribution (with or without modification), making available to the
public, and in some countries other activities as well.

  To "convey" a work means any kind of propagation that enables other
parties to make or receive copies.  Mere interaction with a user through
a computer network, with no transfer of a copy, is not conveying.

  An interactive user interface displays "Appropriate Legal Notices"
to the extent that it includes a convenient and prominently visible
feature that (1) displays an appropriate copyright notice, and (2)
tells the user that there is no warranty for the work (except to the
extent that warranties are provided), that licensees may convey the
work under this License, and how to view a copy of this License.  If
the interface presents a list of user commands or options, such as a
menu, a prominent item in the list meets this criterion.

  1. Source Code.

  The "source code" for a work means the preferred form of the work
for making modifications to it.  "Object code" means any non-source
form of a work.

  A "Standard Interface" means an interface that either is an official
standard defined by a recognized standards body, or, in the case of
interfaces specified for a particular programming language, one that
is widely used among developers working in that language.

  The "System Libraries" of an executable work include anything, other
than the work as a whole, that (a) is included in the normal form of
packaging a Major Component, but which is not part of that Major
Component, and (b) serves only to enable use of the work with that
Major Component, or to implement a Standard Interface for which an
implementation is available to the public in source code form.  A
"Major Component", in this context, means a major essential component
(kernel, window system, and so on) of the specific operating system
(if any) on which the executable work runs, or a compiler used to
produce the work, or an object code interpreter used to run it.

  The "Corresponding Source" for a work in object code form means all
the source code needed to generate, install, and (for an executable
work) run the object code and to modify the work, including scripts to
control those activities.  However, it does not include the work's
```

(continues on next page)

(continued from previous page)

```
System Libraries, or general-purpose tools or generally available free
programs which are used unmodified in performing those activities but
which are not part of the work.  For example, Corresponding Source
includes interface definition files associated with source files for
the work, and the source code for shared libraries and dynamically
linked subprograms that the work is specifically designed to require,
such as by intimate data communication or control flow between those
subprograms and other parts of the work.

  The Corresponding Source need not include anything that users
can regenerate automatically from other parts of the Corresponding
Source.

  The Corresponding Source for a work in source code form is that
same work.

  2. Basic Permissions.

  All rights granted under this License are granted for the term of
copyright on the Program, and are irrevocable provided the stated
conditions are met.  This License explicitly affirms your unlimited
permission to run the unmodified Program.  The output from running a
covered work is covered by this License only if the output, given its
content, constitutes a covered work.  This License acknowledges your
rights of fair use or other equivalent, as provided by copyright law.

  You may make, run and propagate covered works that you do not
convey, without conditions so long as your license otherwise remains
in force.  You may convey covered works to others for the sole purpose
of having them make modifications exclusively for you, or provide you
with facilities for running those works, provided that you comply with
the terms of this License in conveying all material for which you do
not control copyright.  Those thus making or running the covered works
for you must do so exclusively on your behalf, under your direction
and control, on terms that prohibit them from making any copies of
your copyrighted material outside their relationship with you.

  Conveying under any other circumstances is permitted solely under
the conditions stated below.  Sublicensing is not allowed; section 10
makes it unnecessary.

  3. Protecting Users' Legal Rights From Anti-Circumvention Law.

  No covered work shall be deemed part of an effective technological
measure under any applicable law fulfilling obligations under article
11 of the WIPO copyright treaty adopted on 20 December 1996, or
similar laws prohibiting or restricting circumvention of such
measures.

  When you convey a covered work, you waive any legal power to forbid
circumvention of technological measures to the extent such circumvention
is effected by exercising rights under this License with respect to
```

(continues on next page)

(continued from previous page)

```
the covered work, and you disclaim any intention to limit operation or
modification of the work as a means of enforcing, against the work's
users, your or third parties' legal rights to forbid circumvention of
technological measures.

  4. Conveying Verbatim Copies.

  You may convey verbatim copies of the Program's source code as you
receive it, in any medium, provided that you conspicuously and
appropriately publish on each copy an appropriate copyright notice;
keep intact all notices stating that this License and any
non-permissive terms added in accord with section 7 apply to the code;
keep intact all notices of the absence of any warranty; and give all
recipients a copy of this License along with the Program.

  You may charge any price or no price for each copy that you convey,
and you may offer support or warranty protection for a fee.

  5. Conveying Modified Source Versions.

  You may convey a work based on the Program, or the modifications to
produce it from the Program, in the form of source code under the
terms of section 4, provided that you also meet all of these conditions:

    a) The work must carry prominent notices stating that you modified
    it, and giving a relevant date.

    b) The work must carry prominent notices stating that it is
    released under this License and any conditions added under section
    7.  This requirement modifies the requirement in section 4 to
    "keep intact all notices".

    c) You must license the entire work, as a whole, under this
    License to anyone who comes into possession of a copy.  This
    License will therefore apply, along with any applicable section 7
    additional terms, to the whole of the work, and all its parts,
    regardless of how they are packaged.  This License gives no
    permission to license the work in any other way, but it does not
    invalidate such permission if you have separately received it.

    d) If the work has interactive user interfaces, each must display
    Appropriate Legal Notices; however, if the Program has interactive
    interfaces that do not display Appropriate Legal Notices, your
    work need not make them do so.

  A compilation of a covered work with other separate and independent
works, which are not by their nature extensions of the covered work,
and which are not combined with it such as to form a larger program,
in or on a volume of a storage or distribution medium, is called an
"aggregate" if the compilation and its resulting copyright are not
used to limit the access or legal rights of the compilation's users
beyond what the individual works permit.  Inclusion of a covered work
```

(continues on next page)

(continued from previous page)

```
in an aggregate does not cause this License to apply to the other
parts of the aggregate.

  6. Conveying Non-Source Forms.

  You may convey a covered work in object code form under the terms
of sections 4 and 5, provided that you also convey the
machine-readable Corresponding Source under the terms of this License,
in one of these ways:

    a) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by the
    Corresponding Source fixed on a durable physical medium
    customarily used for software interchange.

    b) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by a
    written offer, valid for at least three years and valid for as
    long as you offer spare parts or customer support for that product
    model, to give anyone who possesses the object code either (1) a
    copy of the Corresponding Source for all the software in the
    product that is covered by this License, on a durable physical
    medium customarily used for software interchange, for a price no
    more than your reasonable cost of physically performing this
    conveying of source, or (2) access to copy the
    Corresponding Source from a network server at no charge.

    c) Convey individual copies of the object code with a copy of the
    written offer to provide the Corresponding Source.  This
    alternative is allowed only occasionally and noncommercially, and
    only if you received the object code with such an offer, in accord
    with subsection 6b.

    d) Convey the object code by offering access from a designated
    place (gratis or for a charge), and offer equivalent access to the
    Corresponding Source in the same way through the same place at no
    further charge.  You need not require recipients to copy the
    Corresponding Source along with the object code.  If the place to
    copy the object code is a network server, the Corresponding Source
    may be on a different server (operated by you or a third party)
    that supports equivalent copying facilities, provided you maintain
    clear directions next to the object code saying where to find the
    Corresponding Source.  Regardless of what server hosts the
    Corresponding Source, you remain obligated to ensure that it is
    available for as long as needed to satisfy these requirements.

    e) Convey the object code using peer-to-peer transmission, provided
    you inform other peers where the object code and Corresponding
    Source of the work are being offered to the general public at no
    charge under subsection 6d.

  A separable portion of the object code, whose source code is excluded
```

(continues on next page)

(continued from previous page)

```
from the Corresponding Source as a System Library, need not be
included in conveying the object code work.

  A "User Product" is either (1) a "consumer product", which means any
tangible personal property which is normally used for personal, family,
or household purposes, or (2) anything designed or sold for incorporation
into a dwelling.  In determining whether a product is a consumer product,
doubtful cases shall be resolved in favor of coverage.  For a particular
product received by a particular user, "normally used" refers to a
typical or common use of that class of product, regardless of the status
of the particular user or of the way in which the particular user
actually uses, or expects or is expected to use, the product.  A product
is a consumer product regardless of whether the product has substantial
commercial, industrial or non-consumer uses, unless such uses represent
the only significant mode of use of the product.

  "Installation Information" for a User Product means any methods,
procedures, authorization keys, or other information required to install
and execute modified versions of a covered work in that User Product from
a modified version of its Corresponding Source.  The information must
suffice to ensure that the continued functioning of the modified object
code is in no case prevented or interfered with solely because
modification has been made.

  If you convey an object code work under this section in, or with, or
specifically for use in, a User Product, and the conveying occurs as
part of a transaction in which the right of possession and use of the
User Product is transferred to the recipient in perpetuity or for a
fixed term (regardless of how the transaction is characterized), the
Corresponding Source conveyed under this section must be accompanied
by the Installation Information.  But this requirement does not apply
if neither you nor any third party retains the ability to install
modified object code on the User Product (for example, the work has
been installed in ROM).

  The requirement to provide Installation Information does not include a
requirement to continue to provide support service, warranty, or updates
for a work that has been modified or installed by the recipient, or for
the User Product in which it has been modified or installed.  Access to a
network may be denied when the modification itself materially and
adversely affects the operation of the network or violates the rules and
protocols for communication across the network.

  Corresponding Source conveyed, and Installation Information provided,
in accord with this section must be in a format that is publicly
documented (and with an implementation available to the public in
source code form), and must require no special password or key for
unpacking, reading or copying.

  7. Additional Terms.

  "Additional permissions" are terms that supplement the terms of this
```

(continues on next page)

(continued from previous page)

```
License by making exceptions from one or more of its conditions.
Additional permissions that are applicable to the entire Program shall
be treated as though they were included in this License, to the extent
that they are valid under applicable law.  If additional permissions
apply only to part of the Program, that part may be used separately
under those permissions, but the entire Program remains governed by
this License without regard to the additional permissions.

  When you convey a copy of a covered work, you may at your option
remove any additional permissions from that copy, or from any part of
it.  (Additional permissions may be written to require their own
removal in certain cases when you modify the work.)  You may place
additional permissions on material, added by you to a covered work,
for which you have or can give appropriate copyright permission.

  Notwithstanding any other provision of this License, for material you
add to a covered work, you may (if authorized by the copyright holders of
that material) supplement the terms of this License with terms:

    a) Disclaiming warranty or limiting liability differently from the
    terms of sections 15 and 16 of this License; or

    b) Requiring preservation of specified reasonable legal notices or
    author attributions in that material or in the Appropriate Legal
    Notices displayed by works containing it; or

    c) Prohibiting misrepresentation of the origin of that material, or
    requiring that modified versions of such material be marked in
    reasonable ways as different from the original version; or

    d) Limiting the use for publicity purposes of names of licensors or
    authors of the material; or

    e) Declining to grant rights under trademark law for use of some
    trade names, trademarks, or service marks; or

    f) Requiring indemnification of licensors and authors of that
    material by anyone who conveys the material (or modified versions of
    it) with contractual assumptions of liability to the recipient, for
    any liability that these contractual assumptions directly impose on
    those licensors and authors.

  All other non-permissive additional terms are considered "further
restrictions" within the meaning of section 10.  If the Program as you
received it, or any part of it, contains a notice stating that it is
governed by this License along with a term that is a further
restriction, you may remove that term.  If a license document contains
a further restriction but permits relicensing or conveying under this
License, you may add to a covered work material governed by the terms
of that license document, provided that the further restriction does
not survive such relicensing or conveying.
```

(continues on next page)

(continued from previous page)

```
  If you add terms to a covered work in accord with this section, you
must place, in the relevant source files, a statement of the
additional terms that apply to those files, or a notice indicating
where to find the applicable terms.

  Additional terms, permissive or non-permissive, may be stated in the
form of a separately written license, or stated as exceptions;
the above requirements apply either way.

  8. Termination.

  You may not propagate or modify a covered work except as expressly
provided under this License.  Any attempt otherwise to propagate or
modify it is void, and will automatically terminate your rights under
this License (including any patent licenses granted under the third
paragraph of section 11).

  However, if you cease all violation of this License, then your
license from a particular copyright holder is reinstated (a)
provisionally, unless and until the copyright holder explicitly and
finally terminates your license, and (b) permanently, if the copyright
holder fails to notify you of the violation by some reasonable means
prior to 60 days after the cessation.

  Moreover, your license from a particular copyright holder is
reinstated permanently if the copyright holder notifies you of the
violation by some reasonable means, this is the first time you have
received notice of violation of this License (for any work) from that
copyright holder, and you cure the violation prior to 30 days after
your receipt of the notice.

  Termination of your rights under this section does not terminate the
licenses of parties who have received copies or rights from you under
this License.  If your rights have been terminated and not permanently
reinstated, you do not qualify to receive new licenses for the same
material under section 10.

  9. Acceptance Not Required for Having Copies.

  You are not required to accept this License in order to receive or
run a copy of the Program.  Ancillary propagation of a covered work
occurring solely as a consequence of using peer-to-peer transmission
to receive a copy likewise does not require acceptance.  However,
nothing other than this License grants you permission to propagate or
modify any covered work.  These actions infringe copyright if you do
not accept this License.  Therefore, by modifying or propagating a
covered work, you indicate your acceptance of this License to do so.

  10. Automatic Licensing of Downstream Recipients.

  Each time you convey a covered work, the recipient automatically
receives a license from the original licensors, to run, modify and
```

(continues on next page)

(continued from previous page)

```
propagate that work, subject to this License.  You are not responsible
for enforcing compliance by third parties with this License.

  An "entity transaction" is a transaction transferring control of an
organization, or substantially all assets of one, or subdividing an
organization, or merging organizations.  If propagation of a covered
work results from an entity transaction, each party to that
transaction who receives a copy of the work also receives whatever
licenses to the work the party's predecessor in interest had or could
give under the previous paragraph, plus a right to possession of the
Corresponding Source of the work from the predecessor in interest, if
the predecessor has it or can get it with reasonable efforts.

  You may not impose any further restrictions on the exercise of the
rights granted or affirmed under this License.  For example, you may
not impose a license fee, royalty, or other charge for exercise of
rights granted under this License, and you may not initiate litigation
(including a cross-claim or counterclaim in a lawsuit) alleging that
any patent claim is infringed by making, using, selling, offering for
sale, or importing the Program or any portion of it.

  11. Patents.

  A "contributor" is a copyright holder who authorizes use under this
License of the Program or a work on which the Program is based.  The
work thus licensed is called the contributor's "contributor version".

  A contributor's "essential patent claims" are all patent claims
owned or controlled by the contributor, whether already acquired or
hereafter acquired, that would be infringed by some manner, permitted
by this License, of making, using, or selling its contributor version,
but do not include claims that would be infringed only as a
consequence of further modification of the contributor version.  For
purposes of this definition, "control" includes the right to grant
patent sublicenses in a manner consistent with the requirements of
this License.

  Each contributor grants you a non-exclusive, worldwide, royalty-free
patent license under the contributor's essential patent claims, to
make, use, sell, offer for sale, import and otherwise run, modify and
propagate the contents of its contributor version.

  In the following three paragraphs, a "patent license" is any express
agreement or commitment, however denominated, not to enforce a patent
(such as an express permission to practice a patent or covenant not to
sue for patent infringement).  To "grant" such a patent license to a
party means to make such an agreement or commitment not to enforce a
patent against the party.

  If you convey a covered work, knowingly relying on a patent license,
and the Corresponding Source of the work is not available for anyone
to copy, free of charge and under the terms of this License, through a
```

(continues on next page)

(continued from previous page)

```
publicly available network server or other readily accessible means,
then you must either (1) cause the Corresponding Source to be so
available, or (2) arrange to deprive yourself of the benefit of the
patent license for this particular work, or (3) arrange, in a manner
consistent with the requirements of this License, to extend the patent
license to downstream recipients.  "Knowingly relying" means you have
actual knowledge that, but for the patent license, your conveying the
covered work in a country, or your recipient's use of the covered work
in a country, would infringe one or more identifiable patents in that
country that you have reason to believe are valid.

  If, pursuant to or in connection with a single transaction or
arrangement, you convey, or propagate by procuring conveyance of, a
covered work, and grant a patent license to some of the parties
receiving the covered work authorizing them to use, propagate, modify
or convey a specific copy of the covered work, then the patent license
you grant is automatically extended to all recipients of the covered
work and works based on it.

  A patent license is "discriminatory" if it does not include within
the scope of its coverage, prohibits the exercise of, or is
conditioned on the non-exercise of one or more of the rights that are
specifically granted under this License.  You may not convey a covered
work if you are a party to an arrangement with a third party that is
in the business of distributing software, under which you make payment
to the third party based on the extent of your activity of conveying
the work, and under which the third party grants, to any of the
parties who would receive the covered work from you, a discriminatory
patent license (a) in connection with copies of the covered work
conveyed by you (or copies made from those copies), or (b) primarily
for and in connection with specific products or compilations that
contain the covered work, unless you entered into that arrangement,
or that patent license was granted, prior to 28 March 2007.

  Nothing in this License shall be construed as excluding or limiting
any implied license or other defenses to infringement that may
otherwise be available to you under applicable patent law.

  12. No Surrender of Others' Freedom.

  If conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot convey a
covered work so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you may
not convey it at all.  For example, if you agree to terms that obligate you
to collect a royalty for further conveying from those to whom you convey
the Program, the only way you could satisfy both those terms and this
License would be to refrain entirely from conveying the Program.

  13. Use with the GNU Affero General Public License.
```

(continues on next page)

(continued from previous page)

```
  Notwithstanding any other provision of this License, you have
permission to link or combine any covered work with a work licensed
under version 3 of the GNU Affero General Public License into a single
combined work, and to convey the resulting work.  The terms of this
License will continue to apply to the part which is the covered work,
but the special requirements of the GNU Affero General Public License,
section 13, concerning interaction through a network will apply to the
combination as such.

  14. Revised Versions of this License.

  The Free Software Foundation may publish revised and/or new versions of
the GNU General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

  Each version is given a distinguishing version number.  If the
Program specifies that a certain numbered version of the GNU General
Public License "or any later version" applies to it, you have the
option of following the terms and conditions either of that numbered
version or of any later version published by the Free Software
Foundation.  If the Program does not specify a version number of the
GNU General Public License, you may choose any version ever published
by the Free Software Foundation.

  If the Program specifies that a proxy can decide which future
versions of the GNU General Public License can be used, that proxy's
public statement of acceptance of a version permanently authorizes you
to choose that version for the Program.

  Later license versions may give you additional or different
permissions.  However, no additional obligations are imposed on any
author or copyright holder as a result of your choosing to follow a
later version.

  15. Disclaimer of Warranty.

  THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY
APPLICABLE LAW.  EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT
HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY
OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM
IS WITH YOU.  SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF
ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. Limitation of Liability.

  IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS
THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY
GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE
```

(continues on next page)

(continued from previous page)

```
USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF
DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD
PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS),
EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF
SUCH DAMAGES.

  17. Interpretation of Sections 15 and 16.

  If the disclaimer of warranty and limitation of liability provided
above cannot be given local legal effect according to their terms,
reviewing courts shall apply local law that most closely approximates
an absolute waiver of all civil liability in connection with the
Program, unless a warranty or assumption of liability accompanies a
copy of the Program in return for a fee.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
state the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software: you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation, either version 3 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program.  If not, see <https://www.gnu.org/licenses/>.

Also add information on how to contact you by electronic and paper mail.

  If the program does terminal interaction, make it output a short
notice like this when it starts in an interactive mode:

    <program>  Copyright (C) <year>  <name of author>
    This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
```

(continues on next page)

(continued from previous page)

```
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, your program's commands
might be different; for a GUI interface, you would use an "about box".

  You should also get your employer (if you work as a programmer) or school,
if any, to sign a "copyright disclaimer" for the program, if necessary.
For more information on this, and how to apply and follow the GNU GPL, see
<https://www.gnu.org/licenses/>.

  The GNU General Public License does not permit incorporating your program
into proprietary programs.  If your program is a subroutine library, you
may consider it more useful to permit linking proprietary applications with
the library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.  But first, please read
<https://www.gnu.org/philosophy/why-not-lgpl.html>.
```

40.24 LGPL21+ License Text for gettext-runtime

The following license text applies to gettext-runtime:

Listing 24: Download: LGPL21+ for gettext-runtime

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
     51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations
below.
```

(continues on next page)

(continued from previous page)

```
  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be
introduced by others.
^L
  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
```

(continues on next page)

(continued from previous page)

```
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it
becomes
a de-facto standard.  To achieve this, non-free programs must be
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must
be combined with the library in order to run.
^L
                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
```

(continues on next page)

(continued from previous page)

```
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control
compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
```

(continues on next page)

(continued from previous page)

```
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in
these notices.
^L
  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
```

(continues on next page)

(continued from previous page)

```
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,
whether or not they are linked directly with the Library itself.
^L
  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
```

(continues on next page)

(continued from previous page)

```
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you
distribute.
^L
  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
```

(continues on next page)

(continued from previous page)

```
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with
this License.
^L
  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.
```

(continues on next page)

(continued from previous page)

```
If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply, and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License
may add an explicit geographical distribution limitation excluding those
countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by
the Free Software Foundation.
^L
  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
```

(continues on next page)

(continued from previous page)

```
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.

                     END OF TERMS AND CONDITIONS
^L
           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms
of the ordinary General Public License).

  To apply these terms, attach the following notices to the library.
It is safest to attach them to the start of each source file to most
effectively convey the exclusion of warranty; and each file should
have at least the "copyright" line and a pointer to where the full
notice is found.


    <one line to give the library's name and a brief idea of what it
does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
```

(continues on next page)

(continued from previous page)

```
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston,
    MA 02110-1301, USA

Also add information on how to contact you by electronic and paper
mail.

You should also get your employer (if you work as a programmer) or
your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James
Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That's all there is to it!
```

40.25 LGPL20 License Text for glib

The following license text applies to glib:

Listing 25: Download: LGPL20 for glib

```
                  GNU LIBRARY GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1991 Free Software Foundation, Inc.
 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the library GPL.  It is
 numbered 2 because it goes with version 2 of the ordinary GPL.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Library General Public License, applies to some
specially designated Free Software Foundation software, and to any
other libraries whose authors decide to use it.  You can use it for
your libraries, too.
```

(continues on next page)

(continued from previous page)

```
  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if
you distribute copies of the library, or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link a program with the library, you must provide
complete object files to the recipients so that they can relink them
with the library, after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  Our method of protecting your rights has two steps: (1) copyright
the library, and (2) offer you this license which gives you legal
permission to copy, distribute and/or modify the library.

  Also, for each distributor's protection, we want to make certain
that everyone understands that there is no warranty for this free
library.  If the library is modified by someone else and passed on, we
want its recipients to know that what they have is not the original
version, so that any problems introduced by others will not reflect on

the original authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that companies distributing free
software will individually obtain patent licenses, thus in effect
transforming the program into proprietary software.  To prevent this,
we have made it clear that any patent must be licensed for everyone's
free use or not licensed at all.

  Most GNU software, including some libraries, is covered by the ordinary
GNU General Public License, which was designed for utility programs.  This
license, the GNU Library General Public License, applies to certain
designated libraries.  This license is quite different from the ordinary
one; be sure to read it in full, and don't assume that anything in it is
the same as in the ordinary license.

  The reason we have a separate public license for some libraries is that
they blur the distinction we usually make between modifying or adding to a
program and simply using it.  Linking a program with a library, without
changing the library, is in some sense simply using the library, and is
analogous to running a utility program or application program.  However, in
a textual and legal sense, the linked executable is a combined work, a
```

(continues on next page)

(continued from previous page)

```
derivative of the original library, and the ordinary General Public License
treats it as such.

  Because of this blurred distinction, using the ordinary General
Public License for libraries did not effectively promote software
sharing, because most developers did not use the libraries.  We
concluded that weaker conditions might promote sharing better.

  However, unrestricted linking of non-free programs would deprive the
users of those programs of all benefit from the free status of the
libraries themselves.  This Library General Public License is intended to
permit developers of non-free programs to use free libraries, while
preserving your freedom as a user of such programs to change the free
libraries that are incorporated in them.  (We have not seen how to achieve
this as regards changes in header files, but we have achieved it as regards
changes in the actual functions of the Library.)  The hope is that this
will lead to faster development of free libraries.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, while the latter only
works together with the library.

  Note that it is possible for a library to be covered by the ordinary

General Public License rather than by this special one.

                  GNU LIBRARY GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library which
contains a notice placed by the copyright holder or other authorized
party saying it may be distributed under the terms of this Library
General Public License (also called "this License").  Each licensee is
addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
```

(continues on next page)

(continued from previous page)

```
interface definition files, plus the scripts used to control compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)
```

(continues on next page)

(continued from previous page)

```
These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
```

(continues on next page)

(continued from previous page)

```
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also compile or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
```

(continues on next page)

(continued from previous page)

```
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    c) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    d) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the source code distributed need not include anything that is normally
distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you

distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
```

(continues on next page)

(continued from previous page)

```
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.
```

(continues on next page)

(continued from previous page)

```
This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded.  In such case, this License incorporates the limitation as if
written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Library General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
```

(continues on next page)

(continued from previous page)

```
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.


                     END OF TERMS AND CONDITIONS

           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms of the
ordinary General Public License).

  To apply these terms, attach the following notices to the library.  It is
safest to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Library General Public
    License as published by the Free Software Foundation; either
    version 2 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Library General Public License for more details.

    You should have received a copy of the GNU Library General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That's all there is to it!
```

40.26 ICU License Text for icu

The following license text applies to icu:

Listing 26: Download: ICU for icu

```
COPYRIGHT AND PERMISSION NOTICE (ICU 58 and later)

Copyright © 1991-2020 Unicode, Inc. All rights reserved.
Distributed under the Terms of Use in https://www.unicode.org/copyright.html.

Permission is hereby granted, free of charge, to any person obtaining
a copy of the Unicode data files and any associated documentation
(the "Data Files") or Unicode software and any associated documentation
(the "Software") to deal in the Data Files or Software
without restriction, including without limitation the rights to use,
copy, modify, merge, publish, distribute, and/or sell copies of
the Data Files or Software, and to permit persons to whom the Data Files
or Software are furnished to do so, provided that either
(a) this copyright and permission notice appear with all copies
of the Data Files or Software, or
(b) this copyright and permission notice appear in associated
Documentation.

THE DATA FILES AND SOFTWARE ARE PROVIDED "AS IS", WITHOUT WARRANTY OF
ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE
WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND
NONINFRINGEMENT OF THIRD PARTY RIGHTS.
IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS
NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL
DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE,
DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER
TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR
PERFORMANCE OF THE DATA FILES OR SOFTWARE.

Except as contained in this notice, the name of a copyright holder
shall not be used in advertising or otherwise to promote the sale,
use or other dealings in these Data Files or Software without prior
written authorization of the copyright holder.

---------------------

Third-Party Software Licenses

This section contains third-party software notices and/or additional
terms for licensed third-party software components included within ICU
libraries.

1. ICU License - ICU 1.8.1 to ICU 57.1

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1995-2016 International Business Machines Corporation and others
```

(continues on next page)

(continued from previous page)

```
All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
"Software"), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, and/or sell copies of the Software, and to permit persons
to whom the Software is furnished to do so, provided that the above
copyright notice(s) and this permission notice appear in all copies of
the Software and that both the above copyright notice(s) and this
permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT
OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR
HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY
SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER
RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF
CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN
CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder
shall not be used in advertising or otherwise to promote the sale, use
or other dealings in this Software without prior written authorization
of the copyright holder.

All trademarks and registered trademarks mentioned herein are the
property of their respective owners.

2. Chinese/Japanese Word Break Dictionary Data (cjdict.txt)

 #     The Google Chrome software developed by Google is licensed under
 # the BSD license. Other software included in this distribution is
 # provided under other licenses, as set forth below.
 #
 #  The BSD License
 #  http://opensource.org/licenses/bsd-license.php
 #  Copyright (C) 2006-2008, Google Inc.
 #
 #  All rights reserved.
 #
 #  Redistribution and use in source and binary forms, with or without
 # modification, are permitted provided that the following conditions are met:
 #
 #  Redistributions of source code must retain the above copyright notice,
 # this list of conditions and the following disclaimer.
 #  Redistributions in binary form must reproduce the above
 # copyright notice, this list of conditions and the following
 # disclaimer in the documentation and/or other materials provided with
 # the distribution.
 #  Neither the name of  Google Inc. nor the names of its
```

(continues on next page)

(continued from previous page)

```
# contributors may be used to endorse or promote products derived from
# this software without specific prior written permission.
#
#
#  THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND
# CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES,
# INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
# MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
# DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
# LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
# CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
# SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR
# BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
# LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
# NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
# SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
#
#
#  The word list in cjdict.txt are generated by combining three word lists
# listed below with further processing for compound word breaking. The
# frequency is generated with an iterative training against Google web
# corpora.
#
#  * Libtabe (Chinese)
#    - https://sourceforge.net/project/?group_id=1519
#    - Its license terms and conditions are shown below.
#
#  * IPADIC (Japanese)
#    - http://chasen.aist-nara.ac.jp/chasen/distribution.html
#    - Its license terms and conditions are shown below.
#
#  ---------COPYING.libtabe ---- BEGIN--------------------
#
#  /*
#   * Copyright (c) 1999 TaBE Project.
#   * Copyright (c) 1999 Pai-Hsiang Hsiao.
#   * All rights reserved.
#   *
#   * Redistribution and use in source and binary forms, with or without
#   * modification, are permitted provided that the following conditions
#   * are met:
#   *
#   * . Redistributions of source code must retain the above copyright
#   *   notice, this list of conditions and the following disclaimer.
#   * . Redistributions in binary form must reproduce the above copyright
#   *   notice, this list of conditions and the following disclaimer in
#   *   the documentation and/or other materials provided with the
#   *   distribution.
#   * . Neither the name of the TaBE Project nor the names of its
#   *   contributors may be used to endorse or promote products derived
#   *   from this software without specific prior written permission.
#   *
```

(continues on next page)

(continued from previous page)

```
#   * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
#   * "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
#   * LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
#   * FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE
#   * REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
#   * INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
#   * (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
#   * SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
#   * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
#   * STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
#   * ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
#   * OF THE POSSIBILITY OF SUCH DAMAGE.
#   */
#
#  /*
#   * Copyright (c) 1999 Computer Systems and Communication Lab,
#   *                    Institute of Information Science, Academia
#       *                    Sinica. All rights reserved.
#   *
#   * Redistribution and use in source and binary forms, with or without
#   * modification, are permitted provided that the following conditions
#   * are met:
#   *
#   * . Redistributions of source code must retain the above copyright
#   *   notice, this list of conditions and the following disclaimer.
#   * . Redistributions in binary form must reproduce the above copyright
#   *   notice, this list of conditions and the following disclaimer in
#   *   the documentation and/or other materials provided with the
#   *   distribution.
#   * . Neither the name of the Computer Systems and Communication Lab
#   *   nor the names of its contributors may be used to endorse or
#   *   promote products derived from this software without specific
#   *   prior written permission.
#   *
#   * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
#   * "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
#   * LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
#   * FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE
#   * REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
#   * INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
#   * (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
#   * SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
#   * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
#   * STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
#   * ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
#   * OF THE POSSIBILITY OF SUCH DAMAGE.
#   */
#
#  Copyright 1996 Chih-Hao Tsai @ Beckman Institute,
#      University of Illinois
#  c-tsai4@uiuc.edu  http://casper.beckman.uiuc.edu/~c-tsai4
#
```

(continues on next page)

(continued from previous page)

```
#  ---------------COPYING.libtabe-----END--------------------------------
#
#
#  ---------------COPYING.ipadic-----BEGIN-------------------------------
#
#  Copyright 2000, 2001, 2002, 2003 Nara Institute of Science
#  and Technology.  All Rights Reserved.
#
#  Use, reproduction, and distribution of this software is permitted.
#  Any copy of this software, whether in its original form or modified,
#  must include both the above copyright notice and the following
#  paragraphs.
#
#  Nara Institute of Science and Technology (NAIST),
#  the copyright holders, disclaims all warranties with regard to this
#  software, including all implied warranties of merchantability and
#  fitness, in no event shall NAIST be liable for
#  any special, indirect or consequential damages or any damages
#  whatsoever resulting from loss of use, data or profits, whether in an
#  action of contract, negligence or other tortuous action, arising out
#  of or in connection with the use or performance of this software.
#
#  A large portion of the dictionary entries
#  originate from ICOT Free Software.  The following conditions for ICOT
#  Free Software applies to the current dictionary as well.
#
#  Each User may also freely distribute the Program, whether in its
#  original form or modified, to any third party or parties, PROVIDED
#  that the provisions of Section 3 ("NO WARRANTY") will ALWAYS appear
#  on, or be attached to, the Program, which is distributed substantially
#  in the same form as set out herein and that such intended
#  distribution, if actually made, will neither violate or otherwise
#  contravene any of the laws and regulations of the countries having
#  jurisdiction over the User or the intended distribution itself.
#
#  NO WARRANTY
#
#  The program was produced on an experimental basis in the course of the
#  research and development conducted during the project and is provided
#  to users as so produced on an experimental basis.  Accordingly, the
#  program is provided without any warranty whatsoever, whether express,
#  implied, statutory or otherwise.  The term "warranty" used herein
#  includes, but is not limited to, any warranty of the quality,
#  performance, merchantability and fitness for a particular purpose of
#  the program and the nonexistence of any infringement or violation of
#  any right of any third party.
#
#  Each user of the program will agree and understand, and be deemed to
#  have agreed and understood, that there is no warranty whatsoever for
#  the program and, accordingly, the entire risk arising from or
#  otherwise connected with the program is assumed by the user.
#
```

(continues on next page)

(continued from previous page)

```
 #  Therefore, neither ICOT, the copyright holder, or any other
 #  organization that participated in or was otherwise related to the
 #  development of the program and their respective officials, directors,
 #  officers and other employees shall be held liable for any and all
 #  damages, including, without limitation, general, special, incidental
 #  and consequential damages, arising out of or otherwise in connection
 #  with the use or inability to use the program or any product, material
 #  or result produced or otherwise obtained by using the program,
 #  regardless of whether they have been advised of, or otherwise had
 #  knowledge of, the possibility of such damages at any time during the
 #  project or thereafter.  Each user will be deemed to have agreed to the
 #  foregoing by his or her commencement of use of the program.  The term
 #  "use" as used herein includes, but is not limited to, the use,
 #  modification, copying and distribution of the program and the
 #  production of secondary products from the program.
 #
 #  In the case where the program, whether in its original form or
 #  modified, was distributed or delivered to or received by a user from
 #  any person, organization or entity other than ICOT, unless it makes or
 #  grants independently of ICOT any specific warranty to the user in
 #  writing, such person, organization or entity, will also be exempted
 #  from and not be held liable to the user for any such damages as noted
 #  above as far as the program is concerned.
 #
 #  ---------------COPYING.ipadic-----END----------------------------------

3. Lao Word Break Dictionary Data (laodict.txt)

 # Copyright (C) 2016 and later: Unicode, Inc. and others.
 # License & terms of use: http://www.unicode.org/copyright.html
 # Copyright (c) 2015 International Business Machines Corporation
 # and others. All Rights Reserved.
 #
 # Project: https://github.com/rober42539/lao-dictionary
 # Dictionary: https://github.com/rober42539/lao-dictionary/laodict.txt
 # License: https://github.com/rober42539/lao-dictionary/LICENSE.txt
 #          (copied below)
 #
 #        This file is derived from the above dictionary version of Nov 22, 2020
 #  --------------------------------------------------------------------
 #  Copyright (C) 2013 Brian Eugene Wilson, Robert Martin Campbell.
 #  All rights reserved.
 #
 #  Redistribution and use in source and binary forms, with or without
 #  modification, are permitted provided that the following conditions are met:
 #
 #  Redistributions of source code must retain the above copyright notice, this
 #  list of conditions and the following disclaimer. Redistributions in binary
 #  form must reproduce the above copyright notice, this list of conditions and
 #  the following disclaimer in the documentation and/or other materials
 #  provided with the distribution.
 #
```

(continues on next page)

(continued from previous page)

```
 # THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
 # "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
 # LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
 # FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE
 # COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
 # INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
 # (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
 # SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 # HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
 # STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
 # ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
 # OF THE POSSIBILITY OF SUCH DAMAGE.
 #  --------------------------------------------------------------------------

4. Burmese Word Break Dictionary Data (burmesedict.txt)

 #  Copyright (c) 2014 International Business Machines Corporation
 #  and others. All Rights Reserved.
 #
 #  This list is part of a project hosted at:
 #    github.com/kanyawtech/myanmar-karen-word-lists
 #
 #  --------------------------------------------------------------------------
 #  Copyright (c) 2013, LeRoy Benjamin Sharon
 #  All rights reserved.
 #
 #  Redistribution and use in source and binary forms, with or without
 #  modification, are permitted provided that the following conditions
 #  are met: Redistributions of source code must retain the above
 #  copyright notice, this list of conditions and the following
 #  disclaimer.  Redistributions in binary form must reproduce the
 #  above copyright notice, this list of conditions and the following
 #  disclaimer in the documentation and/or other materials provided
 #  with the distribution.
 #
 #    Neither the name Myanmar Karen Word Lists, nor the names of its
 #    contributors may be used to endorse or promote products derived
 #    from this software without specific prior written permission.
 #
 #  THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND
 #  CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES,
 #  INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
 #  MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
 #  DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS
 #  BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
 #  EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED
 #  TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
 #  DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON
 #  ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR
 #  TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF
 #  THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 #  SUCH DAMAGE.
```

(continues on next page)

(continued from previous page)

```
 #  ------------------------------------------------------------------------------

5. Time Zone Database

  ICU uses the public domain data and code derived from Time Zone
Database for its time zone support. The ownership of the TZ database
is explained in BCP 175: Procedure for Maintaining the Time Zone
Database section 7.

 # 7.  Database Ownership
 #
 #    The TZ database itself is not an IETF Contribution or an IETF
 #    document.  Rather it is a pre-existing and regularly updated work
 #    that is in the public domain, and is intended to remain in the
 #    public domain.  Therefore, BCPs 78 [RFC5378] and 79 [RFC3979] do
 #    not apply to the TZ Database or contributions that individuals make
 #    to it.  Should any claims be made and substantiated against the TZ
 #    Database, the organization that is providing the IANA
 #    Considerations defined in this RFC, under the memorandum of
 #    understanding with the IETF, currently ICANN, may act in accordance
 #    with all competent court orders.  No ownership claims will be made
 #    by ICANN or the IETF Trust on the database or the code.  Any person
 #    making a contribution to the database or code waives all rights to
 #    future claims in that contribution or in the TZ Database.

6. Google double-conversion

Copyright 2006-2011, the V8 project authors. All rights reserved.
Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are
met:

    * Redistributions of source code must retain the above copyright
      notice, this list of conditions and the following disclaimer.
    * Redistributions in binary form must reproduce the above
      copyright notice, this list of conditions and the following
      disclaimer in the documentation and/or other materials provided
      with the distribution.
    * Neither the name of Google Inc. nor the names of its
      contributors may be used to endorse or promote products derived
      from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
```

(continues on next page)

(continued from previous page)

```
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.27 GPLv2 License Text for iftop

The following license text applies to iftop:

Listing 27: Download: GPLv2 for iftop

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.
```

(continues on next page)

(continued from previous page)

```
  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:
```

(continues on next page)

(continued from previous page)

```
    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
```

(continues on next page)

(continued from previous page)

```
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
```

(continues on next page)

(continued from previous page)

```
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
```

(continues on next page)

(continued from previous page)

```
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
```

(continues on next page)

(continued from previous page)

```
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.28 GPLv2+ License Text for igmpproxy

The following license text applies to igmpproxy:

Listing 28: Download: GPLv2+ for igmpproxy

```
igmpproxy - IGMP proxy based multicast router
Copyright (C) 2005 Johnny Egeland <johnny@rlo.org>

This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; either version 2 of the License, or
```

(continues on next page)

(continued from previous page)

```
(at your option) any later version.

This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
GNU General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA


This software is derived work from the following software. The original
source code has been modified from it's original state by the author
of igmpproxy.

smcroute 0.92 - Copyright (C) 2001 Carsten Schill <carsten@cschill.de>
- Licensed under the GNU General Public License, either version 2 or
  any later version.

mrouted 3.9-beta3 - Copyright (C) 2002 by The Board of Trustees of
Leland Stanford Junior University.
- Licensed under the 3-clause BSD license, see Stanford.txt file.


Since 2017-03-25 igmpproxy is GPLv2+ compatible, mrouted licensed
was switched from the proprietary Stanford to 3-clause BSD. New
igmpproxy contributions and patches must be under GPLv2+ license,
old proprietary Stanford is not accepted.
```

40.29 BSD2CLAUSE License Text for indexinfo

The following license text applies to indexinfo:

Listing 29: Download: BSD2CLAUSE for indexinfo

```
The compilation of software known as indexinfo is distributed under the
following terms:

Copyright (c) 2014 Baptiste Daroussin <bapt@FreeBSD.org>.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
```

(continues on next page)

(continued from previous page)

```
THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.
```

40.30 BSD3CLAUSE License Text for ipmitool

The following license text applies to ipmitool:

Listing 30: Download: BSD3CLAUSE for ipmitool

```
Copyright (c) 2003 Sun Microsystems, Inc.  All Rights Reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

Redistribution of source code must retain the above copyright
notice, this list of conditions and the following disclaimer.

Redistribution in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in the
documentation and/or other materials provided with the distribution.

Neither the name of Sun Microsystems, Inc. or the names of
contributors may be used to endorse or promote products derived
from this software without specific prior written permission.

This software is provided "AS IS," without a warranty of any kind.
ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES,
INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A
PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE HEREBY EXCLUDED.
SUN MICROSYSTEMS, INC. ("SUN") AND ITS LICENSORS SHALL NOT BE LIABLE
FOR ANY DAMAGES SUFFERED BY LICENSEE AS A RESULT OF USING, MODIFYING
OR DISTRIBUTING THIS SOFTWARE OR ITS DERIVATIVES.  IN NO EVENT WILL
SUN OR ITS LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT OR DATA,
OR FOR DIRECT, INDIRECT, SPECIAL, CONSEQUENTIAL, INCIDENTAL OR
PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF
LIABILITY, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS SOFTWARE,
EVEN IF SUN HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.
```

40.31 MPL20 License Text for isc-dhcp44-client

The following license text applies to isc-dhcp44-client:

Listing 31: Download: MPL20 for isc-dhcp44-client

```
Mozilla Public License Version 2.0
==================================

1. Definitions
--------------

1.1. "Contributor"
    means each individual or legal entity that creates, contributes to
    the creation of, or owns Covered Software.

1.2. "Contributor Version"
    means the combination of the Contributions of others (if any) used
    by a Contributor and that particular Contributor's Contribution.

1.3. "Contribution"
    means Covered Software of a particular Contributor.

1.4. "Covered Software"
    means Source Code Form to which the initial Contributor has attached
    the notice in Exhibit A, the Executable Form of such Source Code
    Form, and Modifications of such Source Code Form, in each case
    including portions thereof.

1.5. "Incompatible With Secondary Licenses"
    means

    (a) that the initial Contributor has attached the notice described
        in Exhibit B to the Covered Software; or

    (b) that the Covered Software was made available under the terms of
        version 1.1 or earlier of the License, but not also under the
        terms of a Secondary License.

1.6. "Executable Form"
    means any form of the work other than Source Code Form.

1.7. "Larger Work"
    means a work that combines Covered Software with other material, in
    a separate file or files, that is not Covered Software.

1.8. "License"
    means this document.

1.9. "Licensable"
    means having the right to grant, to the maximum extent possible,
    whether at the time of the initial grant or subsequently, any and
    all of the rights conveyed by this License.
```

(continues on next page)

(continued from previous page)

```
1.10. "Modifications"
    means any of the following:

    (a) any file in Source Code Form that results from an addition to,
        deletion from, or modification of the contents of Covered
        Software; or

    (b) any new file in Source Code Form that contains any Covered
        Software.

1.11. "Patent Claims" of a Contributor
    means any patent claim(s), including without limitation, method,
    process, and apparatus claims, in any patent Licensable by such
    Contributor that would be infringed, but for the grant of the
    License, by the making, using, selling, offering for sale, having
    made, import, or transfer of either its Contributions or its
    Contributor Version.

1.12. "Secondary License"
    means either the GNU General Public License, Version 2.0, the GNU
    Lesser General Public License, Version 2.1, the GNU Affero General
    Public License, Version 3.0, or any later versions of those
    licenses.

1.13. "Source Code Form"
    means the form of the work preferred for making modifications.

1.14. "You" (or "Your")
    means an individual or a legal entity exercising rights under this
    License. For legal entities, "You" includes any entity that
    controls, is controlled by, or is under common control with You. For
    purposes of this definition, "control" means (a) the power, direct
    or indirect, to cause the direction or management of such entity,
    whether by contract or otherwise, or (b) ownership of more than
    fifty percent (50%) of the outstanding shares or beneficial
    ownership of such entity.

2. License Grants and Conditions
--------------------------------

2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free,
non-exclusive license:

(a) under intellectual property rights (other than patent or trademark)
    Licensable by such Contributor to use, reproduce, make available,
    modify, display, perform, distribute, and otherwise exploit its
    Contributions, either on an unmodified basis, with Modifications, or
    as part of a Larger Work; and
```

(continues on next page)

(continued from previous page)

```
(b) under Patent Claims of such Contributor to make, use, sell, offer
    for sale, have made, import, and otherwise transfer either its
    Contributions or its Contributor Version.

2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution
become effective for each Contribution on the date the Contributor first
distributes such Contribution.

2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under
this License. No additional rights or licenses will be implied from the
distribution or licensing of Covered Software under this License.
Notwithstanding Section 2.1(b) above, no patent license is granted by a
Contributor:

(a) for any code that a Contributor has removed from Covered Software;
    or

(b) for infringements caused by: (i) Your and any other third party's
    modifications of Covered Software, or (ii) the combination of its
    Contributions with other software (except as part of its Contributor
    Version); or

(c) under Patent Claims infringed by Covered Software in the absence of
    its Contributions.

This License does not grant any rights in the trademarks, service marks,
or logos of any Contributor (except as may be necessary to comply with
the notice requirements in Section 3.4).

2.4. Subsequent Licenses

No Contributor makes additional grants as a result of Your choice to
distribute the Covered Software under a subsequent version of this
License (see Section 10.2) or under the terms of a Secondary License (if
permitted under the terms of Section 3.3).

2.5. Representation

Each Contributor represents that the Contributor believes its
Contributions are its original creation(s) or it has sufficient rights
to grant the rights to its Contributions conveyed by this License.

2.6. Fair Use

This License is not intended to limit any rights You have under
applicable copyright doctrines of fair use, fair dealing, or other
equivalents.
```

(continues on next page)

(continued from previous page)

```
2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted
in Section 2.1.

3. Responsibilities
-------------------

3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any
Modifications that You create or to which You contribute, must be under
the terms of this License. You must inform recipients that the Source
Code Form of the Covered Software is governed by the terms of this
License, and how they can obtain a copy of this License. You may not
attempt to alter or restrict the recipients' rights in the Source Code
Form.

3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:

(a) such Covered Software must also be made available in Source Code
    Form, as described in Section 3.1, and You must inform recipients of
    the Executable Form how they can obtain a copy of such Source Code
    Form by reasonable means in a timely manner, at a charge no more
    than the cost of distribution to the recipient; and

(b) You may distribute such Executable Form under the terms of this
    License, or sublicense it under different terms, provided that the
    license for the Executable Form does not attempt to limit or alter
    the recipients' rights in the Source Code Form under this License.

3.3. Distribution of a Larger Work

You may create and distribute a Larger Work under terms of Your choice,
provided that You also comply with the requirements of this License for
the Covered Software. If the Larger Work is a combination of Covered
Software with a work governed by one or more Secondary Licenses, and the
Covered Software is not Incompatible With Secondary Licenses, this
License permits You to additionally distribute such Covered Software
under the terms of such Secondary License(s), so that the recipient of
the Larger Work may, at their option, further distribute the Covered
Software under the terms of either this License or such Secondary
License(s).

3.4. Notices

You may not remove or alter the substance of any license notices
(including copyright notices, patent notices, disclaimers of warranty,
or limitations of liability) contained within the Source Code Form of
the Covered Software, except that You may alter any license notices to
```

(continues on next page)

(continued from previous page)

```
the extent required to remedy known factual inaccuracies.

3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support,
indemnity or liability obligations to one or more recipients of Covered
Software. However, You may do so only on Your own behalf, and not on
behalf of any Contributor. You must make it absolutely clear that any
such warranty, support, indemnity, or liability obligation is offered by
You alone, and You hereby agree to indemnify every Contributor for any
liability incurred by such Contributor as a result of warranty, support,
indemnity or liability terms You offer. You may include additional
disclaimers of warranty and limitations of liability specific to any
jurisdiction.

4. Inability to Comply Due to Statute or Regulation
---------------------------------------------------

If it is impossible for You to comply with any of the terms of this
License with respect to some or all of the Covered Software due to
statute, judicial order, or regulation then You must: (a) comply with
the terms of this License to the maximum extent possible; and (b)
describe the limitations and the code they affect. Such description must
be placed in a text file included with all distributions of the Covered
Software under this License. Except to the extent prohibited by statute
or regulation, such description must be sufficiently detailed for a
recipient of ordinary skill to be able to understand it.

5. Termination
--------------

5.1. The rights granted under this License will terminate automatically
if You fail to comply with any of its terms. However, if You become
compliant, then the rights granted under this License from a particular
Contributor are reinstated (a) provisionally, unless and until such
Contributor explicitly and finally terminates Your grants, and (b) on an
ongoing basis, if such Contributor fails to notify You of the
non-compliance by some reasonable means prior to 60 days after You have
come back into compliance. Moreover, Your grants from a particular
Contributor are reinstated on an ongoing basis if such Contributor
notifies You of the non-compliance by some reasonable means, this is the
first time You have received notice of non-compliance with this License
from such Contributor, and You become compliant prior to 30 days after
Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent
infringement claim (excluding declaratory judgment actions,
counter-claims, and cross-claims) alleging that a Contributor Version
directly or indirectly infringes any patent, then the rights granted to
You by any and all Contributors for the Covered Software under Section
2.1 of this License shall terminate.
```

(continues on next page)

(continued from previous page)

```
5.3. In the event of termination under Sections 5.1 or 5.2 above, all
end user license agreements (excluding distributors and resellers) which
have been validly granted by You or Your distributors under this License
prior to termination shall survive termination.

************************************************************************
*                                                                      *
*  6. Disclaimer of Warranty                                           *
*  -------------------------                                           *
*                                                                      *
*  Covered Software is provided under this License on an "as is"       *
*  basis, without warranty of any kind, either expressed, implied, or  *
*  statutory, including, without limitation, warranties that the       *
*  Covered Software is free of defects, merchantable, fit for a        *
*  particular purpose or non-infringing. The entire risk as to the     *
*  quality and performance of the Covered Software is with You.        *
*  Should any Covered Software prove defective in any respect, You     *
*  (not any Contributor) assume the cost of any necessary servicing,   *
*  repair, or correction. This disclaimer of warranty constitutes an   *
*  essential part of this License. No use of any Covered Software is   *
*  authorized under this License except under this disclaimer.         *
*                                                                      *
************************************************************************

************************************************************************
*                                                                      *
*  7. Limitation of Liability                                          *
*  --------------------------                                          *
*                                                                      *
*  Under no circumstances and under no legal theory, whether tort      *
*  (including negligence), contract, or otherwise, shall any           *
*  Contributor, or anyone who distributes Covered Software as          *
*  permitted above, be liable to You for any direct, indirect,         *
*  special, incidental, or consequential damages of any character      *
*  including, without limitation, damages for lost profits, loss of    *
*  goodwill, work stoppage, computer failure or malfunction, or any    *
*  and all other commercial damages or losses, even if such party      *
*  shall have been informed of the possibility of such damages. This   *
*  limitation of liability shall not apply to liability for death or   *
*  personal injury resulting from such party's negligence to the       *
*  extent applicable law prohibits such limitation. Some               *
*  jurisdictions do not allow the exclusion or limitation of           *
*  incidental or consequential damages, so this exclusion and          *
*  limitation may not apply to You.                                    *
*                                                                      *
************************************************************************

8. Litigation
-------------

Any litigation relating to this License may be brought only in the
courts of a jurisdiction where the defendant maintains its principal
```

(continues on next page)

(continued from previous page)

```
place of business and such litigation shall be governed by laws of that
jurisdiction, without reference to its conflict-of-law provisions.
Nothing in this Section shall prevent a party's ability to bring
cross-claims or counter-claims.

9. Miscellaneous
----------------

This License represents the complete agreement concerning the subject
matter hereof. If any provision of this License is held to be
unenforceable, such provision shall be reformed only to the extent
necessary to make it enforceable. Any law or regulation which provides
that the language of a contract shall be construed against the drafter
shall not be used to construe this License against a Contributor.

10. Versions of the License
---------------------------

10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section
10.3, no one other than the license steward has the right to modify or
publish new versions of this License. Each version will be given a
distinguishing version number.

10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version
of the License under which You originally received the Covered Software,
or under the terms of any subsequent version published by the license
steward.

10.3. Modified Versions

If you create software not governed by this License, and you want to
create a new license for such software, you may create and use a
modified version of this License if you rename the license and remove
any references to the name of the license steward (except to note that
such modified license differs from this License).

10.4. Distributing Source Code Form that is Incompatible With Secondary
Licenses

If You choose to distribute Source Code Form that is Incompatible With
Secondary Licenses under the terms of this version of the License, the
notice described in Exhibit B of this License must be attached.

Exhibit A - Source Code Form License Notice
-------------------------------------------

  This Source Code Form is subject to the terms of the Mozilla Public
  License, v. 2.0. If a copy of the MPL was not distributed with this
```

(continues on next page)

(continued from previous page)

```
  file, You can obtain one at http://mozilla.org/MPL/2.0/.

If it is not possible or desirable to put the notice in a particular
file, then You may include the notice in a location (such as a LICENSE
file in a relevant directory) where a recipient would be likely to look
for such a notice.

You may add additional accurate notices of copyright ownership.

Exhibit B - "Incompatible With Secondary Licenses" Notice
---------------------------------------------------------

  This Source Code Form is "Incompatible With Secondary Licenses", as
  defined by the Mozilla Public License, v. 2.0.
```

40.32 MPL20 License Text for isc-dhcp44-relay

The following license text applies to isc-dhcp44-relay:

Listing 32: Download: `MPL20` for isc-dhcp44-relay

```
Mozilla Public License Version 2.0
==================================

1. Definitions
--------------

1.1. "Contributor"
    means each individual or legal entity that creates, contributes to
    the creation of, or owns Covered Software.

1.2. "Contributor Version"
    means the combination of the Contributions of others (if any) used
    by a Contributor and that particular Contributor's Contribution.

1.3. "Contribution"
    means Covered Software of a particular Contributor.

1.4. "Covered Software"
    means Source Code Form to which the initial Contributor has attached
    the notice in Exhibit A, the Executable Form of such Source Code
    Form, and Modifications of such Source Code Form, in each case
    including portions thereof.

1.5. "Incompatible With Secondary Licenses"
    means

    (a) that the initial Contributor has attached the notice described
        in Exhibit B to the Covered Software; or
```

(continues on next page)

(continued from previous page)

```
    (b) that the Covered Software was made available under the terms of
        version 1.1 or earlier of the License, but not also under the
        terms of a Secondary License.

1.6. "Executable Form"
    means any form of the work other than Source Code Form.

1.7. "Larger Work"
    means a work that combines Covered Software with other material, in
    a separate file or files, that is not Covered Software.

1.8. "License"
    means this document.

1.9. "Licensable"
    means having the right to grant, to the maximum extent possible,
    whether at the time of the initial grant or subsequently, any and
    all of the rights conveyed by this License.

1.10. "Modifications"
    means any of the following:

    (a) any file in Source Code Form that results from an addition to,
        deletion from, or modification of the contents of Covered
        Software; or

    (b) any new file in Source Code Form that contains any Covered
        Software.

1.11. "Patent Claims" of a Contributor
    means any patent claim(s), including without limitation, method,
    process, and apparatus claims, in any patent Licensable by such
    Contributor that would be infringed, but for the grant of the
    License, by the making, using, selling, offering for sale, having
    made, import, or transfer of either its Contributions or its
    Contributor Version.

1.12. "Secondary License"
    means either the GNU General Public License, Version 2.0, the GNU
    Lesser General Public License, Version 2.1, the GNU Affero General
    Public License, Version 3.0, or any later versions of those
    licenses.

1.13. "Source Code Form"
    means the form of the work preferred for making modifications.

1.14. "You" (or "Your")
    means an individual or a legal entity exercising rights under this
    License. For legal entities, "You" includes any entity that
    controls, is controlled by, or is under common control with You. For
    purposes of this definition, "control" means (a) the power, direct
    or indirect, to cause the direction or management of such entity,
```

(continues on next page)

(continued from previous page)

```
    whether by contract or otherwise, or (b) ownership of more than
    fifty percent (50%) of the outstanding shares or beneficial
    ownership of such entity.

2. License Grants and Conditions
--------------------------------

2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free,
non-exclusive license:

(a) under intellectual property rights (other than patent or trademark)
    Licensable by such Contributor to use, reproduce, make available,
    modify, display, perform, distribute, and otherwise exploit its
    Contributions, either on an unmodified basis, with Modifications, or
    as part of a Larger Work; and

(b) under Patent Claims of such Contributor to make, use, sell, offer
    for sale, have made, import, and otherwise transfer either its
    Contributions or its Contributor Version.

2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution
become effective for each Contribution on the date the Contributor first
distributes such Contribution.

2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under
this License. No additional rights or licenses will be implied from the
distribution or licensing of Covered Software under this License.
Notwithstanding Section 2.1(b) above, no patent license is granted by a
Contributor:

(a) for any code that a Contributor has removed from Covered Software;
    or

(b) for infringements caused by: (i) Your and any other third party's
    modifications of Covered Software, or (ii) the combination of its
    Contributions with other software (except as part of its Contributor
    Version); or

(c) under Patent Claims infringed by Covered Software in the absence of
    its Contributions.

This License does not grant any rights in the trademarks, service marks,
or logos of any Contributor (except as may be necessary to comply with
the notice requirements in Section 3.4).

2.4. Subsequent Licenses
```

(continues on next page)

(continued from previous page)

```
No Contributor makes additional grants as a result of Your choice to
distribute the Covered Software under a subsequent version of this
License (see Section 10.2) or under the terms of a Secondary License (if
permitted under the terms of Section 3.3).

2.5. Representation

Each Contributor represents that the Contributor believes its
Contributions are its original creation(s) or it has sufficient rights
to grant the rights to its Contributions conveyed by this License.

2.6. Fair Use

This License is not intended to limit any rights You have under
applicable copyright doctrines of fair use, fair dealing, or other
equivalents.

2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted
in Section 2.1.

3. Responsibilities
-------------------

3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any
Modifications that You create or to which You contribute, must be under
the terms of this License. You must inform recipients that the Source
Code Form of the Covered Software is governed by the terms of this
License, and how they can obtain a copy of this License. You may not
attempt to alter or restrict the recipients' rights in the Source Code
Form.

3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:

(a) such Covered Software must also be made available in Source Code
    Form, as described in Section 3.1, and You must inform recipients of
    the Executable Form how they can obtain a copy of such Source Code
    Form by reasonable means in a timely manner, at a charge no more
    than the cost of distribution to the recipient; and

(b) You may distribute such Executable Form under the terms of this
    License, or sublicense it under different terms, provided that the
    license for the Executable Form does not attempt to limit or alter
    the recipients' rights in the Source Code Form under this License.

3.3. Distribution of a Larger Work
```

(continues on next page)

(continued from previous page)

```
You may create and distribute a Larger Work under terms of Your choice,
provided that You also comply with the requirements of this License for
the Covered Software. If the Larger Work is a combination of Covered
Software with a work governed by one or more Secondary Licenses, and the
Covered Software is not Incompatible With Secondary Licenses, this
License permits You to additionally distribute such Covered Software
under the terms of such Secondary License(s), so that the recipient of
the Larger Work may, at their option, further distribute the Covered
Software under the terms of either this License or such Secondary
License(s).

3.4. Notices

You may not remove or alter the substance of any license notices
(including copyright notices, patent notices, disclaimers of warranty,
or limitations of liability) contained within the Source Code Form of
the Covered Software, except that You may alter any license notices to
the extent required to remedy known factual inaccuracies.

3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support,
indemnity or liability obligations to one or more recipients of Covered
Software. However, You may do so only on Your own behalf, and not on
behalf of any Contributor. You must make it absolutely clear that any
such warranty, support, indemnity, or liability obligation is offered by
You alone, and You hereby agree to indemnify every Contributor for any
liability incurred by such Contributor as a result of warranty, support,
indemnity or liability terms You offer. You may include additional
disclaimers of warranty and limitations of liability specific to any
jurisdiction.

4. Inability to Comply Due to Statute or Regulation
---------------------------------------------------

If it is impossible for You to comply with any of the terms of this
License with respect to some or all of the Covered Software due to
statute, judicial order, or regulation then You must: (a) comply with
the terms of this License to the maximum extent possible; and (b)
describe the limitations and the code they affect. Such description must
be placed in a text file included with all distributions of the Covered
Software under this License. Except to the extent prohibited by statute
or regulation, such description must be sufficiently detailed for a
recipient of ordinary skill to be able to understand it.

5. Termination
--------------

5.1. The rights granted under this License will terminate automatically
if You fail to comply with any of its terms. However, if You become
compliant, then the rights granted under this License from a particular
```

(continues on next page)

(continued from previous page)

```
Contributor are reinstated (a) provisionally, unless and until such
Contributor explicitly and finally terminates Your grants, and (b) on an
ongoing basis, if such Contributor fails to notify You of the
non-compliance by some reasonable means prior to 60 days after You have
come back into compliance. Moreover, Your grants from a particular
Contributor are reinstated on an ongoing basis if such Contributor
notifies You of the non-compliance by some reasonable means, this is the
first time You have received notice of non-compliance with this License
from such Contributor, and You become compliant prior to 30 days after
Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent
infringement claim (excluding declaratory judgment actions,
counter-claims, and cross-claims) alleging that a Contributor Version
directly or indirectly infringes any patent, then the rights granted to
You by any and all Contributors for the Covered Software under Section
2.1 of this License shall terminate.

5.3. In the event of termination under Sections 5.1 or 5.2 above, all
end user license agreements (excluding distributors and resellers) which
have been validly granted by You or Your distributors under this License
prior to termination shall survive termination.

************************************************************************
*                                                                      *
*  6. Disclaimer of Warranty                                           *
*  -------------------------                                           *
*                                                                      *
*  Covered Software is provided under this License on an "as is"       *
*  basis, without warranty of any kind, either expressed, implied, or  *
*  statutory, including, without limitation, warranties that the       *
*  Covered Software is free of defects, merchantable, fit for a        *
*  particular purpose or non-infringing. The entire risk as to the     *
*  quality and performance of the Covered Software is with You.        *
*  Should any Covered Software prove defective in any respect, You     *
*  (not any Contributor) assume the cost of any necessary servicing,   *
*  repair, or correction. This disclaimer of warranty constitutes an   *
*  essential part of this License. No use of any Covered Software is   *
*  authorized under this License except under this disclaimer.         *
*                                                                      *
************************************************************************

************************************************************************
*                                                                      *
*  7. Limitation of Liability                                          *
*  --------------------------                                          *
*                                                                      *
*  Under no circumstances and under no legal theory, whether tort      *
*  (including negligence), contract, or otherwise, shall any           *
*  Contributor, or anyone who distributes Covered Software as          *
*  permitted above, be liable to You for any direct, indirect,         *
*  special, incidental, or consequential damages of any character      *
```

(continues on next page)

(continued from previous page)

```
*  including, without limitation, damages for lost profits, loss of    *
*  goodwill, work stoppage, computer failure or malfunction, or any    *
*  and all other commercial damages or losses, even if such party      *
*  shall have been informed of the possibility of such damages. This   *
*  limitation of liability shall not apply to liability for death or   *
*  personal injury resulting from such party's negligence to the       *
*  extent applicable law prohibits such limitation. Some               *
*  jurisdictions do not allow the exclusion or limitation of           *
*  incidental or consequential damages, so this exclusion and          *
*  limitation may not apply to You.                                    *
*                                                                      *
************************************************************************

8. Litigation
-------------

Any litigation relating to this License may be brought only in the
courts of a jurisdiction where the defendant maintains its principal
place of business and such litigation shall be governed by laws of that
jurisdiction, without reference to its conflict-of-law provisions.
Nothing in this Section shall prevent a party's ability to bring
cross-claims or counter-claims.

9. Miscellaneous
----------------

This License represents the complete agreement concerning the subject
matter hereof. If any provision of this License is held to be
unenforceable, such provision shall be reformed only to the extent
necessary to make it enforceable. Any law or regulation which provides
that the language of a contract shall be construed against the drafter
shall not be used to construe this License against a Contributor.

10. Versions of the License
---------------------------

10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section
10.3, no one other than the license steward has the right to modify or
publish new versions of this License. Each version will be given a
distinguishing version number.

10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version
of the License under which You originally received the Covered Software,
or under the terms of any subsequent version published by the license
steward.

10.3. Modified Versions
```

(continues on next page)

(continued from previous page)

```
If you create software not governed by this License, and you want to
create a new license for such software, you may create and use a
modified version of this License if you rename the license and remove
any references to the name of the license steward (except to note that
such modified license differs from this License).

10.4. Distributing Source Code Form that is Incompatible With Secondary
Licenses

If You choose to distribute Source Code Form that is Incompatible With
Secondary Licenses under the terms of this version of the License, the
notice described in Exhibit B of this License must be attached.

Exhibit A - Source Code Form License Notice
-------------------------------------------

  This Source Code Form is subject to the terms of the Mozilla Public
  License, v. 2.0. If a copy of the MPL was not distributed with this
  file, You can obtain one at http://mozilla.org/MPL/2.0/.

If it is not possible or desirable to put the notice in a particular
file, then You may include the notice in a location (such as a LICENSE
file in a relevant directory) where a recipient would be likely to look
for such a notice.

You may add additional accurate notices of copyright ownership.

Exhibit B - "Incompatible With Secondary Licenses" Notice
---------------------------------------------------------

  This Source Code Form is "Incompatible With Secondary Licenses", as
  defined by the Mozilla Public License, v. 2.0.
```

40.33 MPL20 License Text for isc-dhcp44-server

The following license text applies to isc-dhcp44-server:

Listing 33: Download: `MPL20` for isc-dhcp44-server

```
Mozilla Public License Version 2.0
==================================

1. Definitions
--------------

1.1. "Contributor"
    means each individual or legal entity that creates, contributes to
    the creation of, or owns Covered Software.

1.2. "Contributor Version"
```

(continues on next page)

(continued from previous page)

```
    means the combination of the Contributions of others (if any) used
    by a Contributor and that particular Contributor's Contribution.

1.3. "Contribution"
    means Covered Software of a particular Contributor.

1.4. "Covered Software"
    means Source Code Form to which the initial Contributor has attached
    the notice in Exhibit A, the Executable Form of such Source Code
    Form, and Modifications of such Source Code Form, in each case
    including portions thereof.

1.5. "Incompatible With Secondary Licenses"
    means

    (a) that the initial Contributor has attached the notice described
        in Exhibit B to the Covered Software; or

    (b) that the Covered Software was made available under the terms of
        version 1.1 or earlier of the License, but not also under the
        terms of a Secondary License.

1.6. "Executable Form"
    means any form of the work other than Source Code Form.

1.7. "Larger Work"
    means a work that combines Covered Software with other material, in
    a separate file or files, that is not Covered Software.

1.8. "License"
    means this document.

1.9. "Licensable"
    means having the right to grant, to the maximum extent possible,
    whether at the time of the initial grant or subsequently, any and
    all of the rights conveyed by this License.

1.10. "Modifications"
    means any of the following:

    (a) any file in Source Code Form that results from an addition to,
        deletion from, or modification of the contents of Covered
        Software; or

    (b) any new file in Source Code Form that contains any Covered
        Software.

1.11. "Patent Claims" of a Contributor
    means any patent claim(s), including without limitation, method,
    process, and apparatus claims, in any patent Licensable by such
    Contributor that would be infringed, but for the grant of the
    License, by the making, using, selling, offering for sale, having
```

(continues on next page)

(continued from previous page)

```
    made, import, or transfer of either its Contributions or its
    Contributor Version.

1.12. "Secondary License"
    means either the GNU General Public License, Version 2.0, the GNU
    Lesser General Public License, Version 2.1, the GNU Affero General
    Public License, Version 3.0, or any later versions of those
    licenses.

1.13. "Source Code Form"
    means the form of the work preferred for making modifications.

1.14. "You" (or "Your")
    means an individual or a legal entity exercising rights under this
    License. For legal entities, "You" includes any entity that
    controls, is controlled by, or is under common control with You. For
    purposes of this definition, "control" means (a) the power, direct
    or indirect, to cause the direction or management of such entity,
    whether by contract or otherwise, or (b) ownership of more than
    fifty percent (50%) of the outstanding shares or beneficial
    ownership of such entity.

2. License Grants and Conditions
--------------------------------

2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free,
non-exclusive license:

(a) under intellectual property rights (other than patent or trademark)
    Licensable by such Contributor to use, reproduce, make available,
    modify, display, perform, distribute, and otherwise exploit its
    Contributions, either on an unmodified basis, with Modifications, or
    as part of a Larger Work; and

(b) under Patent Claims of such Contributor to make, use, sell, offer
    for sale, have made, import, and otherwise transfer either its
    Contributions or its Contributor Version.

2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution
become effective for each Contribution on the date the Contributor first
distributes such Contribution.

2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under
this License. No additional rights or licenses will be implied from the
distribution or licensing of Covered Software under this License.
Notwithstanding Section 2.1(b) above, no patent license is granted by a
```

(continues on next page)

(continued from previous page)

```
Contributor:

(a) for any code that a Contributor has removed from Covered Software;
    or

(b) for infringements caused by: (i) Your and any other third party's
    modifications of Covered Software, or (ii) the combination of its
    Contributions with other software (except as part of its Contributor
    Version); or

(c) under Patent Claims infringed by Covered Software in the absence of
    its Contributions.

This License does not grant any rights in the trademarks, service marks,
or logos of any Contributor (except as may be necessary to comply with
the notice requirements in Section 3.4).

2.4. Subsequent Licenses

No Contributor makes additional grants as a result of Your choice to
distribute the Covered Software under a subsequent version of this
License (see Section 10.2) or under the terms of a Secondary License (if
permitted under the terms of Section 3.3).

2.5. Representation

Each Contributor represents that the Contributor believes its
Contributions are its original creation(s) or it has sufficient rights
to grant the rights to its Contributions conveyed by this License.

2.6. Fair Use

This License is not intended to limit any rights You have under
applicable copyright doctrines of fair use, fair dealing, or other
equivalents.

2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted
in Section 2.1.

3. Responsibilities
-------------------

3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any
Modifications that You create or to which You contribute, must be under
the terms of this License. You must inform recipients that the Source
Code Form of the Covered Software is governed by the terms of this
License, and how they can obtain a copy of this License. You may not
attempt to alter or restrict the recipients' rights in the Source Code
```

(continues on next page)

(continued from previous page)

```
Form.

3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:

(a) such Covered Software must also be made available in Source Code
    Form, as described in Section 3.1, and You must inform recipients of
    the Executable Form how they can obtain a copy of such Source Code
    Form by reasonable means in a timely manner, at a charge no more
    than the cost of distribution to the recipient; and

(b) You may distribute such Executable Form under the terms of this
    License, or sublicense it under different terms, provided that the
    license for the Executable Form does not attempt to limit or alter
    the recipients' rights in the Source Code Form under this License.

3.3. Distribution of a Larger Work

You may create and distribute a Larger Work under terms of Your choice,
provided that You also comply with the requirements of this License for
the Covered Software. If the Larger Work is a combination of Covered
Software with a work governed by one or more Secondary Licenses, and the
Covered Software is not Incompatible With Secondary Licenses, this
License permits You to additionally distribute such Covered Software
under the terms of such Secondary License(s), so that the recipient of
the Larger Work may, at their option, further distribute the Covered
Software under the terms of either this License or such Secondary
License(s).

3.4. Notices

You may not remove or alter the substance of any license notices
(including copyright notices, patent notices, disclaimers of warranty,
or limitations of liability) contained within the Source Code Form of
the Covered Software, except that You may alter any license notices to
the extent required to remedy known factual inaccuracies.

3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support,
indemnity or liability obligations to one or more recipients of Covered
Software. However, You may do so only on Your own behalf, and not on
behalf of any Contributor. You must make it absolutely clear that any
such warranty, support, indemnity, or liability obligation is offered by
You alone, and You hereby agree to indemnify every Contributor for any
liability incurred by such Contributor as a result of warranty, support,
indemnity or liability terms You offer. You may include additional
disclaimers of warranty and limitations of liability specific to any
jurisdiction.

4. Inability to Comply Due to Statute or Regulation
```

(continues on next page)

(continued from previous page)

```
--------------------------------------------------

If it is impossible for You to comply with any of the terms of this
License with respect to some or all of the Covered Software due to
statute, judicial order, or regulation then You must: (a) comply with
the terms of this License to the maximum extent possible; and (b)
describe the limitations and the code they affect. Such description must
be placed in a text file included with all distributions of the Covered
Software under this License. Except to the extent prohibited by statute
or regulation, such description must be sufficiently detailed for a
recipient of ordinary skill to be able to understand it.

5. Termination
--------------

5.1. The rights granted under this License will terminate automatically
if You fail to comply with any of its terms. However, if You become
compliant, then the rights granted under this License from a particular
Contributor are reinstated (a) provisionally, unless and until such
Contributor explicitly and finally terminates Your grants, and (b) on an
ongoing basis, if such Contributor fails to notify You of the
non-compliance by some reasonable means prior to 60 days after You have
come back into compliance. Moreover, Your grants from a particular
Contributor are reinstated on an ongoing basis if such Contributor
notifies You of the non-compliance by some reasonable means, this is the
first time You have received notice of non-compliance with this License
from such Contributor, and You become compliant prior to 30 days after
Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent
infringement claim (excluding declaratory judgment actions,
counter-claims, and cross-claims) alleging that a Contributor Version
directly or indirectly infringes any patent, then the rights granted to
You by any and all Contributors for the Covered Software under Section
2.1 of this License shall terminate.

5.3. In the event of termination under Sections 5.1 or 5.2 above, all
end user license agreements (excluding distributors and resellers) which
have been validly granted by You or Your distributors under this License
prior to termination shall survive termination.

************************************************************************
*                                                                      *
*  6. Disclaimer of Warranty                                           *
*  -------------------------                                           *
*                                                                      *
*  Covered Software is provided under this License on an "as is"       *
*  basis, without warranty of any kind, either expressed, implied, or  *
*  statutory, including, without limitation, warranties that the       *
*  Covered Software is free of defects, merchantable, fit for a        *
*  particular purpose or non-infringing. The entire risk as to the     *
*  quality and performance of the Covered Software is with You.        *
```

(continues on next page)

(continued from previous page)

```
*  Should any Covered Software prove defective in any respect, You     *
*  (not any Contributor) assume the cost of any necessary servicing,   *
*  repair, or correction. This disclaimer of warranty constitutes an   *
*  essential part of this License. No use of any Covered Software is   *
*  authorized under this License except under this disclaimer.         *
*                                                                      *
************************************************************************

************************************************************************
*                                                                      *
*  7. Limitation of Liability                                          *
*  --------------------------                                          *
*                                                                      *
*  Under no circumstances and under no legal theory, whether tort      *
*  (including negligence), contract, or otherwise, shall any           *
*  Contributor, or anyone who distributes Covered Software as          *
*  permitted above, be liable to You for any direct, indirect,         *
*  special, incidental, or consequential damages of any character      *
*  including, without limitation, damages for lost profits, loss of    *
*  goodwill, work stoppage, computer failure or malfunction, or any    *
*  and all other commercial damages or losses, even if such party      *
*  shall have been informed of the possibility of such damages. This   *
*  limitation of liability shall not apply to liability for death or   *
*  personal injury resulting from such party's negligence to the       *
*  extent applicable law prohibits such limitation. Some               *
*  jurisdictions do not allow the exclusion or limitation of           *
*  incidental or consequential damages, so this exclusion and          *
*  limitation may not apply to You.                                    *
*                                                                      *
************************************************************************

8. Litigation
-------------

Any litigation relating to this License may be brought only in the
courts of a jurisdiction where the defendant maintains its principal
place of business and such litigation shall be governed by laws of that
jurisdiction, without reference to its conflict-of-law provisions.
Nothing in this Section shall prevent a party's ability to bring
cross-claims or counter-claims.

9. Miscellaneous
----------------

This License represents the complete agreement concerning the subject
matter hereof. If any provision of this License is held to be
unenforceable, such provision shall be reformed only to the extent
necessary to make it enforceable. Any law or regulation which provides
that the language of a contract shall be construed against the drafter
shall not be used to construe this License against a Contributor.

10. Versions of the License
```

(continues on next page)

(continued from previous page)

```
--------------------------

10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section
10.3, no one other than the license steward has the right to modify or
publish new versions of this License. Each version will be given a
distinguishing version number.

10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version
of the License under which You originally received the Covered Software,
or under the terms of any subsequent version published by the license
steward.

10.3. Modified Versions

If you create software not governed by this License, and you want to
create a new license for such software, you may create and use a
modified version of this License if you rename the license and remove
any references to the name of the license steward (except to note that
such modified license differs from this License).

10.4. Distributing Source Code Form that is Incompatible With Secondary
Licenses

If You choose to distribute Source Code Form that is Incompatible With
Secondary Licenses under the terms of this version of the License, the
notice described in Exhibit B of this License must be attached.

Exhibit A - Source Code Form License Notice
-------------------------------------------

  This Source Code Form is subject to the terms of the Mozilla Public
  License, v. 2.0. If a copy of the MPL was not distributed with this
  file, You can obtain one at http://mozilla.org/MPL/2.0/.

If it is not possible or desirable to put the notice in a particular
file, then You may include the notice in a location (such as a LICENSE
file in a relevant directory) where a recipient would be likely to look
for such a notice.

You may add additional accurate notices of copyright ownership.

Exhibit B - "Incompatible With Secondary Licenses" Notice
---------------------------------------------------------

  This Source Code Form is "Incompatible With Secondary Licenses", as
  defined by the Mozilla Public License, v. 2.0.
```

40.34 MIT License Text for json-c

The following license text applies to json-c:

Listing 34: Download: MIT for json-c

```
Copyright (c) 2009-2012 Eric Haszlakiewicz

Permission is hereby granted, free of charge, to any person obtaining a
copy of this software and associated documentation files (the "Software"),
to deal in the Software without restriction, including without limitation
the rights to use, copy, modify, merge, publish, distribute, sublicense,
and/or sell copies of the Software, and to permit persons to whom the
Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included
in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE
SOFTWARE.

----------------------------------------------------------------------

Copyright (c) 2004, 2005 Metaparadigm Pte Ltd

Permission is hereby granted, free of charge, to any person obtaining a
copy of this software and associated documentation files (the "Software"),
to deal in the Software without restriction, including without limitation
the rights to use, copy, modify, merge, publish, distribute, sublicense,
and/or sell copies of the Software, and to permit persons to whom the
Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included
in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE
SOFTWARE.
```

40.35 BSD3CLAUSE License Text for ldns

The following license text applies to ldns:

Listing 35: Download: BSD3CLAUSE for ldns

```
Copyright (c) 2005,2006, NLnetLabs
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

    * Redistributions of source code must retain the above copyright notice,
      this list of conditions and the following disclaimer.
    * Redistributions in binary form must reproduce the above copyright
      notice, this list of conditions and the following disclaimer in the
      documentation and/or other materials provided with the distribution.
    * Neither the name of NLnetLabs nor the names of its
      contributors may be used to endorse or promote products derived from this
      software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.
```

40.36 CC0-1.0 License Text for libargon2

The following license text applies to libargon2:

Listing 36: Download: CC0-1.0 for libargon2

```
Argon2 reference source code package - reference C implementations

Copyright 2015
Daniel Dinu, Dmitry Khovratovich, Jean-Philippe Aumasson, and Samuel Neves

You may use this work under the terms of a Creative Commons CC0 1.0
License/Waiver or the Apache Public License 2.0, at your option. The terms of
these licenses can be found at:

- CC0 1.0 Universal : http://creativecommons.org/publicdomain/zero/1.0
- Apache 2.0        : http://www.apache.org/licenses/LICENSE-2.0
```

(continues on next page)

(continued from previous page)

```
The terms of the licenses are reproduced below.

-------------------------------------------------------------------------------

Creative Commons Legal Code

CC0 1.0 Universal

    CREATIVE COMMONS CORPORATION IS NOT A LAW FIRM AND DOES NOT PROVIDE
    LEGAL SERVICES. DISTRIBUTION OF THIS DOCUMENT DOES NOT CREATE AN
    ATTORNEY-CLIENT RELATIONSHIP. CREATIVE COMMONS PROVIDES THIS
    INFORMATION ON AN "AS-IS" BASIS. CREATIVE COMMONS MAKES NO WARRANTIES
    REGARDING THE USE OF THIS DOCUMENT OR THE INFORMATION OR WORKS
    PROVIDED HEREUNDER, AND DISCLAIMS LIABILITY FOR DAMAGES RESULTING FROM
    THE USE OF THIS DOCUMENT OR THE INFORMATION OR WORKS PROVIDED
    HEREUNDER.

Statement of Purpose

The laws of most jurisdictions throughout the world automatically confer
exclusive Copyright and Related Rights (defined below) upon the creator
and subsequent owner(s) (each and all, an "owner") of an original work of
authorship and/or a database (each, a "Work").

Certain owners wish to permanently relinquish those rights to a Work for
the purpose of contributing to a commons of creative, cultural and
scientific works ("Commons") that the public can reliably and without fear
of later claims of infringement build upon, modify, incorporate in other
works, reuse and redistribute as freely as possible in any form whatsoever
and for any purposes, including without limitation commercial purposes.
These owners may contribute to the Commons to promote the ideal of a free
culture and the further production of creative, cultural and scientific
works, or to gain reputation or greater distribution for their Work in
part through the use and efforts of others.

For these and/or other purposes and motivations, and without any
expectation of additional consideration or compensation, the person
associating CC0 with a Work (the "Affirmer"), to the extent that he or she
is an owner of Copyright and Related Rights in the Work, voluntarily
elects to apply CC0 to the Work and publicly distribute the Work under its
terms, with knowledge of his or her Copyright and Related Rights in the
Work and the meaning and intended legal effect of CC0 on those rights.

1. Copyright and Related Rights. A Work made available under CC0 may be
protected by copyright and related or neighboring rights ("Copyright and
Related Rights"). Copyright and Related Rights include, but are not
limited to, the following:

  i. the right to reproduce, adapt, distribute, perform, display,
     communicate, and translate a Work;
 ii. moral rights retained by the original author(s) and/or performer(s);
iii. publicity and privacy rights pertaining to a person's image or
```

(continues on next page)

(continued from previous page)

```
     likeness depicted in a Work;
 iv. rights protecting against unfair competition in regards to a Work,
     subject to the limitations in paragraph 4(a), below;
  v. rights protecting the extraction, dissemination, use and reuse of data
     in a Work;
 vi. database rights (such as those arising under Directive 96/9/EC of the
     European Parliament and of the Council of 11 March 1996 on the legal
     protection of databases, and under any national implementation
     thereof, including any amended or successor version of such
     directive); and
vii. other similar, equivalent or corresponding rights throughout the
     world based on applicable law or treaty, and any national
     implementations thereof.

2. Waiver. To the greatest extent permitted by, but not in contravention
of, applicable law, Affirmer hereby overtly, fully, permanently,
irrevocably and unconditionally waives, abandons, and surrenders all of
Affirmer's Copyright and Related Rights and associated claims and causes
of action, whether now known or unknown (including existing as well as
future claims and causes of action), in the Work (i) in all territories
worldwide, (ii) for the maximum duration provided by applicable law or
treaty (including future time extensions), (iii) in any current or future
medium and for any number of copies, and (iv) for any purpose whatsoever,
including without limitation commercial, advertising or promotional
purposes (the "Waiver"). Affirmer makes the Waiver for the benefit of each
member of the public at large and to the detriment of Affirmer's heirs and
successors, fully intending that such Waiver shall not be subject to
revocation, rescission, cancellation, termination, or any other legal or
equitable action to disrupt the quiet enjoyment of the Work by the public
as contemplated by Affirmer's express Statement of Purpose.

3. Public License Fallback. Should any part of the Waiver for any reason
be judged legally invalid or ineffective under applicable law, then the
Waiver shall be preserved to the maximum extent permitted taking into
account Affirmer's express Statement of Purpose. In addition, to the
extent the Waiver is so judged Affirmer hereby grants to each affected
person a royalty-free, non transferable, non sublicensable, non exclusive,
irrevocable and unconditional license to exercise Affirmer's Copyright and
Related Rights in the Work (i) in all territories worldwide, (ii) for the
maximum duration provided by applicable law or treaty (including future
time extensions), (iii) in any current or future medium and for any number
of copies, and (iv) for any purpose whatsoever, including without
limitation commercial, advertising or promotional purposes (the
"License"). The License shall be deemed effective as of the date CC0 was
applied by Affirmer to the Work. Should any part of the License for any
reason be judged legally invalid or ineffective under applicable law, such
partial invalidity or ineffectiveness shall not invalidate the remainder
of the License, and in such case Affirmer hereby affirms that he or she
will not (i) exercise any of his or her remaining Copyright and Related
Rights in the Work or (ii) assert any associated claims and causes of
action with respect to the Work, in either case contrary to Affirmer's
express Statement of Purpose.
```

(continues on next page)

(continued from previous page)

```
4. Limitations and Disclaimers.

 a. No trademark or patent rights held by Affirmer are waived, abandoned,
    surrendered, licensed or otherwise affected by this document.
 b. Affirmer offers the Work as-is and makes no representations or
    warranties of any kind concerning the Work, express, implied,
    statutory or otherwise, including without limitation warranties of
    title, merchantability, fitness for a particular purpose, non
    infringement, or the absence of latent or other defects, accuracy, or
    the present or absence of errors, whether or not discoverable, all to
    the greatest extent permissible under applicable law.
 c. Affirmer disclaims responsibility for clearing rights of other persons
    that may apply to the Work or any use thereof, including without
    limitation any person's Copyright and Related Rights in the Work.
    Further, Affirmer disclaims responsibility for obtaining any necessary
    consents, permissions or other rights required for any use of the
    Work.
 d. Affirmer understands and acknowledges that Creative Commons is not a
    party to this document and has no duty or obligation with respect to
    this CC0 or use of the Work.

--------------------------------------------------------------------------------

                                 Apache License
                           Version 2.0, January 2004
                        http://www.apache.org/licenses/

   TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

   1. Definitions.

      "License" shall mean the terms and conditions for use, reproduction,
      and distribution as defined by Sections 1 through 9 of this document.

      "Licensor" shall mean the copyright owner or entity authorized by
      the copyright owner that is granting the License.

      "Legal Entity" shall mean the union of the acting entity and all
      other entities that control, are controlled by, or are under common
      control with that entity. For the purposes of this definition,
      "control" means (i) the power, direct or indirect, to cause the
      direction or management of such entity, whether by contract or
      otherwise, or (ii) ownership of fifty percent (50%) or more of the
      outstanding shares, or (iii) beneficial ownership of such entity.

      "You" (or "Your") shall mean an individual or Legal Entity
      exercising permissions granted by this License.

      "Source" form shall mean the preferred form for making modifications,
      including but not limited to software source code, documentation
      source, and configuration files.
```

(continues on next page)

(continued from previous page)

```
    "Object" form shall mean any form resulting from mechanical
    transformation or translation of a Source form, including but
    not limited to compiled object code, generated documentation,
    and conversions to other media types.

    "Work" shall mean the work of authorship, whether in Source or
    Object form, made available under the License, as indicated by a
    copyright notice that is included in or attached to the work
    (an example is provided in the Appendix below).

    "Derivative Works" shall mean any work, whether in Source or Object
    form, that is based on (or derived from) the Work and for which the
    editorial revisions, annotations, elaborations, or other modifications
    represent, as a whole, an original work of authorship. For the purposes
    of this License, Derivative Works shall not include works that remain
    separable from, or merely link (or bind by name) to the interfaces of,
    the Work and Derivative Works thereof.

    "Contribution" shall mean any work of authorship, including
    the original version of the Work and any modifications or additions
    to that Work or Derivative Works thereof, that is intentionally
    submitted to Licensor for inclusion in the Work by the copyright owner
    or by an individual or Legal Entity authorized to submit on behalf of
    the copyright owner. For the purposes of this definition, "submitted"
    means any form of electronic, verbal, or written communication sent
    to the Licensor or its representatives, including but not limited to
    communication on electronic mailing lists, source code control systems,
    and issue tracking systems that are managed by, or on behalf of, the
    Licensor for the purpose of discussing and improving the Work, but
    excluding communication that is conspicuously marked or otherwise
    designated in writing by the copyright owner as "Not a Contribution."

    "Contributor" shall mean Licensor and any individual or Legal Entity
    on behalf of whom a Contribution has been received by Licensor and
    subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of
    this License, each Contributor hereby grants to You a perpetual,
    worldwide, non-exclusive, no-charge, royalty-free, irrevocable
    copyright license to reproduce, prepare Derivative Works of,
    publicly display, publicly perform, sublicense, and distribute the
    Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of
    this License, each Contributor hereby grants to You a perpetual,
    worldwide, non-exclusive, no-charge, royalty-free, irrevocable
    (except as stated in this section) patent license to make, have made,
    use, offer to sell, sell, import, and otherwise transfer the Work,
    where such license applies only to those patent claims licensable
    by such Contributor that are necessarily infringed by their
    Contribution(s) alone or by combination of their Contribution(s)
```

(continues on next page)

(continued from previous page)

```
   with the Work to which such Contribution(s) was submitted. If You
   institute patent litigation against any entity (including a
   cross-claim or counterclaim in a lawsuit) alleging that the Work
   or a Contribution incorporated within the Work constitutes direct
   or contributory patent infringement, then any patent licenses
   granted to You under this License for that Work shall terminate
   as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the
   Work or Derivative Works thereof in any medium, with or without
   modifications, and in Source or Object form, provided that You
   meet the following conditions:

   (a) You must give any other recipients of the Work or
       Derivative Works a copy of this License; and

   (b) You must cause any modified files to carry prominent notices
       stating that You changed the files; and

   (c) You must retain, in the Source form of any Derivative Works
       that You distribute, all copyright, patent, trademark, and
       attribution notices from the Source form of the Work,
       excluding those notices that do not pertain to any part of
       the Derivative Works; and

   (d) If the Work includes a "NOTICE" text file as part of its
       distribution, then any Derivative Works that You distribute must
       include a readable copy of the attribution notices contained
       within such NOTICE file, excluding those notices that do not
       pertain to any part of the Derivative Works, in at least one
       of the following places: within a NOTICE text file distributed
       as part of the Derivative Works; within the Source form or
       documentation, if provided along with the Derivative Works; or,
       within a display generated by the Derivative Works, if and
       wherever such third-party notices normally appear. The contents
       of the NOTICE file are for informational purposes only and
       do not modify the License. You may add Your own attribution
       notices within Derivative Works that You distribute, alongside
       or as an addendum to the NOTICE text from the Work, provided
       that such additional attribution notices cannot be construed
       as modifying the License.

   You may add Your own copyright statement to Your modifications and
   may provide additional or different license terms and conditions
   for use, reproduction, or distribution of Your modifications, or
   for any such Derivative Works as a whole, provided Your use,
   reproduction, and distribution of the Work otherwise complies with
   the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise,
   any Contribution intentionally submitted for inclusion in the Work
   by You to the Licensor shall be under the terms and conditions of
```

(continues on next page)

(continued from previous page)

```
   this License, without any additional terms or conditions.
   Notwithstanding the above, nothing herein shall supersede or modify
   the terms of any separate license agreement you may have executed
   with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade
   names, trademarks, service marks, or product names of the Licensor,
   except as required for reasonable and customary use in describing the
   origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or
   agreed to in writing, Licensor provides the Work (and each
   Contributor provides its Contributions) on an "AS IS" BASIS,
   WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
   implied, including, without limitation, any warranties or conditions
   of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A
   PARTICULAR PURPOSE. You are solely responsible for determining the
   appropriateness of using or redistributing the Work and assume any
   risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory,
   whether in tort (including negligence), contract, or otherwise,
   unless required by applicable law (such as deliberate and grossly
   negligent acts) or agreed to in writing, shall any Contributor be
   liable to You for damages, including any direct, indirect, special,
   incidental, or consequential damages of any character arising as a
   result of this License or out of the use or inability to use the
   Work (including but not limited to damages for loss of goodwill,
   work stoppage, computer failure or malfunction, or any and all
   other commercial damages or losses), even if such Contributor
   has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing
   the Work or Derivative Works thereof, You may choose to offer,
   and charge a fee for, acceptance of support, warranty, indemnity,
   or other liability obligations and/or rights consistent with this
   License. However, in accepting such obligations, You may act only
   on Your own behalf and on Your sole responsibility, not on behalf
   of any other Contributor, and only if You agree to indemnify,
   defend, and hold each Contributor harmless for any liability
   incurred by, or claims asserted against, such Contributor by reason
   of your accepting any such warranty or additional liability.
```

40.37 BSD3CLAUSE License Text for libevent

The following license text applies to libevent:

Listing 37: Download: BSD3CLAUSE for libevent

```
Libevent is available for use under the following license, commonly known
as the 3-clause (or "modified") BSD license:

==============================
Copyright (c) 2000-2007 Niels Provos <provos@citi.umich.edu>
Copyright (c) 2007-2012 Niels Provos and Nick Mathewson

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products
   derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR
IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
==============================

Portions of Libevent are based on works by others, also made available by
them under the three-clause BSD license above.  The copyright notices are
available in the corresponding source files; the license is as above.  Here's
a list:

log.c:
   Copyright (c) 2000 Dug Song <dugsong@monkey.org>
   Copyright (c) 1993 The Regents of the University of California.

strlcpy.c:
   Copyright (c) 1998 Todd C. Miller <Todd.Miller@courtesan.com>

win32select.c:
   Copyright (c) 2003 Michael A. Davis <mike@datanerds.net>

evport.c:
   Copyright (c) 2007 Sun Microsystems
```

(continues on next page)

(continued from previous page)

```
ht-internal.h:
   Copyright (c) 2002 Christopher Clark

minheap-internal.h:
   Copyright (c) 2006 Maxim Yegorushkin <maxim.yegorushkin@gmail.com>

==============================

The arc4module is available under the following, sometimes called the
"OpenBSD" license:

   Copyright (c) 1996, David Mazieres <dm@uun.org>
   Copyright (c) 2008, Damien Miller <djm@openbsd.org>

   Permission to use, copy, modify, and distribute this software for any
   purpose with or without fee is hereby granted, provided that the above
   copyright notice and this permission notice appear in all copies.

   THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES
   WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
   MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR
   ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
   WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
   ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF
   OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

==============================

The Windows timer code is based on code from libutp, which is
distributed under this license, sometimes called the "MIT" license.


Copyright (c) 2010 BitTorrent, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN
THE SOFTWARE.
```

40.38 MIT License Text for libffi

The following license text applies to libffi:

Listing 38: Download: MIT for libffi

```
libffi - Copyright (c) 1996-2019  Anthony Green, Red Hat, Inc and others.
See source files for details.

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
``Software''), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject to
the following conditions:

The above copyright notice and this permission notice shall be
included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED ``AS IS'', WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.
IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY
CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT,
TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE
SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

40.39 GPLv2 License Text for libgpg-error

The following license text applies to libgpg-error:

Listing 39: Download: GPLv2 for libgpg-error

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.
     51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
```

(continues on next page)

(continued from previous page)

```
the GNU Library General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and

modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
```

(continues on next page)

(continued from previous page)

```
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on

    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
```

(continues on next page)

(continued from previous page)

```
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not

compelled to copy the source along with the object code.
```

(continues on next page)

(continued from previous page)

```
  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
```

(continues on next page)

(continued from previous page)

```
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to

be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
```

(continues on next page)

(continued from previous page)

```
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.


                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA


Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year  name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.
```

(continues on next page)

(continued from previous page)

```
You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Library General
Public License instead of this License.
```

40.40 LGPL21 License Text for libgpg-error

The following license text applies to libgpg-error:

Listing 40: Download: LGPL21 for libgpg-error

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
        51 Franklin St, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations
below.
```

(continues on next page)

(continued from previous page)

```
  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be

introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
```

(continues on next page)

(continued from previous page)

```
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it
becomes a de-facto standard.  To achieve this, non-free programs must
be allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must

be combined with the library in order to run.

                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
```

(continues on next page)

(continued from previous page)

```
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control
compilation and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
```

(continues on next page)

(continued from previous page)

```
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
```

(continues on next page)

(continued from previous page)

```
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
```

(continues on next page)

(continued from previous page)

```
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at least
    three years, to give the same user the materials specified in
    Subsection 6a, above, for a charge no more than the cost of
    performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you

distribute.
```

(continues on next page)

(continued from previous page)

```
  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
```

(continues on next page)

(continued from previous page)

```
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply, and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License
may add an explicit geographical distribution limitation excluding those
countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.
```

(continues on next page)

(continued from previous page)

```
                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.


                     END OF TERMS AND CONDITIONS

           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms
of the ordinary General Public License).

  To apply these terms, attach the following notices to the library.
It is safest to attach them to the start of each source file to most
effectively convey the exclusion of warranty; and each file should
have at least the "copyright" line and a pointer to where the full
notice is found.


    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
```

(continues on next page)

(continued from previous page)

```
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA  02110-1301  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or
your school, if any, to sign a "copyright disclaimer" for the library,
if necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James
  Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That's all there is to it!
```

40.41 GPLv3 License Text for libiconv

The following license text applies to libiconv:

Listing 41: Download: GPLv3 for libiconv

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 3, 29 June 2007

 Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The GNU General Public License is a free, copyleft license for
software and other kinds of works.

  The licenses for most software and other practical works are designed
to take away your freedom to share and change the works.  By contrast,
the GNU General Public License is intended to guarantee your freedom to
share and change all versions of a program--to make sure it remains free
software for all its users.  We, the Free Software Foundation, use the
GNU General Public License for most of our software; it applies also to
any other work released this way by its authors.  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
```

(continues on next page)

(continued from previous page)

```
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
them if you wish), that you receive source code or can get it if you
want it, that you can change the software or use pieces of it in new
free programs, and that you know you can do these things.

  To protect your rights, we need to prevent others from denying you
these rights or asking you to surrender the rights.  Therefore, you have
certain responsibilities if you distribute copies of the software, or if
you modify it: responsibilities to respect the freedom of others.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must pass on to the recipients the same
freedoms that you received.  You must make sure that they, too, receive
or can get the source code.  And you must show them these terms so they
know their rights.

  Developers that use the GNU GPL protect your rights with two steps:
(1) assert copyright on the software, and (2) offer you this License
giving you legal permission to copy, distribute and/or modify it.

  For the developers' and authors' protection, the GPL clearly explains
that there is no warranty for this free software.  For both users' and
authors' sake, the GPL requires that modified versions be marked as
changed, so that their problems will not be attributed erroneously to
authors of previous versions.

  Some devices are designed to deny users access to install or run
modified versions of the software inside them, although the manufacturer
can do so.  This is fundamentally incompatible with the aim of
protecting users' freedom to change the software.  The systematic
pattern of such abuse occurs in the area of products for individuals to
use, which is precisely where it is most unacceptable.  Therefore, we
have designed this version of the GPL to prohibit the practice for those
products.  If such problems arise substantially in other domains, we
stand ready to extend this provision to those domains in future versions
of the GPL, as needed to protect the freedom of users.

  Finally, every program is threatened constantly by software patents.
States should not allow patents to restrict development and use of
software on general-purpose computers, but in those that do, we wish to
avoid the special danger that patents applied to a free program could
make it effectively proprietary.  To prevent this, the GPL assures that
patents cannot be used to render the program non-free.

  The precise terms and conditions for copying, distribution and
modification follow.

                       TERMS AND CONDITIONS

  0. Definitions.
```

(continues on next page)

(continued from previous page)

```
  "This License" refers to version 3 of the GNU General Public License.

  "Copyright" also means copyright-like laws that apply to other kinds of
works, such as semiconductor masks.

  "The Program" refers to any copyrightable work licensed under this
License.  Each licensee is addressed as "you".  "Licensees" and
"recipients" may be individuals or organizations.

  To "modify" a work means to copy from or adapt all or part of the work
in a fashion requiring copyright permission, other than the making of an
exact copy.  The resulting work is called a "modified version" of the
earlier work or a work "based on" the earlier work.

  A "covered work" means either the unmodified Program or a work based
on the Program.

  To "propagate" a work means to do anything with it that, without
permission, would make you directly or secondarily liable for
infringement under applicable copyright law, except executing it on a
computer or modifying a private copy.  Propagation includes copying,
distribution (with or without modification), making available to the
public, and in some countries other activities as well.

  To "convey" a work means any kind of propagation that enables other
parties to make or receive copies.  Mere interaction with a user through
a computer network, with no transfer of a copy, is not conveying.

  An interactive user interface displays "Appropriate Legal Notices"
to the extent that it includes a convenient and prominently visible
feature that (1) displays an appropriate copyright notice, and (2)
tells the user that there is no warranty for the work (except to the
extent that warranties are provided), that licensees may convey the
work under this License, and how to view a copy of this License.  If
the interface presents a list of user commands or options, such as a
menu, a prominent item in the list meets this criterion.

  1. Source Code.

  The "source code" for a work means the preferred form of the work
for making modifications to it.  "Object code" means any non-source
form of a work.

  A "Standard Interface" means an interface that either is an official
standard defined by a recognized standards body, or, in the case of
interfaces specified for a particular programming language, one that
is widely used among developers working in that language.

  The "System Libraries" of an executable work include anything, other
than the work as a whole, that (a) is included in the normal form of
packaging a Major Component, but which is not part of that Major
Component, and (b) serves only to enable use of the work with that
```

(continues on next page)

(continued from previous page)

```
Major Component, or to implement a Standard Interface for which an
implementation is available to the public in source code form.  A
"Major Component", in this context, means a major essential component
(kernel, window system, and so on) of the specific operating system
(if any) on which the executable work runs, or a compiler used to
produce the work, or an object code interpreter used to run it.

  The "Corresponding Source" for a work in object code form means all
the source code needed to generate, install, and (for an executable
work) run the object code and to modify the work, including scripts to
control those activities.  However, it does not include the work's
System Libraries, or general-purpose tools or generally available free
programs which are used unmodified in performing those activities but
which are not part of the work.  For example, Corresponding Source
includes interface definition files associated with source files for
the work, and the source code for shared libraries and dynamically
linked subprograms that the work is specifically designed to require,
such as by intimate data communication or control flow between those
subprograms and other parts of the work.

  The Corresponding Source need not include anything that users
can regenerate automatically from other parts of the Corresponding
Source.

  The Corresponding Source for a work in source code form is that
same work.

  2. Basic Permissions.

  All rights granted under this License are granted for the term of
copyright on the Program, and are irrevocable provided the stated
conditions are met.  This License explicitly affirms your unlimited
permission to run the unmodified Program.  The output from running a
covered work is covered by this License only if the output, given its
content, constitutes a covered work.  This License acknowledges your
rights of fair use or other equivalent, as provided by copyright law.

  You may make, run and propagate covered works that you do not
convey, without conditions so long as your license otherwise remains
in force.  You may convey covered works to others for the sole purpose
of having them make modifications exclusively for you, or provide you
with facilities for running those works, provided that you comply with
the terms of this License in conveying all material for which you do
not control copyright.  Those thus making or running the covered works
for you must do so exclusively on your behalf, under your direction
and control, on terms that prohibit them from making any copies of
your copyrighted material outside their relationship with you.

  Conveying under any other circumstances is permitted solely under
the conditions stated below.  Sublicensing is not allowed; section 10
makes it unnecessary.
```

(continues on next page)

(continued from previous page)

```
  3. Protecting Users' Legal Rights From Anti-Circumvention Law.

  No covered work shall be deemed part of an effective technological
measure under any applicable law fulfilling obligations under article
11 of the WIPO copyright treaty adopted on 20 December 1996, or
similar laws prohibiting or restricting circumvention of such
measures.

  When you convey a covered work, you waive any legal power to forbid
circumvention of technological measures to the extent such circumvention
is effected by exercising rights under this License with respect to
the covered work, and you disclaim any intention to limit operation or
modification of the work as a means of enforcing, against the work's
users, your or third parties' legal rights to forbid circumvention of
technological measures.

  4. Conveying Verbatim Copies.

  You may convey verbatim copies of the Program's source code as you
receive it, in any medium, provided that you conspicuously and
appropriately publish on each copy an appropriate copyright notice;
keep intact all notices stating that this License and any
non-permissive terms added in accord with section 7 apply to the code;
keep intact all notices of the absence of any warranty; and give all
recipients a copy of this License along with the Program.

  You may charge any price or no price for each copy that you convey,
and you may offer support or warranty protection for a fee.

  5. Conveying Modified Source Versions.

  You may convey a work based on the Program, or the modifications to
produce it from the Program, in the form of source code under the
terms of section 4, provided that you also meet all of these conditions:

    a) The work must carry prominent notices stating that you modified
    it, and giving a relevant date.

    b) The work must carry prominent notices stating that it is
    released under this License and any conditions added under section
    7.  This requirement modifies the requirement in section 4 to
    "keep intact all notices".

    c) You must license the entire work, as a whole, under this
    License to anyone who comes into possession of a copy.  This
    License will therefore apply, along with any applicable section 7
    additional terms, to the whole of the work, and all its parts,
    regardless of how they are packaged.  This License gives no
    permission to license the work in any other way, but it does not
    invalidate such permission if you have separately received it.

    d) If the work has interactive user interfaces, each must display
```

(continues on next page)

(continued from previous page)

```
    Appropriate Legal Notices; however, if the Program has interactive
    interfaces that do not display Appropriate Legal Notices, your
    work need not make them do so.

  A compilation of a covered work with other separate and independent
works, which are not by their nature extensions of the covered work,
and which are not combined with it such as to form a larger program,
in or on a volume of a storage or distribution medium, is called an
"aggregate" if the compilation and its resulting copyright are not
used to limit the access or legal rights of the compilation's users
beyond what the individual works permit.  Inclusion of a covered work
in an aggregate does not cause this License to apply to the other
parts of the aggregate.

  6. Conveying Non-Source Forms.

  You may convey a covered work in object code form under the terms
of sections 4 and 5, provided that you also convey the
machine-readable Corresponding Source under the terms of this License,
in one of these ways:

    a) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by the
    Corresponding Source fixed on a durable physical medium
    customarily used for software interchange.

    b) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by a
    written offer, valid for at least three years and valid for as
    long as you offer spare parts or customer support for that product
    model, to give anyone who possesses the object code either (1) a
    copy of the Corresponding Source for all the software in the
    product that is covered by this License, on a durable physical
    medium customarily used for software interchange, for a price no
    more than your reasonable cost of physically performing this
    conveying of source, or (2) access to copy the
    Corresponding Source from a network server at no charge.

    c) Convey individual copies of the object code with a copy of the
    written offer to provide the Corresponding Source.  This
    alternative is allowed only occasionally and noncommercially, and
    only if you received the object code with such an offer, in accord
    with subsection 6b.

    d) Convey the object code by offering access from a designated
    place (gratis or for a charge), and offer equivalent access to the
    Corresponding Source in the same way through the same place at no
    further charge.  You need not require recipients to copy the
    Corresponding Source along with the object code.  If the place to
    copy the object code is a network server, the Corresponding Source
    may be on a different server (operated by you or a third party)
    that supports equivalent copying facilities, provided you maintain
```

(continues on next page)

(continued from previous page)

```
    clear directions next to the object code saying where to find the
    Corresponding Source.  Regardless of what server hosts the
    Corresponding Source, you remain obligated to ensure that it is
    available for as long as needed to satisfy these requirements.

    e) Convey the object code using peer-to-peer transmission, provided
    you inform other peers where the object code and Corresponding
    Source of the work are being offered to the general public at no
    charge under subsection 6d.

  A separable portion of the object code, whose source code is excluded
from the Corresponding Source as a System Library, need not be
included in conveying the object code work.

  A "User Product" is either (1) a "consumer product", which means any
tangible personal property which is normally used for personal, family,
or household purposes, or (2) anything designed or sold for incorporation
into a dwelling.  In determining whether a product is a consumer product,
doubtful cases shall be resolved in favor of coverage.  For a particular
product received by a particular user, "normally used" refers to a
typical or common use of that class of product, regardless of the status
of the particular user or of the way in which the particular user
actually uses, or expects or is expected to use, the product.  A product
is a consumer product regardless of whether the product has substantial
commercial, industrial or non-consumer uses, unless such uses represent
the only significant mode of use of the product.

  "Installation Information" for a User Product means any methods,
procedures, authorization keys, or other information required to install
and execute modified versions of a covered work in that User Product from
a modified version of its Corresponding Source.  The information must
suffice to ensure that the continued functioning of the modified object
code is in no case prevented or interfered with solely because
modification has been made.

  If you convey an object code work under this section in, or with, or
specifically for use in, a User Product, and the conveying occurs as
part of a transaction in which the right of possession and use of the
User Product is transferred to the recipient in perpetuity or for a
fixed term (regardless of how the transaction is characterized), the
Corresponding Source conveyed under this section must be accompanied
by the Installation Information.  But this requirement does not apply
if neither you nor any third party retains the ability to install
modified object code on the User Product (for example, the work has
been installed in ROM).

  The requirement to provide Installation Information does not include a
requirement to continue to provide support service, warranty, or updates
for a work that has been modified or installed by the recipient, or for
the User Product in which it has been modified or installed.  Access to a
network may be denied when the modification itself materially and
adversely affects the operation of the network or violates the rules and
```

(continues on next page)

(continued from previous page)

```
protocols for communication across the network.

  Corresponding Source conveyed, and Installation Information provided,
in accord with this section must be in a format that is publicly
documented (and with an implementation available to the public in
source code form), and must require no special password or key for
unpacking, reading or copying.

  7. Additional Terms.

  "Additional permissions" are terms that supplement the terms of this
License by making exceptions from one or more of its conditions.
Additional permissions that are applicable to the entire Program shall
be treated as though they were included in this License, to the extent
that they are valid under applicable law.  If additional permissions
apply only to part of the Program, that part may be used separately
under those permissions, but the entire Program remains governed by
this License without regard to the additional permissions.

  When you convey a copy of a covered work, you may at your option
remove any additional permissions from that copy, or from any part of
it.  (Additional permissions may be written to require their own
removal in certain cases when you modify the work.)  You may place
additional permissions on material, added by you to a covered work,
for which you have or can give appropriate copyright permission.

  Notwithstanding any other provision of this License, for material you
add to a covered work, you may (if authorized by the copyright holders of
that material) supplement the terms of this License with terms:

    a) Disclaiming warranty or limiting liability differently from the
    terms of sections 15 and 16 of this License; or

    b) Requiring preservation of specified reasonable legal notices or
    author attributions in that material or in the Appropriate Legal
    Notices displayed by works containing it; or

    c) Prohibiting misrepresentation of the origin of that material, or
    requiring that modified versions of such material be marked in
    reasonable ways as different from the original version; or

    d) Limiting the use for publicity purposes of names of licensors or
    authors of the material; or

    e) Declining to grant rights under trademark law for use of some
    trade names, trademarks, or service marks; or

    f) Requiring indemnification of licensors and authors of that
    material by anyone who conveys the material (or modified versions of
    it) with contractual assumptions of liability to the recipient, for
    any liability that these contractual assumptions directly impose on
    those licensors and authors.
```

(continues on next page)

(continued from previous page)

```
  All other non-permissive additional terms are considered "further
restrictions" within the meaning of section 10.  If the Program as you
received it, or any part of it, contains a notice stating that it is
governed by this License along with a term that is a further
restriction, you may remove that term.  If a license document contains
a further restriction but permits relicensing or conveying under this
License, you may add to a covered work material governed by the terms
of that license document, provided that the further restriction does
not survive such relicensing or conveying.

  If you add terms to a covered work in accord with this section, you
must place, in the relevant source files, a statement of the
additional terms that apply to those files, or a notice indicating
where to find the applicable terms.

  Additional terms, permissive or non-permissive, may be stated in the
form of a separately written license, or stated as exceptions;
the above requirements apply either way.

  8. Termination.

  You may not propagate or modify a covered work except as expressly
provided under this License.  Any attempt otherwise to propagate or
modify it is void, and will automatically terminate your rights under
this License (including any patent licenses granted under the third
paragraph of section 11).

  However, if you cease all violation of this License, then your
license from a particular copyright holder is reinstated (a)
provisionally, unless and until the copyright holder explicitly and
finally terminates your license, and (b) permanently, if the copyright
holder fails to notify you of the violation by some reasonable means
prior to 60 days after the cessation.

  Moreover, your license from a particular copyright holder is
reinstated permanently if the copyright holder notifies you of the
violation by some reasonable means, this is the first time you have
received notice of violation of this License (for any work) from that
copyright holder, and you cure the violation prior to 30 days after
your receipt of the notice.

  Termination of your rights under this section does not terminate the
licenses of parties who have received copies or rights from you under
this License.  If your rights have been terminated and not permanently
reinstated, you do not qualify to receive new licenses for the same
material under section 10.

  9. Acceptance Not Required for Having Copies.

  You are not required to accept this License in order to receive or
run a copy of the Program.  Ancillary propagation of a covered work
```

(continues on next page)

(continued from previous page)

```
occurring solely as a consequence of using peer-to-peer transmission
to receive a copy likewise does not require acceptance.  However,
nothing other than this License grants you permission to propagate or
modify any covered work.  These actions infringe copyright if you do
not accept this License.  Therefore, by modifying or propagating a
covered work, you indicate your acceptance of this License to do so.

  10. Automatic Licensing of Downstream Recipients.

  Each time you convey a covered work, the recipient automatically
receives a license from the original licensors, to run, modify and
propagate that work, subject to this License.  You are not responsible
for enforcing compliance by third parties with this License.

  An "entity transaction" is a transaction transferring control of an
organization, or substantially all assets of one, or subdividing an
organization, or merging organizations.  If propagation of a covered
work results from an entity transaction, each party to that
transaction who receives a copy of the work also receives whatever
licenses to the work the party's predecessor in interest had or could
give under the previous paragraph, plus a right to possession of the
Corresponding Source of the work from the predecessor in interest, if
the predecessor has it or can get it with reasonable efforts.

  You may not impose any further restrictions on the exercise of the
rights granted or affirmed under this License.  For example, you may
not impose a license fee, royalty, or other charge for exercise of
rights granted under this License, and you may not initiate litigation
(including a cross-claim or counterclaim in a lawsuit) alleging that
any patent claim is infringed by making, using, selling, offering for
sale, or importing the Program or any portion of it.

  11. Patents.

  A "contributor" is a copyright holder who authorizes use under this
License of the Program or a work on which the Program is based.  The
work thus licensed is called the contributor's "contributor version".

  A contributor's "essential patent claims" are all patent claims
owned or controlled by the contributor, whether already acquired or
hereafter acquired, that would be infringed by some manner, permitted
by this License, of making, using, or selling its contributor version,
but do not include claims that would be infringed only as a
consequence of further modification of the contributor version.  For
purposes of this definition, "control" includes the right to grant
patent sublicenses in a manner consistent with the requirements of
this License.

  Each contributor grants you a non-exclusive, worldwide, royalty-free
patent license under the contributor's essential patent claims, to
make, use, sell, offer for sale, import and otherwise run, modify and
propagate the contents of its contributor version.
```

(continues on next page)

(continued from previous page)

```
  In the following three paragraphs, a "patent license" is any express
agreement or commitment, however denominated, not to enforce a patent
(such as an express permission to practice a patent or covenant not to
sue for patent infringement).  To "grant" such a patent license to a
party means to make such an agreement or commitment not to enforce a
patent against the party.

  If you convey a covered work, knowingly relying on a patent license,
and the Corresponding Source of the work is not available for anyone
to copy, free of charge and under the terms of this License, through a
publicly available network server or other readily accessible means,
then you must either (1) cause the Corresponding Source to be so
available, or (2) arrange to deprive yourself of the benefit of the
patent license for this particular work, or (3) arrange, in a manner
consistent with the requirements of this License, to extend the patent
license to downstream recipients.  "Knowingly relying" means you have
actual knowledge that, but for the patent license, your conveying the
covered work in a country, or your recipient's use of the covered work
in a country, would infringe one or more identifiable patents in that
country that you have reason to believe are valid.

  If, pursuant to or in connection with a single transaction or
arrangement, you convey, or propagate by procuring conveyance of, a
covered work, and grant a patent license to some of the parties
receiving the covered work authorizing them to use, propagate, modify
or convey a specific copy of the covered work, then the patent license
you grant is automatically extended to all recipients of the covered
work and works based on it.

  A patent license is "discriminatory" if it does not include within
the scope of its coverage, prohibits the exercise of, or is
conditioned on the non-exercise of one or more of the rights that are
specifically granted under this License.  You may not convey a covered
work if you are a party to an arrangement with a third party that is
in the business of distributing software, under which you make payment
to the third party based on the extent of your activity of conveying
the work, and under which the third party grants, to any of the
parties who would receive the covered work from you, a discriminatory
patent license (a) in connection with copies of the covered work
conveyed by you (or copies made from those copies), or (b) primarily
for and in connection with specific products or compilations that
contain the covered work, unless you entered into that arrangement,
or that patent license was granted, prior to 28 March 2007.

  Nothing in this License shall be construed as excluding or limiting
any implied license or other defenses to infringement that may
otherwise be available to you under applicable patent law.

  12. No Surrender of Others' Freedom.

  If conditions are imposed on you (whether by court order, agreement or
```

(continues on next page)

(continued from previous page)

```
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot convey a
covered work so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you may
not convey it at all.  For example, if you agree to terms that obligate you
to collect a royalty for further conveying from those to whom you convey
the Program, the only way you could satisfy both those terms and this
License would be to refrain entirely from conveying the Program.

  13. Use with the GNU Affero General Public License.

  Notwithstanding any other provision of this License, you have
permission to link or combine any covered work with a work licensed
under version 3 of the GNU Affero General Public License into a single
combined work, and to convey the resulting work.  The terms of this
License will continue to apply to the part which is the covered work,
but the special requirements of the GNU Affero General Public License,
section 13, concerning interaction through a network will apply to the
combination as such.

  14. Revised Versions of this License.

  The Free Software Foundation may publish revised and/or new versions of
the GNU General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

  Each version is given a distinguishing version number.  If the
Program specifies that a certain numbered version of the GNU General
Public License "or any later version" applies to it, you have the
option of following the terms and conditions either of that numbered
version or of any later version published by the Free Software
Foundation.  If the Program does not specify a version number of the
GNU General Public License, you may choose any version ever published
by the Free Software Foundation.

  If the Program specifies that a proxy can decide which future
versions of the GNU General Public License can be used, that proxy's
public statement of acceptance of a version permanently authorizes you
to choose that version for the Program.

  Later license versions may give you additional or different
permissions.  However, no additional obligations are imposed on any
author or copyright holder as a result of your choosing to follow a
later version.

  15. Disclaimer of Warranty.

  THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY
APPLICABLE LAW.  EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT
HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY
OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO,
```

(continues on next page)

(continued from previous page)

```
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM
IS WITH YOU.  SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF
ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. Limitation of Liability.

  IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS
THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY
GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE
USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF
DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD
PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS),
EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF
SUCH DAMAGES.

  17. Interpretation of Sections 15 and 16.

  If the disclaimer of warranty and limitation of liability provided
above cannot be given local legal effect according to their terms,
reviewing courts shall apply local law that most closely approximates
an absolute waiver of all civil liability in connection with the
Program, unless a warranty or assumption of liability accompanies a
copy of the Program in return for a fee.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
state the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software: you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation, either version 3 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.
```

(continues on next page)

(continued from previous page)

```
    You should have received a copy of the GNU General Public License
    along with this program.  If not, see <http://www.gnu.org/licenses/>.

Also add information on how to contact you by electronic and paper mail.

  If the program does terminal interaction, make it output a short
notice like this when it starts in an interactive mode:

    <program>  Copyright (C) <year>  <name of author>
    This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, your program's commands
might be different; for a GUI interface, you would use an "about box".

  You should also get your employer (if you work as a programmer) or school,
if any, to sign a "copyright disclaimer" for the program, if necessary.
For more information on this, and how to apply and follow the GNU GPL, see
<http://www.gnu.org/licenses/>.

  The GNU General Public License does not permit incorporating your program
into proprietary programs.  If your program is a subroutine library, you
may consider it more useful to permit linking proprietary applications with
the library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.  But first, please read
<http://www.gnu.org/philosophy/why-not-lgpl.html>.
```

40.42 GPLv3 License Text for libidn2

The following license text applies to libidn2:

Listing 42: Download: GPLv3 for libidn2

```
Libidn2 COPYING -- Licensing information.                    -*- outline -*-
Copyright (C) 2011-2016 Simon Josefsson
See the end for copying conditions.

The source code for the C library (libidn2.a or libidn.so) are
licensed under the terms of either the GNU General Public License
version 2.0 or later (see the file COPYINGv2) or the GNU Lesser
General Public License version 3.0 or later (see the file
COPYING.LESSERv3), or both in parallel as here.

The command line tool, self tests, examples, and other auxiliary
files, are licensed under the GNU General Public License version 3.0
or later.

The license of the Unicode character data files (which are parsed into
```

(continues on next page)

(continued from previous page)

```
static storage in the library) are documented in COPYING.unicode.

Other files are licensed as indicated in each file.

There may be exceptions to these general rules, see each file for
precise information.

----------------------------------------------------------------------
This file is free software: you can redistribute it and/or modify it
under the terms of the GNU General Public License as published by the
Free Software Foundation, either version 3 of the License, or (at your
option) any later version.

This file is distributed in the hope that it will be useful, but
WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
General Public License for more details.

You should have received a copy of the GNU General Public License
along with this file.  If not, see <http://www.gnu.org/licenses/>.
```

40.43 LGPL21 License Text for libltdl

The following license text applies to libltdl:

Listing 43: Download: LGPL21 for libltdl

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
```

(continues on next page)

(continued from previous page)

```
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be

introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
```

(continues on next page)

(continued from previous page)

```
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it becomes
a de-facto standard.  To achieve this, non-free programs must be
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must

be combined with the library in order to run.

                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
```

(continues on next page)

(continued from previous page)

```
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
```

(continues on next page)

(continued from previous page)

```
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
```

(continues on next page)

(continued from previous page)

```
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
```

(continues on next page)

(continued from previous page)

```
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
```

(continues on next page)

(continued from previous page)

```
use both them and the Library together in an executable that you

distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
```

(continues on next page)

(continued from previous page)

```
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded.  In such case, this License incorporates the limitation as if
written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
```

(continues on next page)

(continued from previous page)

```
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.


                     END OF TERMS AND CONDITIONS

           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms of the
ordinary General Public License).

  To apply these terms, attach the following notices to the library.  It is
safest to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
```

(continues on next page)

(continued from previous page)

```
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That's all there is to it!
```

40.44 GPLv2 License Text for liblz4

The following license text applies to liblz4:

Listing 44: Download: GPLv2 for liblz4

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
```

(continues on next page)

(continued from previous page)

```
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
```

(continues on next page)

(continued from previous page)

```
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
```

(continues on next page)

(continued from previous page)

```
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.
```

(continues on next page)

(continued from previous page)

```
  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
```

(continues on next page)

(continued from previous page)

```
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs
```

(continues on next page)

(continued from previous page)

```
  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice
```

(continues on next page)

(continued from previous page)

```
This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.45 LGPL21+ License Text for libmcrypt

The following license text applies to libmcrypt:

Listing 45: Download: LGPL21+ for libmcrypt

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
     59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations
below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
```

(continues on next page)

(continued from previous page)

```
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be
introduced by others.
^L
  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
```

(continues on next page)

(continued from previous page)

```
encourage the widest possible use of a certain library, so that it
becomes a de-facto standard.  To achieve this, non-free programs must
be allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must
be combined with the library in order to run.
^L
                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control
compilation and installation of the library.
```

(continues on next page)

(continued from previous page)

```
  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
```

(continues on next page)

(continued from previous page)

```
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in
these notices.
^L
  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.
```

(continues on next page)

(continued from previous page)

```
  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,
whether or not they are linked directly with the Library itself.
^L
  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)
```

(continues on next page)

(continued from previous page)

```
    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at least
    three years, to give the same user the materials specified in
    Subsection 6a, above, for a charge no more than the cost of
    performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you
distribute.
^L
  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.
```

(continues on next page)

(continued from previous page)

```
  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with
this License.
^L
  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply, and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
```

(continues on next page)

(continued from previous page)

```
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License
may add an explicit geographical distribution limitation excluding those
countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by
the Free Software Foundation.
^L
  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
```

(continues on next page)

(continued from previous page)

```
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.

                     END OF TERMS AND CONDITIONS
^L
           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms
of the ordinary General Public License).

  To apply these terms, attach the following notices to the library.
It is safest to attach them to the start of each source file to most
effectively convey the exclusion of warranty; and each file should
have at least the "copyright" line and a pointer to where the full
notice is found.


    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or
your school, if any, to sign a "copyright disclaimer" for the library,
if necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James
  Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice
```

(continues on next page)

(continued from previous page)

```
That's all there is to it!
```

40.46 MIT License Text for libnghttp2

The following license text applies to libnghttp2:

Listing 46: Download: MIT for libnghttp2

```
The MIT License

Copyright (c) 2012, 2014, 2015, 2016 Tatsuhiro Tsujikawa
Copyright (c) 2012, 2014, 2015, 2016 nghttp2 contributors

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
"Software"), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject to
the following conditions:

The above copyright notice and this permission notice shall be
included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND
NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE
LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION
OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION
WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

40.47 BSD3CLAUSE License Text for libssh2

The following license text applies to libssh2:

Listing 47: Download: BSD3CLAUSE for libssh2

```
/* Copyright (c) 2004-2007 Sara Golemon <sarag@libssh2.org>
 * Copyright (c) 2005,2006 Mikhail Gusarov <dottedmag@dottedmag.net>
 * Copyright (c) 2006-2007 The Written Word, Inc.
 * Copyright (c) 2007 Eli Fant <elifantu@mail.ru>
 * Copyright (c) 2009-2014 Daniel Stenberg
 * Copyright (C) 2008, 2009 Simon Josefsson
 * All rights reserved.
 *
 * Redistribution and use in source and binary forms,
```

(continues on next page)

(continued from previous page)

```
 * with or without modification, are permitted provided
 * that the following conditions are met:
 *
 *   Redistributions of source code must retain the above
 *   copyright notice, this list of conditions and the
 *   following disclaimer.
 *
 *   Redistributions in binary form must reproduce the above
 *   copyright notice, this list of conditions and the following
 *   disclaimer in the documentation and/or other materials
 *   provided with the distribution.
 *
 *   Neither the name of the copyright holder nor the names
 *   of any other contributors may be used to endorse or
 *   promote products derived from this software without
 *   specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND
 * CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES,
 * INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
 * OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR
 * CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
 * SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
 * BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
 * SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
 * INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
 * WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
 * NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE
 * USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY
 * OF SUCH DAMAGE.
 */
```

40.48 BSD2CLAUSE License Text for libucl

The following license text applies to libucl:

Listing 48: Download: BSD2CLAUSE for libucl

```
Copyright (c) 2013-2014, Vsevolod Stakhov <vsevolod@highsecure.ru>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this
  list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice,
  this list of conditions and the following disclaimer in the documentation
```

(continues on next page)

(continued from previous page)

```
  and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.49 GFDL License Text for libunistring

The following license text applies to libunistring:

Listing 49: Download: GFDL for libunistring

```
\input texinfo          @c -*-texinfo-*-
@comment %**start of header
@setfilename libunistring.info
@documentencoding UTF-8
@settitle GNU libunistring
@finalout
@c Indices:
@c   am = autoconf macro  @amindex
@c   cp = concept         @cindex
@c   fn = function        @findex
@c   tp = type            @tindex
@c Unused predefined indices:
@c   ky = keystroke       @kindex
@c   pg = program         @pindex
@c   vr = variable        @vindex
@defcodeindex am
@syncodeindex am cp
@syncodeindex fn cp
@syncodeindex tp cp
@ifclear texi2html
@firstparagraphindent insert
@end ifclear
@c texi2html-1.76 does not support @arrow{}.
@ifset texi2html
@macro arrow{}
→
@end macro
@end ifset
@comment %**end of header

@include version.texi
```

(continues on next page)

(continued from previous page)

```
@c Location of the POSIX specification on the web.
@set POSIXURL http://pubs.opengroup.org/onlinepubs/9699919799

@c Macro for referencing a POSIX header.
@ifinfo
@macro posixheader{header}
@code{<\header\>}
@end macro
@end ifinfo
@ifnotinfo
@macro posixheader{header}
@uref{@value{POSIXURL}/basedefs/\header\.html,,@code{<\header\>}}
@end macro
@end ifnotinfo

@c Macro for referencing a POSIX function.
@c We don't write it as func(), see section "GNU Manuals" of the
@c GNU coding standards.
@ifinfo
@macro posixfunc{func}
@code{\func\}
@end macro
@end ifinfo
@ifnotinfo
@macro posixfunc{func}
@uref{@value{POSIXURL}/functions/\func\.html,,@code{\func\}}
@end macro
@end ifnotinfo

@c Macro for referencing a normal function.
@c We don't write it as func(), see section "GNU Manuals" of the
@c GNU coding standards.
@macro func{func}
@code{\func\}
@end macro

@c Macro for an advisory ragged line break in TeX mode.
@c Needed because there are long unbreakable pieces of text (such as URLs or
@c formulas), TeX is too shy to move them to a new line. TeX considers only
@c two choices: a line break in aligned mode (which it rejects due to aesthetic
@c reasons) and writing into the margin. What we want in many cases is a line
@c break without filling the first line. Like what @* delivers. But we want it
@c only when needed, so that it disappears when unrelated changes in the same
@c paragraph cause a line break in a nearby position. And we need it only in
@c TeX mode. info and HTML modes are fine.
@c This trick is from Karl Berry.
@iftex
@macro texnl
@hfil@penalty9000@hfilneg
@end macro
@end iftex
```

(continues on next page)

(continued from previous page)

```
@ifnottex
@macro texnl
@end macro
@end ifnottex

@ifinfo
@dircategory Software development
@direntry
* GNU libunistring: (libunistring).     Unicode string library.
@end direntry
@end ifinfo

@ifinfo
This manual is for GNU libunistring.

@ignore
@c This was: @copying but it triggers a makeinfo 4.13 bug
Copyright (C) 2001-2018 Free Software Foundation, Inc.

This manual is free documentation.  It is dually licensed under the
GNU FDL and the GNU GPL.  This means that you can redistribute this
manual under either of these two licenses, at your choice.

This manual is covered by the GNU FDL.  Permission is granted to copy,
distribute and/or modify this document under the terms of the
GNU Free Documentation License (FDL), either version 1.2 of the
License, or (at your option) any later version published by the
Free Software Foundation (FSF); with no Invariant Sections, with no
Front-Cover Text, and with no Back-Cover Texts.
A copy of the license is included in @ref{GNU FDL}.

This manual is covered by the GNU GPL.  You can redistribute it and/or
modify it under the terms of the GNU General Public License (GPL), either
version 3 of the License, or (at your option) any later version published
by the Free Software Foundation (FSF).
A copy of the license is included in @ref{GNU GPL}.
@end ignore
@end ifinfo

@titlepage
@title GNU libunistring, version @value{VERSION}
@subtitle updated @value{UPDATED}
@subtitle Edition @value{EDITION}, @value{UPDATED}
@author Bruno Haible

@ifnothtml
@page
@vskip 0pt plus 1filll
@c @insertcopying
Copyright (C) 2001-2018 Free Software Foundation, Inc.

This manual is free documentation.  It is dually licensed under the
```

(continues on next page)

(continued from previous page)

```
GNU FDL and the GNU GPL.  This means that you can redistribute this
manual under either of these two licenses, at your choice.

This manual is covered by the GNU FDL.  Permission is granted to copy,
distribute and/or modify this document under the terms of the
GNU Free Documentation License (FDL), either version 1.2 of the
License, or (at your option) any later version published by the
Free Software Foundation (FSF); with no Invariant Sections, with no
Front-Cover Text, and with no Back-Cover Texts.
A copy of the license is included in @ref{GNU FDL}.

This manual is covered by the GNU GPL.  You can redistribute it and/or
modify it under the terms of the GNU General Public License (GPL), either
version 3 of the License, or (at your option) any later version published
by the Free Software Foundation (FSF).
A copy of the license is included in @ref{GNU GPL}.
@end ifnothtml
@end titlepage

@c Table of Contents
@contents

@ifnottex
@node Top
@top GNU libunistring
@end ifnottex

@menu
* Introduction::                Who may need Unicode strings?
* Conventions::                 Conventions used in this manual
* unitypes.h::                  Elementary types
* unistr.h::                    Elementary Unicode string functions
* uniconv.h::                   Conversions between Unicode and encodings
* unistdio.h::                  Output with Unicode strings
* uniname.h::                   Names of Unicode characters
* unictype.h::                  Unicode character classification and properties
* uniwidth.h::                  Display width
* unigbrk.h::                   Grapheme cluster breaking
* uniwbrk.h::                   Word breaks in strings
* unilbrk.h::                   Line breaking
* uninorm.h::                   Normalization forms
* unicase.h::                   Case mappings
* uniregex.h::                  Regular expressions
* Using the library::           How to link with the library and use it?
* More functionality::          More advanced functionality
* The wchar_t mess::            Why @code{wchar_t *} strings are useless
* Licenses::                    Licenses

* Index::                       General Index

@detailmenu
 --- The Detailed Node Listing ---
```

(continues on next page)

(continued from previous page)

```
Introduction

* Unicode::                      What is Unicode?
* Unicode and i18n::             Unicode and internationalization
* Locale encodings::             What is a locale encoding?
* In-memory representation::     How to represent strings in memory?
* char * strings::               What to keep in mind with @code{char *} strings
* Unicode strings::              How are Unicode strings represented?

unistr.h

* Elementary string checks::
* Elementary string conversions::
* Elementary string functions::
* Elementary string functions with memory allocation::
* Elementary string functions on NUL terminated strings::

Elementary string functions

* Iterating::
* Creating Unicode strings::
* Copying Unicode strings::
* Comparing Unicode strings::
* Searching for a character::
* Counting characters::

Elementary string functions on NUL terminated strings

* Iterating over a NUL terminated Unicode string::
* Length::
* Copying a NUL terminated Unicode string::
* Comparing NUL terminated Unicode strings::
* Duplicating a NUL terminated Unicode string::
* Searching for a character in a NUL terminated Unicode string::
* Searching for a substring::
* Tokenizing::

unictype.h

* General category::
* Canonical combining class::
* Bidi class::
* Decimal digit value::
* Digit value::
* Numeric value::
* Mirrored character::
* Arabic shaping::
* Properties::
* Scripts::
* Blocks::
* ISO C and Java syntax::
```

(continues on next page)

(continued from previous page)

```
* Classifications like in ISO C::

General category

* Object oriented API::
* Bit mask API::

Properties

* Properties as objects::
* Properties as functions::

unigbrk.h

* Grapheme cluster breaks in a string::
* Grapheme cluster break property::

uniwbrk.h

* Word breaks in a string::
* Word break property::

uninorm.h

* Decomposition of characters::
* Composition of characters::
* Normalization of strings::
* Normalizing comparisons::
* Normalization of streams::

unicase,h

* Case mappings of characters::
* Case mappings of strings::
* Case mappings of substrings::
* Case insensitive comparison::
* Case detection::

Using the library

* Installation::
* Compiler options::
* Include files::
* Autoconf macro::
* Reporting problems::

Licenses

* GNU GPL::                      GNU General Public License
* GNU LGPL::                     GNU Lesser General Public License
* GNU FDL::                      GNU Free Documentation License
```

(continues on next page)

(continued from previous page)

```
@end detailmenu
@end menu

@node Introduction
@chapter Introduction

This library provides functions for manipulating Unicode strings and
for manipulating C strings according to the Unicode standard.

It consists of the following parts:

@table @code
@item <unistr.h>
elementary string functions
@item <uniconv.h>
conversion from/to legacy encodings
@item <unistdio.h>
formatted output to strings
@item <uniname.h>
character names
@item <unictype.h>
character classification and properties
@item <uniwidth.h>
string width when using nonproportional fonts
@item <unigbrk.h>
grapheme cluster breaks
@item <uniwbrk.h>
word breaks
@item <unilbrk.h>
line breaking algorithm
@item <uninorm.h>
normalization (composition and decomposition)
@item <unicase.h>
case folding
@item <uniregex.h>
regular expressions (not yet implemented)
@end table

@cindex use cases
@cindex value, of libunistring
libunistring is for you if your application involves non-trivial text
processing, such as upper/lower case conversions, line breaking, operations
on words, or more advanced analysis of text.  Text provided by the user can,
in general, contain characters of all kinds of scripts.  The text processing
functions provided by this library handle all scripts and all languages.

libunistring is for you if your application already uses the ISO C / POSIX
@posixheader{ctype.h}, @posixheader{wctype.h} functions and the text it
operates on is provided by the user and can be in any language.

libunistring is also for you if your application uses Unicode strings as
internal in-memory representation.
```

(continues on next page)

(continued from previous page)

```
@menu
* Unicode::                     What is Unicode?
* Unicode and i18n::            Unicode and internationalization
* Locale encodings::            What is a locale encoding?
* In-memory representation::    How to represent strings in memory?
* char * strings::              What to keep in mind with @code{char *} strings
* Unicode strings::             How are Unicode strings represented?
@end menu

@node Unicode
@section Unicode

@cindex Unicode
Unicode is a standardized repertoire of characters that contains characters
from all scripts of the world, from Latin letters to Chinese ideographs
and Babylonian cuneiform glyphs.  It also specifies how these characters
are to be rendered on a screen or on paper, and how common text processing
(word selection, line breaking, uppercasing of page titles etc.) is supposed
to behave on Unicode text.

Unicode also specifies three ways of storing sequences of Unicode
characters in a computer whose basic unit of data is an 8-bit byte:
@cindex UTF-8
@cindex UTF-16
@cindex UTF-32
@cindex UCS-4
@table @asis
@item UTF-8
Every character is represented as 1 to 4 bytes.
@item UTF-16
Every character is represented as 1 to 2 units of 16 bits.
@item UTF-32, a.k.a@. UCS-4
Every character is represented as 1 unit of 32 bits.
@end table

For encoding Unicode text in a file, UTF-8 is usually used.  For encoding
Unicode strings in memory for a program, either of the three encoding forms
can be reasonably used.

Unicode is widely used on the web.  Prior to the use of Unicode, web pages
were in many different encodings (ISO-8859-1 for English, French, Spanish,
ISO-8859-2 for Polish, ISO-8859-7 for Greek, KOI8-R for Russian, GB2312 or
BIG5 for Chinese, ISO-2022-JP-2 or EUC-JP or Shift_JIS for Japanese, and many
many others).  It was next to impossible to create a document that contained
Chinese and Polish text in the same document.  Due to the many encodings for
Japanese, even the processing of pure Japanese text was error prone.

References:
@itemize @bullet
@item
The Unicode standard:@texnl{} @url{http://www.unicode.org/}
```

(continues on next page)

(continued from previous page)

```
@item
Definition of UTF-8:@texnl{} @url{http://www.rfc-editor.org/rfc/rfc3629.txt}
@item
Definition of UTF-16:@texnl{} @url{http://www.rfc-editor.org/rfc/rfc2781.txt}
@item
Markus Kuhn's UTF-8 and Unicode FAQ:@texnl{}
@url{http://www.cl.cam.ac.uk/~mgk25/unicode.html}
@end itemize

@node Unicode and i18n
@section Unicode and Internationalization

@cindex internationalization
Internationalization is the process of changing the source code of a program
so that it can meet the expectations of users in any culture, if culture
specific data (translations, images etc.) are provided.

Use of Unicode is not strictly required for internationalization, but it
makes internationalization much easier, because operations that need to
look at specific characters (like hyphenation, spell checking, or the
automatic conversion of double-quotes to opening and closing double-quote
characters) don't need to consider multiple possible encodings of the text.

Use of Unicode also enables multilingualization: the ability of having text
in multiple languages present in the same document or even in the same line
of text.

But use of Unicode is not everything.  Internationalization usually consists
of four features:
@itemize @bullet
@item
Use of Unicode where needed for text processing.  This is what this library
is for.
@item
Use of message catalogs for messages shown to the user, This is what
GNU gettext is about.
@item
Use of locale specific conventions for date and time formats, for numeric
formatting, or for sorting of text.  This can be done adequately with the
POSIX APIs and the implementation of locales in the GNU C library.
@item
In graphical user interfaces, adapting the GUI to the default text direction
of the current locale (see
@url{https://en.wikipedia.org/wiki/Right-to-left,right-to-left languages}).
@end itemize

@node Locale encodings
@section Locale encodings

@cindex locale
A locale is a set of cultural conventions.  According to POSIX, for a program,
at any moment, there is one locale being designated as the ``current locale''.
```

(continues on next page)

(continued from previous page)

```
(Actually, POSIX supports also one locale per thread, but this feature is not
yet universally implemented and not widely used.)
@cindex locale categories
The locale is partitioned into several aspects, called the ``categories''
of the locale.  The main various aspects are:
@itemize @bullet
@item
The character encoding and the character properties.  This is the
@code{LC_CTYPE} category.
@item
The sorting rules for text.  This is the @code{LC_COLLATE} category.
@item
The language specific translations of messages.  This is the
@code{LC_MESSAGES} category.
@item
The formatting rules for numbers, such as the decimal separator.  This is
the @code{LC_NUMERIC} category.
@item
The formatting rules for amounts of money.  This is the @code{LC_MONETARY}
category.
@item
The formatting of date and time.  This is the @code{LC_TIME} category.
@end itemize

@cindex locale encoding
In particular, the @code{LC_CTYPE} category of the current locale determines
the character encoding.  This is the encoding of @samp{char *} strings.
We also call it the ``locale encoding''.  GNU libunistring has a function,
@func{locale_charset}, that returns a standardized (platform independent)
name for this encoding.

All locale encodings used on glibc systems are essentially ASCII compatible:
Most graphic ASCII characters have the same representation, as a single byte,
in that encoding as in ASCII.

Among the possible locale encodings are UTF-8 and GB18030.  Both allow
to represent any Unicode character as a sequence of bytes.  UTF-8 is used in
most of the world, whereas GB18030 is used in the People's Republic of China,
because it is backward compatible with the GB2312 encoding that was used in
this country earlier.

The legacy locale encodings, ISO-8859-15 (which supplanted ISO-8859-1 in
most of Europe), ISO-8859-2, KOI8-R, EUC-JP, etc., are still in use in
some places, though.

UTF-16 and UTF-32 are not used as locale encodings, because they are not
ASCII compatible.

@node In-memory representation
@section Choice of in-memory representation of strings

There are three ways of representing strings in memory of a running
```

(continues on next page)

(continued from previous page)

```
program.
@itemize @bullet
@item
As @samp{char *} strings.  Such strings are represented in locale encoding.
This approach is employed when not much text processing is done by the
program.  When some Unicode aware processing is to be done, a string is
converted to Unicode on the fly and back to locale encoding afterwards.
@item
As UTF-8 or UTF-16 or UTF-32 strings.  This implies that conversion from
locale encoding to Unicode is performed on input, and in the opposite
direction on output.  This approach is employed when the program does
a significant amount of text processing, or when the program has multiple
threads operating on the same data but in different locales.
@item
As @samp{wchar_t *}, a.k.a@. ``wide strings''.  This approach is misguided,
see @ref{The wchar_t mess}.
@end itemize

Of course, a @samp{char *} string can, in some cases, be encoded in UTF-8.
You will use the data type depending on what you can guarantee about how
it's encoded: If a string is encoded in the locale encoding, or if you
don't know how it's encoded, use @samp{char *}.  If, on the other hand,
you can @emph{guarantee} that it is UTF-8 encoded, then you can use the
UTF-8 string type, @code{uint8_t *}, for it.

The five types @code{char *}, @code{uint8_t *}, @code{uint16_t *},
@code{uint32_t *}, and @code{wchar_t *} are incompatible types at the C
level.  Therefore, @samp{gcc -Wall} will produce a warning if, by mistake,
your code contains a mismatch between these types.  In the context of
using GNU libunistring, even a warning about a mismatch between
@code{char *} and @code{uint8_t *} is a sign of a bug in your code
that you should not try to silence through a cast.

@node char * strings
@section @samp{char *} strings

@cindex C string functions
The classical C strings, with its C library support standardized by
ISO C and POSIX, can be used in internationalized programs with some
precautions.  The problem with this API is that many of the C library
functions for strings don't work correctly on strings in locale
encodings, leading to bugs that only people in some cultures of the
world will experience.

@cindex locale, multibyte
The first problem with the C library API is the support of multibyte
locales.  According to the locale encoding, in general, every character
is represented by one or more bytes (up to 4 bytes in practice --- but
use @code{MB_LEN_MAX} instead of the number 4 in the code).
When every character is represented by only 1 byte, we speak of an
``unibyte locale'', otherwise of a ``multibyte locale''.  It is important
to realize that the majority of Unix installations nowadays use UTF-8
```

(continues on next page)

(continued from previous page)

```
or GB18030 as locale encoding; therefore, the majority of users are
using multibyte locales.

@cindex char, type
The important fact to remember is:
@cartouche
@emph{A @samp{char} is a byte, not a character.}
@end cartouche

As a consequence:
@itemize @bullet
@item
The @posixheader{ctype.h} API is useless in this context; it does not work in
multibyte locales.
@item
The @posixfunc{strlen} function does not return the number of characters
in a string.  Nor does it return the number of screen columns occupied
by a string after it is output.  It merely returns the number of
@emph{bytes} occupied by a string.
@item
Truncating a string, for example, with @posixfunc{strncpy}, can have the
effect of truncating it in the middle of a multibyte character.  Such
a string will, when output, have a garbled character at its end, often
represented by a hollow box.
@item
@posixfunc{strchr} and @posixfunc{strrchr} do not work with multibyte strings
if the locale encoding is GB18030 and the character to be searched is
a digit.
@item
@posixfunc{strstr} does not work with multibyte strings if the locale encoding
is different from UTF-8.
@item
@posixfunc{strcspn}, @posixfunc{strpbrk}, @posixfunc{strspn} cannot work
correctly in multibyte locales: they assume the second argument is a list of
single-byte characters.  Even in this simple case, they do not work with
multibyte strings if the locale encoding is GB18030 and one of the
characters to be searched is a digit.
@item
@posixfunc{strsep} and @posixfunc{strtok_r} do not work with multibyte strings
unless all of the delimiter characters are ASCII characters < 0x30.
@item
The @posixfunc{strcasecmp}, @posixfunc{strncasecmp}, and @posixfunc{strcasestr}
functions do not work with multibyte strings.
@end itemize

The workarounds can be found in GNU gnulib
@url{http://www.gnu.org/software/gnulib/}.
@itemize @bullet
@item
gnulib has modules @samp{mbchar}, @samp{mbiter}, @samp{mbuiter} that
represent multibyte characters and allow to iterate across a multibyte
string with the same ease as through a unibyte string.
```

(continues on next page)

(continued from previous page)

```
@item
gnulib has functions @func{mbslen} and @func{mbswidth} that can be
used instead of @posixfunc{strlen} when the number of characters or the
number of screen columns of a string is requested.
@item
gnulib has functions @func{mbschr} and @func{mbsrrchr} that are
like @posixfunc{strchr} and @posixfunc{strrchr}, but work in multibyte locales.
@item
gnulib has a function @func{mbsstr}, like @posixfunc{strstr}, but works
in multibyte locales.
@item
gnulib has functions @func{mbscspn}, @func{mbspbrk}, @func{mbsspn}
that are like @posixfunc{strcspn}, @posixfunc{strpbrk}, @posixfunc{strspn}, but
work in multibyte locales.
@item
gnulib has functions @func{mbssep} and @func{mbstok_r} that are
like @posixfunc{strsep} and @posixfunc{strtok_r} but work in multibyte locales.
@item
gnulib has functions @func{mbscasecmp}, @func{mbsncasecmp},
@func{mbspcasecmp}, and @func{mbscasestr} that are like @posixfunc{strcasecmp},
@posixfunc{strncasecmp}, and @posixfunc{strcasestr}, but
work in multibyte locales.  Still, the function @code{ulc_casecmp} is
preferable to these functions; see below.
@end itemize

The second problem with the C library API is that it has some assumptions built-in that␣
↪are not valid in some languages:
@itemize @bullet
@item
It assumes that there are only two forms of every character: uppercase
and lowercase.  This is not true for Croatian, where the character
@sc{LETTER DZ WITH CARON} comes in three forms:
@sc{LATIN CAPITAL LETTER DZ WITH CARON} (DZ),
@sc{LATIN CAPITAL LETTER D WITH SMALL LETTER Z WITH CARON} (Dz),
@sc{LATIN SMALL LETTER DZ WITH CARON} (dz).
@item
It assumes that uppercasing of 1 character leads to 1 character.  This
is not true for German, where the @sc{LATIN SMALL LETTER SHARP S}, when
uppercased, becomes @samp{SS}.
@item
It assumes that there is 1:1 mapping between uppercase and lowercase forms.
This is not true for the Greek sigma: @sc{GREEK CAPITAL LETTER SIGMA} is
the uppercase of both @sc{GREEK SMALL LETTER SIGMA} and
@sc{GREEK SMALL LETTER FINAL SIGMA}.
@item
It assumes that the upper/lowercase mappings are position independent.
This is not true for the Greek sigma and the Lithuanian i.
@end itemize

The correct way to deal with this problem is
@enumerate
@item
```

(continues on next page)

(continued from previous page)

```
to provide functions for titlecasing, as well as for upper- and
lowercasing,
@item
to view case transformations as functions that operates on strings,
rather than on characters.
@end enumerate

This is implemented in this library, through the functions declared in @code{<unicase.h>}
↪, see @ref{unicase.h}.

@node Unicode strings
@section Unicode strings

libunistring supports Unicode strings in three representations:
@cindex UTF-8, strings
@cindex UTF-16, strings
@cindex UTF-32, strings
@itemize @bullet
@item
UTF-8 strings, through the type @samp{uint8_t *}.  The units are bytes
(@code{uint8_t}).
@item
UTF-16 strings, through the type @samp{uint16_t *},  The units are 16-bit
memory words (@code{uint16_t}).
@item
UTF-32 strings, through the type @samp{uint32_t *}.  The units are 32-bit
memory words (@code{uint32_t}).
@end itemize

As with C strings, there are two variants:
@itemize @bullet
@item
Unicode strings with a terminating NUL character are represented as
a pointer to the first unit of the string.  There is a unit containing
a 0 value at the end.  It is considered part of the string for all
memory allocation purposes, but is not considered part of the string
for all other logical purposes.
@item
Unicode strings where embedded NUL characters are allowed.  These
are represented by a pointer to the first unit and the number of units
(not bytes!) of the string.  In this setting, there is no trailing
zero-valued unit used as ``end marker''.
@end itemize

@node Conventions
@chapter Conventions

This chapter explains conventions valid throughout the libunistring library.

@cindex argument conventions
Variables of type @code{char *} denote C strings in locale encoding.
See @ref{Locale encodings}.
```

(continues on next page)

(continued from previous page)

```
Variables of type @code{uint8_t *} denote UTF-8 strings.  Their units
are bytes.

Variables of type @code{uint16_t *} denote UTF-16 strings, without byte
order mark.  Their units are 2-byte words.

Variables of type @code{uint32_t *} denote UTF-32 strings, without byte
order mark.  Their units are 4-byte words.

Argument pairs @code{(@var{s}, @var{n})} denote a string
@code{@var{s}[0..@var{n}-1]} with exactly @var{n} units.

All functions with prefix @samp{ulc_} operate on C strings in locale
encoding.

All functions with prefix @samp{u8_} operate on UTF-8 strings.

All functions with prefix @samp{u16_} operate on UTF-16 strings.

All functions with prefix @samp{u32_} operate on UTF-32 strings.

For every function with prefix @samp{u8_}, operating on UTF-8 strings,
there is also a corresponding function with prefix @samp{u16_},
operating on UTF-16 strings, and a corresponding function with prefix
@samp{u32_}, operating on UTF-32 strings.  Their description is
analogous; in this documentation we describe only the function that
operates on UTF-8 strings, for brevity.

A declaration with a variable @var{n} denotes the three concrete
declarations with @var{n} = 8, @var{n} = 16, @var{n} = 32.

All parameters starting with @samp{str} and the parameters of
functions starting with @code{u8_str}/@code{u16_str}/@code{u32_str}
denote a NUL terminated string.

@cindex return value conventions
Error values are always returned through the @code{errno} variable,
usually with a return value that indicates the presence of an error
(NULL for functions that return an pointer, or -1 for functions that
return an @code{int}).

Functions returning a string result take a
@code{(@var{resultbuf}, @var{lengthp})}
argument pair.  If @var{resultbuf} is not NULL and the result fits
into @code{*@var{lengthp}} units, it is put in @var{resultbuf}, and
@var{resultbuf} is returned.  Otherwise, a freshly allocated string
is returned.  In both cases, @code{*@var{lengthp}} is set to the
length (number of units) of the returned string.  In case of error,
NULL is returned and @code{errno} is set.

@include unitypes.texi
```

(continues on next page)

(continued from previous page)

```
@include unistr.texi
@include uniconv.texi
@include unistdio.texi
@include uniname.texi
@include unictype.texi
@include uniwidth.texi
@include unigbrk.texi
@include uniwbrk.texi
@include unilbrk.texi
@include uninorm.texi
@include unicase.texi
@include uniregex.texi

@node Using the library
@chapter Using the library

This chapter explains some practical considerations, regarding the
installation and compiler options that are needed in order to use this
library.

@menu
* Installation::
* Compiler options::
* Include files::
* Autoconf macro::
* Reporting problems::
@end menu

@node Installation
@section Installation

@cindex dependencies
Before you can use the library, it must be installed.  First, you have to
make sure all dependencies are installed.  They are listed in the file
@file{DEPENDENCIES}.

@cindex installation
Then you can proceed to build and install the library, as described in the
file @file{INSTALL}.  For installation on Windows systems, please refer to
the file @file{INSTALL.windows}.

@node Compiler options
@section Compiler options

Let's denote as @code{LIBUNISTRING_PREFIX} the value of the @samp{--prefix}
option that you passed to @code{configure} while installing this package.
If you didn't pass any @samp{--prefix} option, then the package is installed
in @file{/usr/local}.

Let's denote as @code{LIBUNISTRING_INCLUDEDIR} the directory where the
include files were installed.  This is usually the same as
@code{$@{LIBUNISTRING_PREFIX@}/include}.  Except that if you passed an
```

(continues on next page)

(continued from previous page)

```
@samp{--includedir} option to @code{configure}, it is the value of that
option.

Let's further denote as @code{LIBUNISTRING_LIBDIR} the directory where
the library itself was installed.  This is the value that you passed
with the @samp{--libdir} option to @code{configure}, or otherwise the
same as @code{$@{LIBUNISTRING_PREFIX@}/lib}.  Recall that when building
in 64-bit mode on a 64-bit GNU/Linux system that supports executables
in either 64-bit mode or 32-bit mode, you should have used the option
@code{--libdir=$@{LIBUNISTRING_PREFIX@}/lib64}.

@cindex compiler options
So that the compiler finds the include files, you have to pass it the
option @code{-I$@{LIBUNISTRING_INCLUDEDIR@}}.

So that the compiler finds the library during its linking pass, you have
to pass it the options @code{-L$@{LIBUNISTRING_LIBDIR@} -lunistring}.
On some systems, in some configurations, you also have to pass options
needed for linking with @code{libiconv}.  The autoconf macro
@code{gl_LIBUNISTRING} (see @ref{Autoconf macro}) deals with this
particularity.

@node Include files
@section Include files

Most of the include files have been presented in the introduction, see
@ref{Introduction}, and subsequent detailed chapters.

Another include file is @code{<unistring/version.h>}. It contains the
version number of the libunistring library.

@deftypevr Macro int _LIBUNISTRING_VERSION
This constant contains the version of libunistring that is being used
at compile time.  It encodes the major and minor parts of the version
number only.  These parts are encoded in the form @code{(major<<8) + minor}.
@end deftypevr

@deftypevr Constant int _libunistring_version
This constant contains the version of libunistring that is being used
at run time.  It encodes the major and minor parts of the version
number only.  These parts are encoded in the form @code{(major<<8) + minor}.
@end deftypevr

It is possible that @code{_libunistring_version} is greater than
@code{_LIBUNISTRING_VERSION}.  This can happen when you use
@code{libunistring} as a shared library, and a newer, binary
backward-compatible version has been installed after your program
that uses @code{libunistring} was installed.

@node Autoconf macro
@section Autoconf macro
```

(continues on next page)

(continued from previous page)

```
@cindex autoconf macro
GNU Gnulib provides an autoconf macro that tests for the availability
of @code{libunistring}.  It is contained in the Gnulib module
@samp{libunistring}, see@texnl{}
@url{http://www.gnu.org/software/gnulib/MODULES.html#module=libunistring}.

@amindex gl_LIBUNISTRING
The macro is called @code{gl_LIBUNISTRING}.  It searches for an installed
libunistring.  If found, it sets and AC_SUBSTs @code{HAVE_LIBUNISTRING=yes}
and the @code{LIBUNISTRING} and @code{LTLIBUNISTRING} variables and augments
the @code{CPPFLAGS} variable, and defines the C macro
@code{HAVE_LIBUNISTRING} to 1.  Otherwise, it sets and AC_SUBSTs
@code{HAVE_LIBUNISTRING=no} and @code{LIBUNISTRING} and @code{LTLIBUNISTRING}
to empty.

The complexities that @code{gl_LIBUNISTRING} deals with are the following:

@itemize @bullet
@item
On some operating systems, in some configurations, libunistring depends
on @code{libiconv}, and the options for linking with libiconv must be
mentioned explicitly on the link command line.

@item
GNU @code{libunistring}, if installed, is not necessarily already in the
search path (@code{CPPFLAGS} for the include file search path,
@code{LDFLAGS} for the library search path).

@item
GNU @code{libunistring}, if installed, is not necessarily already in the
run time library search path.  To avoid the need for setting an environment
variable like @code{LD_LIBRARY_PATH}, the macro adds the appropriate
run time search path options to the @code{LIBUNISTRING} variable.  This works
on most systems.
@end itemize

@node Reporting problems
@section Reporting problems

@cindex bug reports
@cindex bug tracker
@cindex mailing list
If you encounter any problem, please don't hesitate to send a detailed
bug report to the @code{bug-libunistring@@gnu.org} mailing list.  You can
alternatively also use the bug tracker at the project page
@url{https://savannah.gnu.org/projects/libunistring}.

Please always include the version number of this library, and a short
description of your operating system and compilation environment with
corresponding version numbers.

For problems that appear while building and installing @code{libunistring},
```

(continues on next page)

(continued from previous page)

```
for which you don't find the remedy in the @file{INSTALL} file, please include
a description of the options that you passed to the @samp{configure} script.

@node More functionality
@chapter More advanced functionality

@cindex bidirectional reordering
For bidirectional reordering of strings, we recommend the GNU FriBidi library:
@url{http://www.fribidi.org/}.

@cindex rendering
For the rendering of Unicode strings outside of the context of a given toolkit
(KDE/Qt or GNOME/Gtk), we recommend the Pango library:
@url{http://www.pango.org/}.

@include wchar_t.texi

@node Licenses
@appendix Licenses
@cindex Licenses

The files of this package are covered by the licenses indicated in each
particular file or directory.  Here is a summary:

@itemize @bullet
@item
The @code{libunistring} library and its header files are dual-licensed under
"the GNU LGPLv3+ or the GNU GPLv2". This means, you can use it under either
@itemize @bullet
@item @minus{}
the terms of the GNU Lesser General Public License (LGPL) version 3 or
(at your option) any later version, or
@item @minus{}
the terms of the GNU General Public License (GPL) version 2, or
@item @minus{}
the same dual license "the GNU LGPLv3+ or the GNU GPLv2".
@end itemize
You find the GNU LGPL version 3 in @ref{GNU LGPL}.  This license is
based on the GNU GPL version 3, see @ref{GNU GPL}.
@*
You can find the GNU GPL version 2 at
@url{https://www.gnu.org/licenses/old-licenses/gpl-2.0.html}.
@*
Note: This dual license makes it possible for the @code{libunistring} library
to be used by packages under GPLv2 or GPLv2+ licenses, in particular. See
the table in @url{https://www.gnu.org/licenses/gpl-faq.html#AllCompatibility}.


@item
This manual is free documentation.  It is dually licensed under the
GNU FDL and the GNU GPL.  This means that you can redistribute this
manual under either of these two licenses, at your choice.
```

(continues on next page)

(continued from previous page)

```
@*
This manual is covered by the GNU FDL.  Permission is granted to copy,
distribute and/or modify this document under the terms of the
GNU Free Documentation License (FDL), either version 1.2 of the
License, or (at your option) any later version published by the
Free Software Foundation (FSF); with no Invariant Sections, with no
Front-Cover Text, and with no Back-Cover Texts.
A copy of the license is included in @ref{GNU FDL}.
@*
This manual is covered by the GNU GPL.  You can redistribute it and/or
modify it under the terms of the GNU General Public License (GPL), either
version 3 of the License, or (at your option) any later version published
by the Free Software Foundation (FSF).
A copy of the license is included in @ref{GNU GPL}.
@end itemize

@menu
* GNU GPL::                     GNU General Public License
* GNU LGPL::                    GNU Lesser General Public License
* GNU FDL::                     GNU Free Documentation License
@end menu

@page
@node GNU GPL
@appendixsec GNU GENERAL PUBLIC LICENSE
@cindex GPL, GNU General Public License
@cindex License, GNU GPL
@include gpl.texi
@page
@node GNU LGPL
@appendixsec GNU LESSER GENERAL PUBLIC LICENSE
@cindex LGPL, GNU Lesser General Public License
@cindex License, GNU LGPL
@include lgpl.texi
@page
@node GNU FDL
@appendixsec GNU Free Documentation License
@cindex FDL, GNU Free Documentation License
@cindex License, GNU FDL
@include fdl.texi

@node Index
@unnumbered Index

@printindex cp

@bye

@c Local Variables:
@c indent-tabs-mode: nil
@c whitespace-check-buffer-indent: nil
@c End:
```

40.50 GPLv2 License Text for libunistring

The following license text applies to libunistring:

Listing 50: Download: GPLv2 for libunistring

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 3, 29 June 2007

 Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The GNU General Public License is a free, copyleft license for
software and other kinds of works.

  The licenses for most software and other practical works are designed
to take away your freedom to share and change the works.  By contrast,
the GNU General Public License is intended to guarantee your freedom to
share and change all versions of a program--to make sure it remains free
software for all its users.  We, the Free Software Foundation, use the
GNU General Public License for most of our software; it applies also to
any other work released this way by its authors.  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
them if you wish), that you receive source code or can get it if you
want it, that you can change the software or use pieces of it in new
free programs, and that you know you can do these things.

  To protect your rights, we need to prevent others from denying you
these rights or asking you to surrender the rights.  Therefore, you have
certain responsibilities if you distribute copies of the software, or if
you modify it: responsibilities to respect the freedom of others.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must pass on to the recipients the same
freedoms that you received.  You must make sure that they, too, receive
or can get the source code.  And you must show them these terms so they
know their rights.

  Developers that use the GNU GPL protect your rights with two steps:
(1) assert copyright on the software, and (2) offer you this License
giving you legal permission to copy, distribute and/or modify it.

  For the developers' and authors' protection, the GPL clearly explains
that there is no warranty for this free software.  For both users' and
authors' sake, the GPL requires that modified versions be marked as
changed, so that their problems will not be attributed erroneously to
```

(continues on next page)

(continued from previous page)

```
authors of previous versions.

  Some devices are designed to deny users access to install or run
modified versions of the software inside them, although the manufacturer
can do so.  This is fundamentally incompatible with the aim of
protecting users' freedom to change the software.  The systematic
pattern of such abuse occurs in the area of products for individuals to
use, which is precisely where it is most unacceptable.  Therefore, we
have designed this version of the GPL to prohibit the practice for those
products.  If such problems arise substantially in other domains, we
stand ready to extend this provision to those domains in future versions
of the GPL, as needed to protect the freedom of users.

  Finally, every program is threatened constantly by software patents.
States should not allow patents to restrict development and use of
software on general-purpose computers, but in those that do, we wish to
avoid the special danger that patents applied to a free program could
make it effectively proprietary.  To prevent this, the GPL assures that
patents cannot be used to render the program non-free.

  The precise terms and conditions for copying, distribution and
modification follow.

                       TERMS AND CONDITIONS

  0. Definitions.

  "This License" refers to version 3 of the GNU General Public License.

  "Copyright" also means copyright-like laws that apply to other kinds of
works, such as semiconductor masks.

  "The Program" refers to any copyrightable work licensed under this
License.  Each licensee is addressed as "you".  "Licensees" and
"recipients" may be individuals or organizations.

  To "modify" a work means to copy from or adapt all or part of the work
in a fashion requiring copyright permission, other than the making of an
exact copy.  The resulting work is called a "modified version" of the
earlier work or a work "based on" the earlier work.

  A "covered work" means either the unmodified Program or a work based
on the Program.

  To "propagate" a work means to do anything with it that, without
permission, would make you directly or secondarily liable for
infringement under applicable copyright law, except executing it on a
computer or modifying a private copy.  Propagation includes copying,
distribution (with or without modification), making available to the
public, and in some countries other activities as well.

  To "convey" a work means any kind of propagation that enables other
```

(continues on next page)

(continued from previous page)

```
parties to make or receive copies.  Mere interaction with a user through
a computer network, with no transfer of a copy, is not conveying.

  An interactive user interface displays "Appropriate Legal Notices"
to the extent that it includes a convenient and prominently visible
feature that (1) displays an appropriate copyright notice, and (2)
tells the user that there is no warranty for the work (except to the
extent that warranties are provided), that licensees may convey the
work under this License, and how to view a copy of this License.  If
the interface presents a list of user commands or options, such as a
menu, a prominent item in the list meets this criterion.

  1. Source Code.

  The "source code" for a work means the preferred form of the work
for making modifications to it.  "Object code" means any non-source
form of a work.

  A "Standard Interface" means an interface that either is an official
standard defined by a recognized standards body, or, in the case of
interfaces specified for a particular programming language, one that
is widely used among developers working in that language.

  The "System Libraries" of an executable work include anything, other
than the work as a whole, that (a) is included in the normal form of
packaging a Major Component, but which is not part of that Major
Component, and (b) serves only to enable use of the work with that
Major Component, or to implement a Standard Interface for which an
implementation is available to the public in source code form.  A
"Major Component", in this context, means a major essential component
(kernel, window system, and so on) of the specific operating system
(if any) on which the executable work runs, or a compiler used to
produce the work, or an object code interpreter used to run it.

  The "Corresponding Source" for a work in object code form means all
the source code needed to generate, install, and (for an executable
work) run the object code and to modify the work, including scripts to
control those activities.  However, it does not include the work's
System Libraries, or general-purpose tools or generally available free
programs which are used unmodified in performing those activities but
which are not part of the work.  For example, Corresponding Source
includes interface definition files associated with source files for
the work, and the source code for shared libraries and dynamically
linked subprograms that the work is specifically designed to require,
such as by intimate data communication or control flow between those
subprograms and other parts of the work.

  The Corresponding Source need not include anything that users
can regenerate automatically from other parts of the Corresponding
Source.

  The Corresponding Source for a work in source code form is that
```

(continues on next page)

(continued from previous page)

```
same work.

  2. Basic Permissions.

  All rights granted under this License are granted for the term of
copyright on the Program, and are irrevocable provided the stated
conditions are met.  This License explicitly affirms your unlimited
permission to run the unmodified Program.  The output from running a
covered work is covered by this License only if the output, given its
content, constitutes a covered work.  This License acknowledges your
rights of fair use or other equivalent, as provided by copyright law.

  You may make, run and propagate covered works that you do not
convey, without conditions so long as your license otherwise remains
in force.  You may convey covered works to others for the sole purpose
of having them make modifications exclusively for you, or provide you
with facilities for running those works, provided that you comply with
the terms of this License in conveying all material for which you do
not control copyright.  Those thus making or running the covered works
for you must do so exclusively on your behalf, under your direction
and control, on terms that prohibit them from making any copies of
your copyrighted material outside their relationship with you.

  Conveying under any other circumstances is permitted solely under
the conditions stated below.  Sublicensing is not allowed; section 10
makes it unnecessary.

  3. Protecting Users' Legal Rights From Anti-Circumvention Law.

  No covered work shall be deemed part of an effective technological
measure under any applicable law fulfilling obligations under article
11 of the WIPO copyright treaty adopted on 20 December 1996, or
similar laws prohibiting or restricting circumvention of such
measures.

  When you convey a covered work, you waive any legal power to forbid
circumvention of technological measures to the extent such circumvention
is effected by exercising rights under this License with respect to
the covered work, and you disclaim any intention to limit operation or
modification of the work as a means of enforcing, against the work's
users, your or third parties' legal rights to forbid circumvention of
technological measures.

  4. Conveying Verbatim Copies.

  You may convey verbatim copies of the Program's source code as you
receive it, in any medium, provided that you conspicuously and
appropriately publish on each copy an appropriate copyright notice;
keep intact all notices stating that this License and any
non-permissive terms added in accord with section 7 apply to the code;
keep intact all notices of the absence of any warranty; and give all
recipients a copy of this License along with the Program.
```

(continues on next page)

(continued from previous page)

```
  You may charge any price or no price for each copy that you convey,
and you may offer support or warranty protection for a fee.

  5. Conveying Modified Source Versions.

  You may convey a work based on the Program, or the modifications to
produce it from the Program, in the form of source code under the
terms of section 4, provided that you also meet all of these conditions:

    a) The work must carry prominent notices stating that you modified
    it, and giving a relevant date.

    b) The work must carry prominent notices stating that it is
    released under this License and any conditions added under section
    7.  This requirement modifies the requirement in section 4 to
    "keep intact all notices".

    c) You must license the entire work, as a whole, under this
    License to anyone who comes into possession of a copy.  This
    License will therefore apply, along with any applicable section 7
    additional terms, to the whole of the work, and all its parts,
    regardless of how they are packaged.  This License gives no
    permission to license the work in any other way, but it does not
    invalidate such permission if you have separately received it.

    d) If the work has interactive user interfaces, each must display
    Appropriate Legal Notices; however, if the Program has interactive
    interfaces that do not display Appropriate Legal Notices, your
    work need not make them do so.

  A compilation of a covered work with other separate and independent
works, which are not by their nature extensions of the covered work,
and which are not combined with it such as to form a larger program,
in or on a volume of a storage or distribution medium, is called an
"aggregate" if the compilation and its resulting copyright are not
used to limit the access or legal rights of the compilation's users
beyond what the individual works permit.  Inclusion of a covered work
in an aggregate does not cause this License to apply to the other
parts of the aggregate.

  6. Conveying Non-Source Forms.

  You may convey a covered work in object code form under the terms
of sections 4 and 5, provided that you also convey the
machine-readable Corresponding Source under the terms of this License,
in one of these ways:

    a) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by the
    Corresponding Source fixed on a durable physical medium
    customarily used for software interchange.
```

(continues on next page)

(continued from previous page)

```
    b) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by a
    written offer, valid for at least three years and valid for as
    long as you offer spare parts or customer support for that product
    model, to give anyone who possesses the object code either (1) a
    copy of the Corresponding Source for all the software in the
    product that is covered by this License, on a durable physical
    medium customarily used for software interchange, for a price no
    more than your reasonable cost of physically performing this
    conveying of source, or (2) access to copy the
    Corresponding Source from a network server at no charge.

    c) Convey individual copies of the object code with a copy of the
    written offer to provide the Corresponding Source.  This
    alternative is allowed only occasionally and noncommercially, and
    only if you received the object code with such an offer, in accord
    with subsection 6b.

    d) Convey the object code by offering access from a designated
    place (gratis or for a charge), and offer equivalent access to the
    Corresponding Source in the same way through the same place at no
    further charge.  You need not require recipients to copy the
    Corresponding Source along with the object code.  If the place to
    copy the object code is a network server, the Corresponding Source
    may be on a different server (operated by you or a third party)
    that supports equivalent copying facilities, provided you maintain
    clear directions next to the object code saying where to find the
    Corresponding Source.  Regardless of what server hosts the
    Corresponding Source, you remain obligated to ensure that it is
    available for as long as needed to satisfy these requirements.

    e) Convey the object code using peer-to-peer transmission, provided
    you inform other peers where the object code and Corresponding
    Source of the work are being offered to the general public at no
    charge under subsection 6d.

  A separable portion of the object code, whose source code is excluded
from the Corresponding Source as a System Library, need not be
included in conveying the object code work.

  A "User Product" is either (1) a "consumer product", which means any
tangible personal property which is normally used for personal, family,
or household purposes, or (2) anything designed or sold for incorporation
into a dwelling.  In determining whether a product is a consumer product,
doubtful cases shall be resolved in favor of coverage.  For a particular
product received by a particular user, "normally used" refers to a
typical or common use of that class of product, regardless of the status
of the particular user or of the way in which the particular user
actually uses, or expects or is expected to use, the product.  A product
is a consumer product regardless of whether the product has substantial
commercial, industrial or non-consumer uses, unless such uses represent
```

(continues on next page)

(continued from previous page)

```
the only significant mode of use of the product.

  "Installation Information" for a User Product means any methods,
procedures, authorization keys, or other information required to install
and execute modified versions of a covered work in that User Product from
a modified version of its Corresponding Source.  The information must
suffice to ensure that the continued functioning of the modified object
code is in no case prevented or interfered with solely because
modification has been made.

  If you convey an object code work under this section in, or with, or
specifically for use in, a User Product, and the conveying occurs as
part of a transaction in which the right of possession and use of the
User Product is transferred to the recipient in perpetuity or for a
fixed term (regardless of how the transaction is characterized), the
Corresponding Source conveyed under this section must be accompanied
by the Installation Information.  But this requirement does not apply
if neither you nor any third party retains the ability to install
modified object code on the User Product (for example, the work has
been installed in ROM).

  The requirement to provide Installation Information does not include a
requirement to continue to provide support service, warranty, or updates
for a work that has been modified or installed by the recipient, or for
the User Product in which it has been modified or installed.  Access to a
network may be denied when the modification itself materially and
adversely affects the operation of the network or violates the rules and
protocols for communication across the network.

  Corresponding Source conveyed, and Installation Information provided,
in accord with this section must be in a format that is publicly
documented (and with an implementation available to the public in
source code form), and must require no special password or key for
unpacking, reading or copying.

  7. Additional Terms.

  "Additional permissions" are terms that supplement the terms of this
License by making exceptions from one or more of its conditions.
Additional permissions that are applicable to the entire Program shall
be treated as though they were included in this License, to the extent
that they are valid under applicable law.  If additional permissions
apply only to part of the Program, that part may be used separately
under those permissions, but the entire Program remains governed by
this License without regard to the additional permissions.

  When you convey a copy of a covered work, you may at your option
remove any additional permissions from that copy, or from any part of
it.  (Additional permissions may be written to require their own
removal in certain cases when you modify the work.)  You may place
additional permissions on material, added by you to a covered work,
for which you have or can give appropriate copyright permission.
```

(continues on next page)

(continued from previous page)

```
  Notwithstanding any other provision of this License, for material you
add to a covered work, you may (if authorized by the copyright holders of
that material) supplement the terms of this License with terms:

    a) Disclaiming warranty or limiting liability differently from the
    terms of sections 15 and 16 of this License; or

    b) Requiring preservation of specified reasonable legal notices or
    author attributions in that material or in the Appropriate Legal
    Notices displayed by works containing it; or

    c) Prohibiting misrepresentation of the origin of that material, or
    requiring that modified versions of such material be marked in
    reasonable ways as different from the original version; or

    d) Limiting the use for publicity purposes of names of licensors or
    authors of the material; or

    e) Declining to grant rights under trademark law for use of some
    trade names, trademarks, or service marks; or

    f) Requiring indemnification of licensors and authors of that
    material by anyone who conveys the material (or modified versions of
    it) with contractual assumptions of liability to the recipient, for
    any liability that these contractual assumptions directly impose on
    those licensors and authors.

  All other non-permissive additional terms are considered "further
restrictions" within the meaning of section 10.  If the Program as you
received it, or any part of it, contains a notice stating that it is
governed by this License along with a term that is a further
restriction, you may remove that term.  If a license document contains
a further restriction but permits relicensing or conveying under this
License, you may add to a covered work material governed by the terms
of that license document, provided that the further restriction does
not survive such relicensing or conveying.

  If you add terms to a covered work in accord with this section, you
must place, in the relevant source files, a statement of the
additional terms that apply to those files, or a notice indicating
where to find the applicable terms.

  Additional terms, permissive or non-permissive, may be stated in the
form of a separately written license, or stated as exceptions;
the above requirements apply either way.

  8. Termination.

  You may not propagate or modify a covered work except as expressly
provided under this License.  Any attempt otherwise to propagate or
modify it is void, and will automatically terminate your rights under
```

(continues on next page)

(continued from previous page)

```
this License (including any patent licenses granted under the third
paragraph of section 11).

  However, if you cease all violation of this License, then your
license from a particular copyright holder is reinstated (a)
provisionally, unless and until the copyright holder explicitly and
finally terminates your license, and (b) permanently, if the copyright
holder fails to notify you of the violation by some reasonable means
prior to 60 days after the cessation.

  Moreover, your license from a particular copyright holder is
reinstated permanently if the copyright holder notifies you of the
violation by some reasonable means, this is the first time you have
received notice of violation of this License (for any work) from that
copyright holder, and you cure the violation prior to 30 days after
your receipt of the notice.

  Termination of your rights under this section does not terminate the
licenses of parties who have received copies or rights from you under
this License.  If your rights have been terminated and not permanently
reinstated, you do not qualify to receive new licenses for the same
material under section 10.

  9. Acceptance Not Required for Having Copies.

  You are not required to accept this License in order to receive or
run a copy of the Program.  Ancillary propagation of a covered work
occurring solely as a consequence of using peer-to-peer transmission
to receive a copy likewise does not require acceptance.  However,
nothing other than this License grants you permission to propagate or
modify any covered work.  These actions infringe copyright if you do
not accept this License.  Therefore, by modifying or propagating a
covered work, you indicate your acceptance of this License to do so.

  10. Automatic Licensing of Downstream Recipients.

  Each time you convey a covered work, the recipient automatically
receives a license from the original licensors, to run, modify and
propagate that work, subject to this License.  You are not responsible
for enforcing compliance by third parties with this License.

  An "entity transaction" is a transaction transferring control of an
organization, or substantially all assets of one, or subdividing an
organization, or merging organizations.  If propagation of a covered
work results from an entity transaction, each party to that
transaction who receives a copy of the work also receives whatever
licenses to the work the party's predecessor in interest had or could
give under the previous paragraph, plus a right to possession of the
Corresponding Source of the work from the predecessor in interest, if
the predecessor has it or can get it with reasonable efforts.

  You may not impose any further restrictions on the exercise of the
```

(continues on next page)

(continued from previous page)

```
rights granted or affirmed under this License.  For example, you may
not impose a license fee, royalty, or other charge for exercise of
rights granted under this License, and you may not initiate litigation
(including a cross-claim or counterclaim in a lawsuit) alleging that
any patent claim is infringed by making, using, selling, offering for
sale, or importing the Program or any portion of it.

  11. Patents.

  A "contributor" is a copyright holder who authorizes use under this
License of the Program or a work on which the Program is based.  The
work thus licensed is called the contributor's "contributor version".

  A contributor's "essential patent claims" are all patent claims
owned or controlled by the contributor, whether already acquired or
hereafter acquired, that would be infringed by some manner, permitted
by this License, of making, using, or selling its contributor version,
but do not include claims that would be infringed only as a
consequence of further modification of the contributor version.  For
purposes of this definition, "control" includes the right to grant
patent sublicenses in a manner consistent with the requirements of
this License.

  Each contributor grants you a non-exclusive, worldwide, royalty-free
patent license under the contributor's essential patent claims, to
make, use, sell, offer for sale, import and otherwise run, modify and
propagate the contents of its contributor version.

  In the following three paragraphs, a "patent license" is any express
agreement or commitment, however denominated, not to enforce a patent
(such as an express permission to practice a patent or covenant not to
sue for patent infringement).  To "grant" such a patent license to a
party means to make such an agreement or commitment not to enforce a
patent against the party.

  If you convey a covered work, knowingly relying on a patent license,
and the Corresponding Source of the work is not available for anyone
to copy, free of charge and under the terms of this License, through a
publicly available network server or other readily accessible means,
then you must either (1) cause the Corresponding Source to be so
available, or (2) arrange to deprive yourself of the benefit of the
patent license for this particular work, or (3) arrange, in a manner
consistent with the requirements of this License, to extend the patent
license to downstream recipients.  "Knowingly relying" means you have
actual knowledge that, but for the patent license, your conveying the
covered work in a country, or your recipient's use of the covered work
in a country, would infringe one or more identifiable patents in that
country that you have reason to believe are valid.

  If, pursuant to or in connection with a single transaction or
arrangement, you convey, or propagate by procuring conveyance of, a
covered work, and grant a patent license to some of the parties
```

(continues on next page)

(continued from previous page)

```
receiving the covered work authorizing them to use, propagate, modify
or convey a specific copy of the covered work, then the patent license
you grant is automatically extended to all recipients of the covered
work and works based on it.

  A patent license is "discriminatory" if it does not include within
the scope of its coverage, prohibits the exercise of, or is
conditioned on the non-exercise of one or more of the rights that are
specifically granted under this License.  You may not convey a covered
work if you are a party to an arrangement with a third party that is
in the business of distributing software, under which you make payment
to the third party based on the extent of your activity of conveying
the work, and under which the third party grants, to any of the
parties who would receive the covered work from you, a discriminatory
patent license (a) in connection with copies of the covered work
conveyed by you (or copies made from those copies), or (b) primarily
for and in connection with specific products or compilations that
contain the covered work, unless you entered into that arrangement,
or that patent license was granted, prior to 28 March 2007.

  Nothing in this License shall be construed as excluding or limiting
any implied license or other defenses to infringement that may
otherwise be available to you under applicable patent law.

  12. No Surrender of Others' Freedom.

  If conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot convey a
covered work so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you may
not convey it at all.  For example, if you agree to terms that obligate you
to collect a royalty for further conveying from those to whom you convey
the Program, the only way you could satisfy both those terms and this
License would be to refrain entirely from conveying the Program.

  13. Use with the GNU Affero General Public License.

  Notwithstanding any other provision of this License, you have
permission to link or combine any covered work with a work licensed
under version 3 of the GNU Affero General Public License into a single
combined work, and to convey the resulting work.  The terms of this
License will continue to apply to the part which is the covered work,
but the special requirements of the GNU Affero General Public License,
section 13, concerning interaction through a network will apply to the
combination as such.

  14. Revised Versions of this License.

  The Free Software Foundation may publish revised and/or new versions of
the GNU General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
```

(continues on next page)

(continued from previous page)

```
address new problems or concerns.

  Each version is given a distinguishing version number.  If the
Program specifies that a certain numbered version of the GNU General
Public License "or any later version" applies to it, you have the
option of following the terms and conditions either of that numbered
version or of any later version published by the Free Software
Foundation.  If the Program does not specify a version number of the
GNU General Public License, you may choose any version ever published
by the Free Software Foundation.

  If the Program specifies that a proxy can decide which future
versions of the GNU General Public License can be used, that proxy's
public statement of acceptance of a version permanently authorizes you
to choose that version for the Program.

  Later license versions may give you additional or different
permissions.  However, no additional obligations are imposed on any
author or copyright holder as a result of your choosing to follow a
later version.

  15. Disclaimer of Warranty.

  THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY
APPLICABLE LAW.  EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT
HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY
OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM
IS WITH YOU.  SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF
ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. Limitation of Liability.

  IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS
THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY
GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE
USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF
DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD
PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS),
EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF
SUCH DAMAGES.

  17. Interpretation of Sections 15 and 16.

  If the disclaimer of warranty and limitation of liability provided
above cannot be given local legal effect according to their terms,
reviewing courts shall apply local law that most closely approximates
an absolute waiver of all civil liability in connection with the
Program, unless a warranty or assumption of liability accompanies a
copy of the Program in return for a fee.
```

(continues on next page)

(continued from previous page)

```
                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
state the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software: you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation, either version 3 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program.  If not, see <http://www.gnu.org/licenses/>.

Also add information on how to contact you by electronic and paper mail.

  If the program does terminal interaction, make it output a short
notice like this when it starts in an interactive mode:

    <program>  Copyright (C) <year>  <name of author>
    This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, your program's commands
might be different; for a GUI interface, you would use an "about box".

  You should also get your employer (if you work as a programmer) or school,
if any, to sign a "copyright disclaimer" for the program, if necessary.
For more information on this, and how to apply and follow the GNU GPL, see
<http://www.gnu.org/licenses/>.

  The GNU General Public License does not permit incorporating your program
into proprietary programs.  If your program is a subroutine library, you
may consider it more useful to permit linking proprietary applications with
```

(continues on next page)

(continued from previous page)

```
the library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.  But first, please read
<http://www.gnu.org/philosophy/why-not-lgpl.html>.
```

40.51 LGPL3+ License Text for libunistring

The following license text applies to libunistring:

Listing 51: Download: LGPL3+ for libunistring

```
                   GNU LESSER GENERAL PUBLIC LICENSE
                       Version 3, 29 June 2007

 Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.


  This version of the GNU Lesser General Public License incorporates
the terms and conditions of version 3 of the GNU General Public
License, supplemented by the additional permissions listed below.

  0. Additional Definitions.

  As used herein, "this License" refers to version 3 of the GNU Lesser
General Public License, and the "GNU GPL" refers to version 3 of the GNU
General Public License.

  "The Library" refers to a covered work governed by this License,
other than an Application or a Combined Work as defined below.

  An "Application" is any work that makes use of an interface provided
by the Library, but which is not otherwise based on the Library.
Defining a subclass of a class defined by the Library is deemed a mode
of using an interface provided by the Library.

  A "Combined Work" is a work produced by combining or linking an
Application with the Library.  The particular version of the Library
with which the Combined Work was made is also called the "Linked
Version".

  The "Minimal Corresponding Source" for a Combined Work means the
Corresponding Source for the Combined Work, excluding any source code
for portions of the Combined Work that, considered in isolation, are
based on the Application, and not on the Linked Version.

  The "Corresponding Application Code" for a Combined Work means the
object code and/or source code for the Application, including any data
and utility programs needed for reproducing the Combined Work from the
Application, but excluding the System Libraries of the Combined Work.
```

(continues on next page)

(continued from previous page)

```
  1. Exception to Section 3 of the GNU GPL.

  You may convey a covered work under sections 3 and 4 of this License
without being bound by section 3 of the GNU GPL.

  2. Conveying Modified Versions.

  If you modify a copy of the Library, and, in your modifications, a
facility refers to a function or data to be supplied by an Application
that uses the facility (other than as an argument passed when the
facility is invoked), then you may convey a copy of the modified
version:

   a) under this License, provided that you make a good faith effort to
   ensure that, in the event an Application does not supply the
   function or data, the facility still operates, and performs
   whatever part of its purpose remains meaningful, or

   b) under the GNU GPL, with none of the additional permissions of
   this License applicable to that copy.

  3. Object Code Incorporating Material from Library Header Files.

  The object code form of an Application may incorporate material from
a header file that is part of the Library.  You may convey such object
code under terms of your choice, provided that, if the incorporated
material is not limited to numerical parameters, data structure
layouts and accessors, or small macros, inline functions and templates
(ten or fewer lines in length), you do both of the following:

   a) Give prominent notice with each copy of the object code that the
   Library is used in it and that the Library and its use are
   covered by this License.

   b) Accompany the object code with a copy of the GNU GPL and this license
   document.

  4. Combined Works.

  You may convey a Combined Work under terms of your choice that,
taken together, effectively do not restrict modification of the
portions of the Library contained in the Combined Work and reverse
engineering for debugging such modifications, if you also do each of
the following:

   a) Give prominent notice with each copy of the Combined Work that
   the Library is used in it and that the Library and its use are
   covered by this License.

   b) Accompany the Combined Work with a copy of the GNU GPL and this license
   document.
```

(continues on next page)

(continued from previous page)

```
   c) For a Combined Work that displays copyright notices during
   execution, include the copyright notice for the Library among
   these notices, as well as a reference directing the user to the
   copies of the GNU GPL and this license document.

   d) Do one of the following:

       0) Convey the Minimal Corresponding Source under the terms of this
       License, and the Corresponding Application Code in a form
       suitable for, and under terms that permit, the user to
       recombine or relink the Application with a modified version of
       the Linked Version to produce a modified Combined Work, in the
       manner specified by section 6 of the GNU GPL for conveying
       Corresponding Source.

       1) Use a suitable shared library mechanism for linking with the
       Library.  A suitable mechanism is one that (a) uses at run time
       a copy of the Library already present on the user's computer
       system, and (b) will operate properly with a modified version
       of the Library that is interface-compatible with the Linked
       Version.

   e) Provide Installation Information, but only if you would otherwise
   be required to provide such information under section 6 of the
   GNU GPL, and only to the extent that such information is
   necessary to install and execute a modified version of the
   Combined Work produced by recombining or relinking the
   Application with a modified version of the Linked Version. (If
   you use option 4d0, the Installation Information must accompany
   the Minimal Corresponding Source and Corresponding Application
   Code. If you use option 4d1, you must provide the Installation
   Information in the manner specified by section 6 of the GNU GPL
   for conveying Corresponding Source.)

  5. Combined Libraries.

  You may place library facilities that are a work based on the
Library side by side in a single library together with other library
facilities that are not Applications and are not covered by this
License, and convey such a combined library under terms of your
choice, if you do both of the following:

   a) Accompany the combined library with a copy of the same work based
   on the Library, uncombined with any other library facilities,
   conveyed under the terms of this License.

   b) Give prominent notice with the combined library that part of it
   is a work based on the Library, and explaining where to find the
   accompanying uncombined form of the same work.

  6. Revised Versions of the GNU Lesser General Public License.
```

(continues on next page)

(continued from previous page)

```
  The Free Software Foundation may publish revised and/or new versions
of the GNU Lesser General Public License from time to time. Such new
versions will be similar in spirit to the present version, but may
differ in detail to address new problems or concerns.

  Each version is given a distinguishing version number. If the
Library as you received it specifies that a certain numbered version
of the GNU Lesser General Public License "or any later version"
applies to it, you have the option of following the terms and
conditions either of that published version or of any later version
published by the Free Software Foundation. If the Library as you
received it does not specify a version number of the GNU Lesser
General Public License, you may choose any version of the GNU Lesser
General Public License ever published by the Free Software Foundation.

  If the Library as you received it specifies that a proxy can decide
whether future versions of the GNU Lesser General Public License shall
apply, that proxy's public statement of acceptance of any version is
permanent authorization for you to choose that version for the
Library.
```

40.52 NODE License Text for libuv

The following license text applies to libuv:

Listing 52: Download: NODE for libuv

```
libuv is licensed for use as follows:

====
Copyright (c) 2015-present libuv project contributors.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to
deal in the Software without restriction, including without limitation the
rights to use, copy, modify, merge, publish, distribute, sublicense, and/or
sell copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS
IN THE SOFTWARE.
```

(continues on next page)

(continued from previous page)

```
====

This license applies to parts of libuv originating from the
https://github.com/joyent/libuv repository:

====

Copyright Joyent, Inc. and other Node contributors. All rights reserved.
Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to
deal in the Software without restriction, including without limitation the
rights to use, copy, modify, merge, publish, distribute, sublicense, and/or
sell copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS
IN THE SOFTWARE.

====

This license applies to all parts of libuv that are not externally
maintained libraries.

The externally maintained libraries used by libuv are:

  - tree.h (from FreeBSD), copyright Niels Provos. Two clause BSD license.

  - inet_pton and inet_ntop implementations, contained in src/inet.c, are
    copyright the Internet Systems Consortium, Inc., and licensed under the ISC
    license.

  - stdint-msvc2008.h (from msinttypes), copyright Alexander Chemeris. Three
    clause BSD license.

  - pthread-fixes.c, copyright Google Inc. and Sony Mobile Communications AB.
    Three clause BSD license.

  - android-ifaddrs.h, android-ifaddrs.c, copyright Berkeley Software Design
    Inc, Kenneth MacKay and Emergya (Cloud4all, FP7/2007-2013, grant agreement
    n° 289016). Three clause BSD license.
```

40.53 MIT License Text for libxml2

The following license text applies to libxml2:

Listing 53: Download: MIT for libxml2

```
Except where otherwise noted in the source code (e.g. the files hash.c,
list.c and the trio files, which are covered by a similar licence but
with different Copyright notices) all the files are:

 Copyright (C) 1998-2012 Daniel Veillard.  All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is fur-
nished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FIT-
NESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN
THE SOFTWARE.
```

40.54 TRIO License Text for libxml2

The following license text applies to libxml2:

Listing 54: Download: TRIO for libxml2

```
(Following sentences are from trio.c of libxml2-2.9.4.tar.gz.)

Copyright (C) 1998 Bjorn Reese and Daniel Stenberg.

Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED
WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF
MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE AUTHORS AND
CONTRIBUTORS ACCEPT NO RESPONSIBILITY IN ANY CONCEIVABLE MANNER.
```

40.55 MIT License Text for libxslt

The following license text applies to libxslt:

Listing 55: Download: MIT for libxslt

```
Licence for libxslt except libexslt
-----------------------------------------------------------------------
 Copyright (C) 2001-2002 Daniel Veillard.  All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is fur-
nished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FIT-
NESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
DANIEL VEILLARD BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER
IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CON-
NECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Daniel Veillard shall not
be used in advertising or otherwise to promote the sale, use or other deal-
ings in this Software without prior written authorization from him.

-----------------------------------------------------------------------

Licence for libexslt
-----------------------------------------------------------------------
 Copyright (C) 2001-2002 Thomas Broyer, Charlie Bozeman and Daniel Veillard.
```

(continues on next page)

(continued from previous page)

```
 All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is fur-
nished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FIT-
NESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
AUTHORS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER
IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CON-
NECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of the authors shall not
be used in advertising or otherwise to promote the sale, use or other deal-
ings in this Software without prior written authorization from him.
------------------------------------------------------------------------
```

40.56 GPLv2 License Text for links

The following license text applies to links:

Listing 56: Download: GPLv2 for links

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.
     59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Library General Public License instead.)  You can apply it to
your programs, too.

```

(continues on next page)

(continued from previous page)

```
  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and

modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".
```

(continues on next page)

(continued from previous page)

```
Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on

    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.
```

(continues on next page)

(continued from previous page)

```
Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not

compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
```

(continues on next page)

(continued from previous page)

```
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.
```

(continues on next page)

(continued from previous page)

```
This section is intended to make thoroughly clear what is believed to

be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
```

(continues on next page)

(continued from previous page)

```
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.


                    END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program; if not, write to the Free Software
    Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA


Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year  name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
```

(continues on next page)

(continued from previous page)

```
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Library General
Public License instead of this License.

  In addition, as a special exception, the copyright holders give
  permission to link the code of portions of this program with the
  OpenSSL library under certain conditions as described in each
  individual source file, and distribute linked combinations
  including the two.
  You must obey the GNU General Public License in all respects
  for all of the code used other than OpenSSL.  If you modify
  file(s) with this exception, you may extend this exception to your
  version of the file(s), but you are not obligated to do so.  If you
  do not wish to do so, delete this exception statement from your
  version.  If you delete this exception statement from all source
  files in the program, then also delete it here.
```

40.57 MIT License Text for luajit-openresty

The following license text applies to luajit-openresty:

Listing 57: Download: MIT for luajit-openresty

```
==============================================================================
LuaJIT -- a Just-In-Time Compiler for Lua. https://luajit.org/

Copyright (C) 2005-2021 Mike Pall. All rights reserved.

Copyright (C) 2017-2018 Yichun Zhang. All rights reserved.

Copyright (C) 2017-2018 OpenResty Inc. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
```

(continues on next page)

(continued from previous page)

```
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN
THE SOFTWARE.

[ MIT license: https://www.opensource.org/licenses/mit-license.php ]

===============================================================================
[ LuaJIT includes code from Lua 5.1/5.2, which has this license statement: ]

Copyright (C) 1994-2012 Lua.org, PUC-Rio.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN
THE SOFTWARE.

===============================================================================
[ LuaJIT includes code from dlmalloc, which has this license statement: ]

This is a version (aka dlmalloc) of malloc/free/realloc written by
Doug Lea and released to the public domain, as explained at
https://creativecommons.org/licenses/publicdomain

===============================================================================
```

40.58 PD License Text for luajit-openresty

The following license text applies to luajit-openresty:

Listing 58: Download: PD for luajit-openresty

```
================================================================================
LuaJIT -- a Just-In-Time Compiler for Lua. https://luajit.org/

Copyright (C) 2005-2021 Mike Pall. All rights reserved.

Copyright (C) 2017-2018 Yichun Zhang. All rights reserved.

Copyright (C) 2017-2018 OpenResty Inc. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN
THE SOFTWARE.

[ MIT license: https://www.opensource.org/licenses/mit-license.php ]

================================================================================
[ LuaJIT includes code from Lua 5.1/5.2, which has this license statement: ]

Copyright (C) 1994-2012 Lua.org, PUC-Rio.

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.  IN NO EVENT SHALL THE
```

(continues on next page)

(continued from previous page)

```
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN
THE SOFTWARE.

================================================================================
[ LuaJIT includes code from dlmalloc, which has this license statement: ]

This is a version (aka dlmalloc) of malloc/free/realloc written by
Doug Lea and released to the public domain, as explained at
https://creativecommons.org/licenses/publicdomain

================================================================================
```

40.59 GPLv2 License Text for lzo2

The following license text applies to lzo2:

Listing 59: Download: GPLv2 for lzo2

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
```

(continues on next page)

(continued from previous page)

```
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
```

(continues on next page)

(continued from previous page)

```
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
```

(continues on next page)

(continued from previous page)

```
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
```

(continues on next page)

(continued from previous page)

```
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
```

(continues on next page)

(continued from previous page)

```
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
```

(continues on next page)

(continued from previous page)

```
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.60 PD License Text for mobile-broadband-provider-info

The following license text applies to mobile-broadband-provider-info:

Listing 60: Download: PD for mobile-broadband-provider-info

```
THIS WORK IS IN PUBLIC DOMAIN:
The person or persons who have associated work with this document
(the "Dedicator" or "Certifier") hereby either (a) certifies that, to the best
of his knowledge, the work of authorship identified is in the public domain of
the country from which the work is published, or (b) hereby dedicates whatever
copyright the dedicators holds in the work of authorship identified below
(the "Work") to the public domain. A certifier, moreover, dedicates any
copyright interest he may have in the associated work, and for these purposes,
is described as a "dedicator" below.

A certifier has taken reasonable steps to verify the copyright status of this
work. Certifier recognizes that his good faith efforts may not shield him from
liability if in fact the work certified is not in the public domain.

Dedicator makes this dedication for the benefit of the public at large and to
the detriment of the Dedicator's heirs and successors. Dedicator intends this
dedication to be an overt act of relinquishment in perpetuity of all present
and future rights under copyright law, whether vested or contingent, in the
Work. Dedicator understands that such relinquishment of all rights includes the
relinquishment of all rights to enforce (by lawsuit or otherwise) those
copyrights in the Work.

Dedicator recognizes that, once placed in the public domain, the Work may be
freely reproduced, distributed, transmitted, used, modified, built upon, or
otherwise exploited by anyone for any purpose, commercial or non-commercial,
and in any way, including by methods that have not yet been invented or
conceived.
```

40.61 BSD2CLAUSE License Text for mpdecimal

The following license text applies to mpdecimal:

Listing 61: Download: BSD2CLAUSE for mpdecimal

```
/*
 * Copyright (c) 2008-2020 Stefan Krah. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
```

(continues on next page)

(continued from previous page)

```
 *    notice, this list of conditions and the following disclaimer in the
 *    documentation and/or other materials provided with the distribution.
 *
 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 */
```

40.62 BSD2CLAUSE License Text for nginx

The following license text applies to nginx:

Listing 62: Download: BSD2CLAUSE for nginx

```
/*
 * Copyright (C) 2002-2021 Igor Sysoev
 * Copyright (C) 2011-2021 Nginx, Inc.
 * All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in the
 *    documentation and/or other materials provided with the distribution.
 *
 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 */
```

40.63 GPLv2 License Text for nss_ldap

The following license text applies to nss_ldap:

Listing 63: Download: GPLv2 for nss_ldap

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
```

(continues on next page)

(continued from previous page)

```
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
```

(continues on next page)

(continued from previous page)

```
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
```

(continues on next page)

(continued from previous page)

```
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
```

(continues on next page)

(continued from previous page)

```
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
```

(continues on next page)

(continued from previous page)

```
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
```

(continues on next page)

(continued from previous page)

```
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.64 BSD2CLAUSE License Text for ntp

The following license text applies to ntp:

Listing 64: Download: BSD2CLAUSE for ntp

```
This file is automatically generated from html/copyright.html
  Copyright Notice

   jpg "Clone me," says Dolly sheepishly.

   Last update: 4-Feb-2020 23:47 UTC
     ---------------------------------------------------------------------

   The following copyright notice applies to all files collectively called
   the Network Time Protocol Version 4 Distribution. Unless specifically
   declared otherwise in an individual file, this entire notice applies as
```

(continues on next page)

(continued from previous page)

```
   if the text was explicitly included in the file.
*********************************************************************
*                                                                   *
* Copyright (c) University of Delaware 1992-2015                    *
*                                                                   *
* Permission to use, copy, modify, and distribute this software and *
* its documentation for any purpose with or without fee is hereby   *
* granted, provided that the above copyright notice appears in all  *
* copies and that both the copyright notice and this permission     *
* notice appear in supporting documentation, and that the name      *
* University of Delaware not be used in advertising or publicity    *
* pertaining to distribution of the software without specific,      *
* written prior permission. The University of Delaware makes no     *
* representations about the suitability this software for any       *
* purpose. It is provided "as is" without express or implied        *
* warranty.                                                         *
*                                                                   *
*********************************************************************

   Content starting in 2011 from Harlan Stenn, Danny Mayer, and Martin
   Burnicki is:
*********************************************************************
*                                                                   *
* Copyright (c) Network Time Foundation 2011-2020                   *
*                                                                   *
* All Rights Reserved                                               *
*                                                                   *
* Redistribution and use in source and binary forms, with or without *
* modification, are permitted provided that the following conditions *
* are met:                                                          *
* 1. Redistributions of source code must retain the above copyright *
*    notice, this list of conditions and the following disclaimer.  *
* 2. Redistributions in binary form must reproduce the above        *
*    copyright notice, this list of conditions and the following    *
*    disclaimer in the documentation and/or other materials provided *
*    with the distribution.                                         *
*                                                                   *
* THIS SOFTWARE IS PROVIDED BY THE AUTHORS ``AS IS'' AND ANY EXPRESS *
* OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED  *
* WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE *
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE   *
* LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR *
* CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT  *
* OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR *
* BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF         *
* LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT          *
* (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE  *
* USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH   *
* DAMAGE.                                                           *
*********************************************************************

   The following individuals contributed in part to the Network Time
```

(continues on next page)

(continued from previous page)

```
Protocol Distribution Version 4 and are acknowledged as authors of this
work.
 1. [1]Takao Abe <takao_abe@xurb.jp> Clock driver for JJY receivers
 2. [2]Mark Andrews <mark_andrews@isc.org> Leitch atomic clock
    controller
 3. [3]Bernd Altmeier <altmeier@atlsoft.de> hopf Elektronik serial line
    and PCI-bus devices
 4. [4]Viraj Bais <vbais@mailman1.intel.com> and [5]Clayton Kirkwood
    <kirkwood@striderfm.intel.com> port to WindowsNT 3.5
 5. [6]Michael Barone <michael,barone@lmco.com> GPSVME fixes
 6. [7]Karl Berry <karl@owl.HQ.ileaf.com> syslog to file option
 7. [8]Greg Brackley <greg.brackley@bigfoot.com> Major rework of WINNT
    port. Clean up recvbuf and iosignal code into separate modules.
 8. [9]Marc Brett <Marc.Brett@westgeo.com> Magnavox GPS clock driver
 9. [10]Piete Brooks <Piete.Brooks@cl.cam.ac.uk> MSF clock driver,
    Trimble PARSE support
10. [11]Nelson B Bolyard <nelson@bolyard.me> update and complete
    broadcast and crypto features in sntp
11. [12]Jean-Francois Boudreault
    <Jean-Francois.Boudreault@viagenie.qc.ca> IPv6 support
12. [13]Reg Clemens <reg@dwf.com> Oncore driver (Current maintainer)
13. [14]Steve Clift <clift@ml.csiro.au> OMEGA clock driver
14. [15]Casey Crellin <casey@csc.co.za> vxWorks (Tornado) port and help
    with target configuration
15. [16]Sven Dietrich <sven_dietrich@trimble.com> Palisade reference
    clock driver, NT adj. residuals, integrated Greg's Winnt port.
16. [17]John A. Dundas III <dundas@salt.jpl.nasa.gov> Apple A/UX port
17. [18]Torsten Duwe <duwe@immd4.informatik.uni-erlangen.de> Linux port
18. [19]Dennis Ferguson <dennis@mrbill.canet.ca> foundation code for
    NTP Version 2 as specified in RFC-1119
19. [20]John Hay <jhay@icomtek.csir.co.za> IPv6 support and testing
20. [21]Dave Hart <davehart@davehart.com> General maintenance, Windows
    port interpolation rewrite
21. [22]Claas Hilbrecht <neoclock4x@linum.com> NeoClock4X clock driver
22. [23]Glenn Hollinger <glenn@herald.usask.ca> GOES clock driver
23. [24]Mike Iglesias <iglesias@uci.edu> DEC Alpha port
24. [25]Jim Jagielski <jim@jagubox.gsfc.nasa.gov> A/UX port
25. [26]Jeff Johnson <jbj@chatham.usdesign.com> massive prototyping
    overhaul
26. [27]Hans Lambermont <Hans.Lambermont@nl.origin-it.com> or
    [28]<H.Lambermont@chello.nl> ntpsweep
27. [29]Poul-Henning Kamp <phk@FreeBSD.ORG> Oncore driver (Original
    author)
28. [30]Frank Kardel [31]<kardel (at) ntp (dot) org> PARSE <GENERIC>
    (driver 14 reference clocks), STREAMS modules for PARSE, support
    scripts, syslog cleanup, dynamic interface handling
29. [32]Johannes Maximilian Kuehn <kuehn@ntp.org> Rewrote sntp to
    comply with NTPv4 specification, ntpq saveconfig
30. [33]William L. Jones <jones@hermes.chpc.utexas.edu> RS/6000 AIX
    modifications, HPUX modifications
31. [34]Dave Katz <dkatz@cisco.com> RS/6000 AIX port
32. [35]Craig Leres <leres@ee.lbl.gov> 4.4BSD port, ppsclock, Magnavox
```

(continues on next page)

(continued from previous page)

```
      GPS clock driver
   33. [36]George Lindholm <lindholm@ucs.ubc.ca> SunOS 5.1 port
   34. [37]Louis A. Mamakos <louie@ni.umd.edu> MD5-based authentication
   35. [38]Lars H. Mathiesen <thorinn@diku.dk> adaptation of foundation
       code for Version 3 as specified in RFC-1305
   36. [39]Danny Mayer <mayer@ntp.org>Network I/O, Windows Port, Code
       Maintenance
   37. [40]David L. Mills <mills@udel.edu> Version 4 foundation, precision
       kernel; clock drivers: 1, 3, 4, 6, 7, 11, 13, 18, 19, 22, 36
   38. [41]Wolfgang Moeller <moeller@gwdgv1.dnet.gwdg.de> VMS port
   39. [42]Jeffrey Mogul <mogul@pa.dec.com> ntptrace utility
   40. [43]Tom Moore <tmoore@fievel.daytonoh.ncr.com> i386 svr4 port
   41. [44]Kamal A Mostafa <kamal@whence.com> SCO OpenServer port
   42. [45]Derek Mulcahy <derek@toybox.demon.co.uk> and [46]Damon
       Hart-Davis <d@hd.org> ARCRON MSF clock driver
   43. [47]Rob Neal <neal@ntp.org> Bancomm refclock and config/parse code
       maintenance
   44. [48]Rainer Pruy <Rainer.Pruy@informatik.uni-erlangen.de>
       monitoring/trap scripts, statistics file handling
   45. [49]Dirce Richards <dirce@zk3.dec.com> Digital UNIX V4.0 port
   46. [50]Wilfredo Sánchez <wsanchez@apple.com> added support for NetInfo
   47. [51]Nick Sayer <mrapple@quack.kfu.com> SunOS streams modules
   48. [52]Jack Sasportas <jack@innovativeinternet.com> Saved a Lot of
       space on the stuff in the html/pic/ subdirectory
   49. [53]Ray Schnitzler <schnitz@unipress.com> Unixware1 port
   50. [54]Michael Shields <shields@tembel.org> USNO clock driver
   51. [55]Jeff Steinman <jss@pebbles.jpl.nasa.gov> Datum PTS clock driver
   52. [56]Harlan Stenn <harlan@pfcs.com> GNU automake/autoconfigure
       makeover, various other bits (see the ChangeLog)
   53. [57]Kenneth Stone <ken@sdd.hp.com> HP-UX port
   54. [58]Ajit Thyagarajan <ajit@ee.udel.edu>IP multicast/anycast support
   55. [59]Tomoaki TSURUOKA <tsuruoka@nc.fukuoka-u.ac.jp>TRAK clock driver
   56. [60]Brian Utterback <brian.utterback@oracle.com> General codebase,
       Solaris issues
   57. [61]Loganaden Velvindron <loganaden@gmail.com> Sandboxing
       (libseccomp) support
   58. [62]Paul A Vixie <vixie@vix.com> TrueTime GPS driver, generic
       TrueTime clock driver
   59. [63]Ulrich Windl <Ulrich.Windl@rz.uni-regensburg.de> corrected and
       validated HTML documents according to the HTML DTD
     ------------------------------------------------------------------------

References

   1. mailto:%20takao_abe@xurb.jp
   2. mailto:%20mark_andrews@isc.org
   3. mailto:%20altmeier@atlsoft.de
   4. mailto:%20vbais@mailman1.intel.co
   5. mailto:%20kirkwood@striderfm.intel.com
   6. mailto:%20michael.barone@lmco.com
   7. mailto:%20karl@owl.HQ.ileaf.com
   8. mailto:%20greg.brackley@bigfoot.com
```

(continues on next page)

(continued from previous page)

```
 9. mailto:%20Marc.Brett@westgeo.com
10. mailto:%20Piete.Brooks@cl.cam.ac.uk
11. mailto:%20nelson@bolyard.me
12. mailto:%20Jean-Francois.Boudreault@viagenie.qc.ca
13. mailto:%20reg@dwf.com
14. mailto:%20clift@ml.csiro.au
15. mailto:%20casey@csc.co.za
16. mailto:%20Sven_Dietrich@trimble.COM
17. mailto:%20dundas@salt.jpl.nasa.gov
18. mailto:%20duwe@immd4.informatik.uni-erlangen.de
19. mailto:%20dennis@mrbill.canet.ca
20. mailto:%20jhay@icomtek.csir.co.za
21. mailto:%20davehart@davehart.com
22. mailto:%20neoclock4x@linum.com
23. mailto:%20glenn@herald.usask.ca
24. mailto:%20iglesias@uci.edu
25. mailto:%20jagubox.gsfc.nasa.gov
26. mailto:%20jbj@chatham.usdesign.com
27. mailto:%20Hans.Lambermont@nl.origin-it.com
28. mailto:H.Lambermont@chello.nl
29. mailto:%20phk@FreeBSD.ORG
30. http://www4.informatik.uni-erlangen.de/%7ekardel
31. mailto:%20kardel%20%28at%29%20ntp%20%28dot%29%20org
32. mailto:kuehn@ntp.org
33. mailto:%20jones@hermes.chpc.utexas.edu
34. mailto:%20dkatz@cisco.com
35. mailto:%20leres@ee.lbl.gov
36. mailto:%20lindholm@ucs.ubc.ca
37. mailto:%20louie@ni.umd.edu
38. mailto:%20thorinn@diku.dk
39. mailto:%20mayer@ntp.org
40. mailto:%20mills@udel.edu
41. mailto:%20moeller@gwdgv1.dnet.gwdg.de
42. mailto:%20mogul@pa.dec.com
43. mailto:%20tmoore@fievel.daytonoh.ncr.com
44. mailto:%20kamal@whence.com
45. mailto:%20derek@toybox.demon.co.uk
46. mailto:%20d@hd.org
47. mailto:%20neal@ntp.org
48. mailto:%20Rainer.Pruy@informatik.uni-erlangen.de
49. mailto:%20dirce@zk3.dec.com
50. mailto:%20wsanchez@apple.com
51. mailto:%20mrapple@quack.kfu.com
52. mailto:%20jack@innovativeinternet.com
53. mailto:%20schnitz@unipress.com
54. mailto:%20shields@tembel.org
55. mailto:%20pebbles.jpl.nasa.gov
56. mailto:%20harlan@pfcs.com
57. mailto:%20ken@sdd.hp.com
58. mailto:%20ajit@ee.udel.edu
59. mailto:%20tsuruoka@nc.fukuoka-u.ac.jp
60. mailto:%20brian.utterback@oracle.com
```

(continues on next page)

(continued from previous page)

```
  61. mailto:%20loganaden@gmail.com
  62. mailto:%20vixie@vix.com
  63. mailto:%20Ulrich.Windl@rz.uni-regensburg.de
```

40.65 BSD2CLAUSE License Text for oniguruma

The following license text applies to oniguruma:

Listing 65: Download: BSD2CLAUSE for oniguruma

```
Oniguruma LICENSE
-----------------

Copyright (c) 2002-2021  K.Kosako  <kkosako0@gmail.com>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.
```

40.66 OPENLDAP License Text for openldap24-client

The following license text applies to openldap24-client:

Listing 66: Download: OPENLDAP for openldap24-client

```
The OpenLDAP Public License
  Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation
("Software"), with or without modification, are permitted provided
```

(continues on next page)

(continued from previous page)

```
that the following conditions are met:

1. Redistributions in source form must retain copyright statements
   and notices,

2. Redistributions in binary form must reproduce applicable copyright
   statements and notices, this list of conditions, and the following
   disclaimer in the documentation and/or other materials provided
   with the distribution, and

3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time.
Each revision is distinguished by a version number.  You may use
this Software under terms of this license revision or under the
terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS
CONTRIBUTORS ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES,
INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY
AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT
SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S)
OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN
ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in
advertising or otherwise to promote the sale, use or other dealing
in this Software without specific, written prior permission.  Title
to copyright in this Software shall at all times remain with copyright
holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City,
California, USA.  All Rights Reserved.  Permission to copy and
distribute verbatim copies of this document is granted.
```

40.67 LGPL21 License Text for opensc

The following license text applies to opensc:

Listing 67: Download: LGPL21 for opensc

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.
```

(continues on next page)

(continued from previous page)

```
  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be

introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it becomes
a de-facto standard.  To achieve this, non-free programs must be
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
```

(continues on next page)

(continued from previous page)

```
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must

be combined with the library in order to run.

                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.
```

(continues on next page)

(continued from previous page)

```
  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.
```

(continues on next page)

(continued from previous page)

```
Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.
```

(continues on next page)

(continued from previous page)

```
  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
```

(continues on next page)

(continued from previous page)

```
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you

distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
```

(continues on next page)

(continued from previous page)

```
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.
```

(continues on next page)

(continued from previous page)

```
  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded.  In such case, this License incorporates the limitation as if
written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
```

(continues on next page)

(continued from previous page)

```
DAMAGES.


                     END OF TERMS AND CONDITIONS

           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms of the
ordinary General Public License).

  To apply these terms, attach the following notices to the library.  It is
safest to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That's all there is to it!
```

40.68 APACHE20 License Text for openvpn-auth-script

The following license text applies to openvpn-auth-script:

Listing 68: Download: APACHE20 for openvpn-auth-script

```
                                 Apache License
                           Version 2.0, January 2004
                        http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

   "License" shall mean the terms and conditions for use, reproduction,
   and distribution as defined by Sections 1 through 9 of this document.

   "Licensor" shall mean the copyright owner or entity authorized by
   the copyright owner that is granting the License.

   "Legal Entity" shall mean the union of the acting entity and all
   other entities that control, are controlled by, or are under common
   control with that entity. For the purposes of this definition,
   "control" means (i) the power, direct or indirect, to cause the
   direction or management of such entity, whether by contract or
   otherwise, or (ii) ownership of fifty percent (50%) or more of the
   outstanding shares, or (iii) beneficial ownership of such entity.

   "You" (or "Your") shall mean an individual or Legal Entity
   exercising permissions granted by this License.

   "Source" form shall mean the preferred form for making modifications,
   including but not limited to software source code, documentation
   source, and configuration files.

   "Object" form shall mean any form resulting from mechanical
   transformation or translation of a Source form, including but
   not limited to compiled object code, generated documentation,
   and conversions to other media types.

   "Work" shall mean the work of authorship, whether in Source or
   Object form, made available under the License, as indicated by a
   copyright notice that is included in or attached to the work
   (an example is provided in the Appendix below).

   "Derivative Works" shall mean any work, whether in Source or Object
   form, that is based on (or derived from) the Work and for which the
   editorial revisions, annotations, elaborations, or other modifications
   represent, as a whole, an original work of authorship. For the purposes
   of this License, Derivative Works shall not include works that remain
   separable from, or merely link (or bind by name) to the interfaces of,
   the Work and Derivative Works thereof.
```

(continues on next page)

(continued from previous page)

```
   "Contribution" shall mean any work of authorship, including
   the original version of the Work and any modifications or additions
   to that Work or Derivative Works thereof, that is intentionally
   submitted to Licensor for inclusion in the Work by the copyright owner
   or by an individual or Legal Entity authorized to submit on behalf of
   the copyright owner. For the purposes of this definition, "submitted"
   means any form of electronic, verbal, or written communication sent
   to the Licensor or its representatives, including but not limited to
   communication on electronic mailing lists, source code control systems,
   and issue tracking systems that are managed by, or on behalf of, the
   Licensor for the purpose of discussing and improving the Work, but
   excluding communication that is conspicuously marked or otherwise
   designated in writing by the copyright owner as "Not a Contribution."

   "Contributor" shall mean Licensor and any individual or Legal Entity
   on behalf of whom a Contribution has been received by Licensor and
   subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   copyright license to reproduce, prepare Derivative Works of,
   publicly display, publicly perform, sublicense, and distribute the
   Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   (except as stated in this section) patent license to make, have made,
   use, offer to sell, sell, import, and otherwise transfer the Work,
   where such license applies only to those patent claims licensable
   by such Contributor that are necessarily infringed by their
   Contribution(s) alone or by combination of their Contribution(s)
   with the Work to which such Contribution(s) was submitted. If You
   institute patent litigation against any entity (including a
   cross-claim or counterclaim in a lawsuit) alleging that the Work
   or a Contribution incorporated within the Work constitutes direct
   or contributory patent infringement, then any patent licenses
   granted to You under this License for that Work shall terminate
   as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the
   Work or Derivative Works thereof in any medium, with or without
   modifications, and in Source or Object form, provided that You
   meet the following conditions:

   (a) You must give any other recipients of the Work or
       Derivative Works a copy of this License; and

   (b) You must cause any modified files to carry prominent notices
       stating that You changed the files; and
```

(continues on next page)

(continued from previous page)

```
   (c) You must retain, in the Source form of any Derivative Works
       that You distribute, all copyright, patent, trademark, and
       attribution notices from the Source form of the Work,
       excluding those notices that do not pertain to any part of
       the Derivative Works; and

   (d) If the Work includes a "NOTICE" text file as part of its
       distribution, then any Derivative Works that You distribute must
       include a readable copy of the attribution notices contained
       within such NOTICE file, excluding those notices that do not
       pertain to any part of the Derivative Works, in at least one
       of the following places: within a NOTICE text file distributed
       as part of the Derivative Works; within the Source form or
       documentation, if provided along with the Derivative Works; or,
       within a display generated by the Derivative Works, if and
       wherever such third-party notices normally appear. The contents
       of the NOTICE file are for informational purposes only and
       do not modify the License. You may add Your own attribution
       notices within Derivative Works that You distribute, alongside
       or as an addendum to the NOTICE text from the Work, provided
       that such additional attribution notices cannot be construed
       as modifying the License.

   You may add Your own copyright statement to Your modifications and
   may provide additional or different license terms and conditions
   for use, reproduction, or distribution of Your modifications, or
   for any such Derivative Works as a whole, provided Your use,
   reproduction, and distribution of the Work otherwise complies with
   the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise,
   any Contribution intentionally submitted for inclusion in the Work
   by You to the Licensor shall be under the terms and conditions of
   this License, without any additional terms or conditions.
   Notwithstanding the above, nothing herein shall supersede or modify
   the terms of any separate license agreement you may have executed
   with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade
   names, trademarks, service marks, or product names of the Licensor,
   except as required for reasonable and customary use in describing the
   origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or
   agreed to in writing, Licensor provides the Work (and each
   Contributor provides its Contributions) on an "AS IS" BASIS,
   WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
   implied, including, without limitation, any warranties or conditions
   of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A
   PARTICULAR PURPOSE. You are solely responsible for determining the
   appropriateness of using or redistributing the Work and assume any
   risks associated with Your exercise of permissions under this License.
```

(continues on next page)

(continued from previous page)

```
8. Limitation of Liability. In no event and under no legal theory,
   whether in tort (including negligence), contract, or otherwise,
   unless required by applicable law (such as deliberate and grossly
   negligent acts) or agreed to in writing, shall any Contributor be
   liable to You for damages, including any direct, indirect, special,
   incidental, or consequential damages of any character arising as a
   result of this License or out of the use or inability to use the
   Work (including but not limited to damages for loss of goodwill,
   work stoppage, computer failure or malfunction, or any and all
   other commercial damages or losses), even if such Contributor
   has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing
   the Work or Derivative Works thereof, You may choose to offer,
   and charge a fee for, acceptance of support, warranty, indemnity,
   or other liability obligations and/or rights consistent with this
   License. However, in accepting such obligations, You may act only
   on Your own behalf and on Your sole responsibility, not on behalf
   of any other Contributor, and only if You agree to indemnify,
   defend, and hold each Contributor harmless for any liability
   incurred by, or claims asserted against, such Contributor by reason
   of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

   To apply the Apache License to your work, attach the following
   boilerplate notice, with the fields enclosed by brackets "[]"
   replaced with your own identifying information. (Don't include
   the brackets!)  The text should be enclosed in the appropriate
   comment syntax for the file format. We also recommend that a
   file or class name and description of purpose be included on the
   same "printed page" as the copyright notice for easier
   identification within third-party archives.

Copyright 2017 FreeAgent Central Ltd

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

    http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```

40.69 GPLv2 License Text for openvpn

The following license text applies to openvpn:

Listing 69: Download: GPLv2 for openvpn

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
```

(continues on next page)

(continued from previous page)

```
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
```

(continues on next page)

(continued from previous page)

```
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
```

(continues on next page)

(continued from previous page)

```
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
```

(continues on next page)

(continued from previous page)

```
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
```

(continues on next page)

(continued from previous page)

```
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License version 2
    as published by the Free Software Foundation.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.
```

(continues on next page)

(continued from previous page)

```
Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.70 GPLv2+ License Text for pam_ldap

The following license text applies to pam_ldap:

Listing 70: Download: GPLv2+ for pam_ldap

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.
     59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
```

(continues on next page)

(continued from previous page)

```
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Library General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and

modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
```

(continues on next page)

(continued from previous page)

```
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on

    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
```

(continues on next page)

(continued from previous page)

```
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
```

(continues on next page)

(continued from previous page)

```
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not

compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
```

(continues on next page)

(continued from previous page)

```
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to

be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.
```

(continues on next page)

(continued from previous page)

```
  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.


                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program; if not, write to the Free Software
    Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA


Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year  name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.
```

(continues on next page)

(continued from previous page)

```
The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Library General
Public License instead of this License.
```

40.71 LGPL20+ License Text for pam_ldap

The following license text applies to pam_ldap:

Listing 71: Download: LGPL20+ for pam_ldap

```
                  GNU LIBRARY GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1991 Free Software Foundation, Inc.
 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the library GPL.  It is
 numbered 2 because it goes with version 2 of the ordinary GPL.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Library General Public License, applies to some
specially designated Free Software Foundation software, and to any
other libraries whose authors decide to use it.  You can use it for
your libraries, too.
```

(continues on next page)

(continued from previous page)

```
  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if
you distribute copies of the library, or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link a program with the library, you must provide
complete object files to the recipients so that they can relink them
with the library, after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  Our method of protecting your rights has two steps: (1) copyright
the library, and (2) offer you this license which gives you legal
permission to copy, distribute and/or modify the library.

  Also, for each distributor's protection, we want to make certain
that everyone understands that there is no warranty for this free
library.  If the library is modified by someone else and passed on, we
want its recipients to know that what they have is not the original
version, so that any problems introduced by others will not reflect on

the original authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that companies distributing free
software will individually obtain patent licenses, thus in effect
transforming the program into proprietary software.  To prevent this,
we have made it clear that any patent must be licensed for everyone's
free use or not licensed at all.

  Most GNU software, including some libraries, is covered by the ordinary
GNU General Public License, which was designed for utility programs.  This
license, the GNU Library General Public License, applies to certain
designated libraries.  This license is quite different from the ordinary
one; be sure to read it in full, and don't assume that anything in it is
the same as in the ordinary license.

  The reason we have a separate public license for some libraries is that
they blur the distinction we usually make between modifying or adding to a
program and simply using it.  Linking a program with a library, without
changing the library, is in some sense simply using the library, and is
analogous to running a utility program or application program.  However, in
```

(continues on next page)

(continued from previous page)

```
a textual and legal sense, the linked executable is a combined work, a
derivative of the original library, and the ordinary General Public License
treats it as such.

  Because of this blurred distinction, using the ordinary General
Public License for libraries did not effectively promote software
sharing, because most developers did not use the libraries.  We
concluded that weaker conditions might promote sharing better.

  However, unrestricted linking of non-free programs would deprive the
users of those programs of all benefit from the free status of the
libraries themselves.  This Library General Public License is intended to
permit developers of non-free programs to use free libraries, while
preserving your freedom as a user of such programs to change the free
libraries that are incorporated in them.  (We have not seen how to achieve
this as regards changes in header files, but we have achieved it as regards
changes in the actual functions of the Library.)  The hope is that this
will lead to faster development of free libraries.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, while the latter only
works together with the library.

  Note that it is possible for a library to be covered by the ordinary

General Public License rather than by this special one.

                  GNU LIBRARY GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library which
contains a notice placed by the copyright holder or other authorized
party saying it may be distributed under the terms of this Library
General Public License (also called "this License").  Each licensee is
addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
```

(continues on next page)

(continued from previous page)

```
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)
```

(continues on next page)

(continued from previous page)

```
These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in

these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.
```

(continues on next page)

(continued from previous page)

```
  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also compile or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
```

(continues on next page)

(continued from previous page)

```
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    c) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    d) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the source code distributed need not include anything that is normally
distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you

distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
```

(continues on next page)

(continued from previous page)

```
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to

this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.
```

(continues on next page)

(continued from previous page)

```
This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded.  In such case, this License incorporates the limitation as if
written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Library General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by

the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
```

(continues on next page)

(continued from previous page)

```
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.


                     END OF TERMS AND CONDITIONS

     Appendix: How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms of the
ordinary General Public License).

  To apply these terms, attach the following notices to the library.  It is
safest to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Library General Public
    License as published by the Free Software Foundation; either
    version 2 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Library General Public License for more details.

    You should have received a copy of the GNU Library General Public
    License along with this library; if not, write to the Free
    Software Foundation, Inc., 59 Temple Place - Suite 330, Boston,
    MA 02111-1307, USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice
```

(continues on next page)

(continued from previous page)

```
That's all there is to it!
```

40.72 BSD4CLAUSE License Text for pam_mkhomedir

The following license text applies to pam_mkhomedir:

Listing 72: Download: BSD4CLAUSE for pam_mkhomedir

```
Copyright (c) 1980, 1987, 1988, 1991, 1993, 1994
     The Regents of the University of California.  All rights reserved.
Copyright (c) 2001 Mark R V Murray
All rights reserved.
Copyright (c) 2001 Networks Associates Technology, Inc.
All rights reserved.
Copyright (c) 2004 Joe R. Doupnik
All rights reserved.
Copyright (c) 2005 Martin Mersberger
All rights reserved.

Portions of this software were developed for the FreeBSD Project by
ThinkSec AS and NAI Labs, the Security Research Division of Network
Associates, Inc.  under DARPA/SPAWAR contract N66001-01-C-8035
("CBOSS"), as part of the DARPA CHATS research program.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:
1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote
   products derived from this software without specific prior written
   permission.
4. Neither the name of the University nor the names of its contributors
   may be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.
```

40.73 BSD3CLAUSE License Text for pcre2

The following license text applies to pcre2:

Listing 73: Download: BSD3CLAUSE for pcre2

```
PCRE2 LICENCE
-------------

PCRE2 is a library of functions to support regular expressions whose syntax
and semantics are as close as possible to those of the Perl 5 language.

Releases 10.00 and above of PCRE2 are distributed under the terms of the "BSD"
licence, as specified below, with one exemption for certain binary
redistributions. The documentation for PCRE2, supplied in the "doc" directory,
is distributed under the same terms as the software itself. The data in the
testdata directory is not copyrighted and is in the public domain.

The basic library functions are written in C and are freestanding. Also
included in the distribution is a just-in-time compiler that can be used to
optimize pattern matching. This is an optional feature that can be omitted when
the library is built.


THE BASIC LIBRARY FUNCTIONS
---------------------------

Written by:       Philip Hazel
Email local part: Philip.Hazel
Email domain:     gmail.com

Retired from University of Cambridge Computing Service,
Cambridge, England.

Copyright (c) 1997-2021 University of Cambridge
All rights reserved.


PCRE2 JUST-IN-TIME COMPILATION SUPPORT
--------------------------------------

Written by:       Zoltan Herczeg
Email local part: hzmester
Email domain:     freemail.hu

Copyright(c) 2010-2021 Zoltan Herczeg
All rights reserved.


STACK-LESS JUST-IN-TIME COMPILER
--------------------------------

Written by:       Zoltan Herczeg
```

(continues on next page)

(continued from previous page)

```
Email local part: hzmester
Email domain:     freemail.hu

Copyright(c) 2009-2021 Zoltan Herczeg
All rights reserved.


THE "BSD" LICENCE
-----------------

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

    * Redistributions of source code must retain the above copyright notices,
      this list of conditions and the following disclaimer.

    * Redistributions in binary form must reproduce the above copyright
      notices, this list of conditions and the following disclaimer in the
      documentation and/or other materials provided with the distribution.

    * Neither the name of the University of Cambridge nor the names of any
      contributors may be used to endorse or promote products derived from this
      software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.


EXEMPTION FOR BINARY LIBRARY-LIKE PACKAGES
------------------------------------------

The second condition in the BSD licence (covering binary redistributions) does
not apply all the way down a chain of software. If binary package A includes
PCRE2, it must respect the condition, but if package B is software that
includes package A, the condition is not imposed on package B unless it uses
PCRE2 independently.

End
```

40.74 BSD3CLAUSE License Text for pcre

The following license text applies to pcre:

Listing 74: Download: BSD3CLAUSE for pcre

```
PCRE LICENCE
------------

PCRE is a library of functions to support regular expressions whose syntax
and semantics are as close as possible to those of the Perl 5 language.

Release 8 of PCRE is distributed under the terms of the "BSD" licence, as
specified below. The documentation for PCRE, supplied in the "doc"
directory, is distributed under the same terms as the software itself. The data
in the testdata directory is not copyrighted and is in the public domain.

The basic library functions are written in C and are freestanding. Also
included in the distribution is a set of C++ wrapper functions, and a
just-in-time compiler that can be used to optimize pattern matching. These
are both optional features that can be omitted when the library is built.


THE BASIC LIBRARY FUNCTIONS
---------------------------

Written by:       Philip Hazel
Email local part: Philip.Hazel
Email domain:     gmail.com

University of Cambridge Computing Service,
Cambridge, England.

Copyright (c) 1997-2021 University of Cambridge
All rights reserved.


PCRE JUST-IN-TIME COMPILATION SUPPORT
-------------------------------------

Written by:       Zoltan Herczeg
Email local part: hzmester
Email domain:     freemail.hu

Copyright(c) 2010-2021 Zoltan Herczeg
All rights reserved.


STACK-LESS JUST-IN-TIME COMPILER
--------------------------------

Written by:       Zoltan Herczeg
Email local part: hzmester
```

(continues on next page)

(continued from previous page)

```
Email domain:     freemail.hu

Copyright(c) 2009-2021 Zoltan Herczeg
All rights reserved.


THE C++ WRAPPER FUNCTIONS
-------------------------

Contributed by:   Google Inc.

Copyright (c) 2007-2012, Google Inc.
All rights reserved.


THE "BSD" LICENCE
-----------------

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

    * Redistributions of source code must retain the above copyright notice,
      this list of conditions and the following disclaimer.

    * Redistributions in binary form must reproduce the above copyright
      notice, this list of conditions and the following disclaimer in the
      documentation and/or other materials provided with the distribution.

    * Neither the name of the University of Cambridge nor the name of Google
      Inc. nor the names of their contributors may be used to endorse or
      promote products derived from this software without specific prior
      written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.

End
```

40.75 BSD3CLAUSE License Text for pcsc-lite

The following license text applies to pcsc-lite:

Listing 75: Download: BSD3CLAUSE for pcsc-lite

```
Copyright (c) 1999-2003 David Corcoran <corcoran@musclecard.com>
Copyright (c) 2001-2011 Ludovic Rousseau <ludovic.rousseau@free.fr>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products
   derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR
IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.


Some files are under GNU GPL v3 or any later version
- doc/example/pcsc_demo.c
- the files in src/spy/
- the files in UnitaryTests/

    Copyright (C) 2003-2014  Ludovic Rousseau

    This program is free software: you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation, either version 3 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program.  If not, see <http://www.gnu.org/licenses/>.
```

(continues on next page)

(continued from previous page)

```
Files src/auth.c and src/auth.h are:
 * Copyright (C) 2013 Red Hat
 *
 * All rights reserved.
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 * notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
 * notice, this list of conditions and the following disclaimer in the
 * documentation and/or other materials provided with the distribution.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
 * "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
 * LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
 * FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE
 * COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
 * INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
 * BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS
 * OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED
 * AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
 * OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF
 * THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH
 * DAMAGE.
 *
 * Author: Nikos Mavrogiannopoulos <nmav@redhat.com>


Files src/simclist.c and src/simclist.h are:
 * Copyright (c) 2007,2008,2009,2010,2011 Mij <mij@bitchx.it>
 *
 * Permission to use, copy, modify, and distribute this software for any
 * purpose with or without fee is hereby granted, provided that the above
 * copyright notice and this permission notice appear in all copies.
 *
 * THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES
 * WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
 * MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR
 * ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
 * WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
 * ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF
 * OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.
```

40.76 GPLv3+ License Text for pcsc-lite

The following license text applies to pcsc-lite:

Listing 76: Download: GPLv3+ for pcsc-lite

```
Copyright (c) 1999-2003 David Corcoran <corcoran@musclecard.com>
Copyright (c) 2001-2011 Ludovic Rousseau <ludovic.rousseau@free.fr>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products
   derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR
IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.


Some files are under GNU GPL v3 or any later version
- doc/example/pcsc_demo.c
- the files in src/spy/
- the files in UnitaryTests/

    Copyright (C) 2003-2014  Ludovic Rousseau

    This program is free software: you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation, either version 3 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program.  If not, see <http://www.gnu.org/licenses/>.
```

(continues on next page)

(continued from previous page)

```
Files src/auth.c and src/auth.h are:
 * Copyright (C) 2013 Red Hat
 *
 * All rights reserved.
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 * notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
 * notice, this list of conditions and the following disclaimer in the
 * documentation and/or other materials provided with the distribution.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
 * "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
 * LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
 * FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE
 * COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
 * INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
 * BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS
 * OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED
 * AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
 * OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF
 * THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH
 * DAMAGE.
 *
 * Author: Nikos Mavrogiannopoulos <nmav@redhat.com>


Files src/simclist.c and src/simclist.h are:
 * Copyright (c) 2007,2008,2009,2010,2011 Mij <mij@bitchx.it>
 *
 * Permission to use, copy, modify, and distribute this software for any
 * purpose with or without fee is hereby granted, provided that the above
 * copyright notice and this permission notice appear in all copies.
 *
 * THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES
 * WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
 * MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR
 * ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
 * WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
 * ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF
 * OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.
```

40.77 BSD3CLAUSE License Text for pear-Crypt_CHAP

The following license text applies to pear-Crypt_CHAP:

Listing 77: Download: BSD3CLAUSE for pear-Crypt_CHAP

```
Copyright (c) 2002-2010, Michael Bretterklieber <michael@bretterklieber.com>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products
   derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY
OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE,
EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This code cannot simply be copied and put under the GNU Public License or
any other GPL-like (LGPL, GPL2) License.
```

40.78 ART10 License Text for perl5

The following license text applies to perl5:

Listing 78: Download: ART10 for perl5

```
                        The "Artistic License"

                                Preamble

The intent of this document is to state the conditions under which a
Package may be copied, such that the Copyright Holder maintains some
semblance of artistic control over the development of the package,
while giving the users of the package the right to use and distribute
the Package in a more-or-less customary fashion, plus the right to make
reasonable modifications.
```

(continues on next page)

(continued from previous page)

```
Definitions:

        "Package" refers to the collection of files distributed by the
        Copyright Holder, and derivatives of that collection of files
        created through textual modification.

        "Standard Version" refers to such a Package if it has not been
        modified, or has been modified in accordance with the wishes
        of the Copyright Holder as specified below.

        "Copyright Holder" is whoever is named in the copyright or
        copyrights for the package.

        "You" is you, if you're thinking about copying or distributing
        this Package.

        "Reasonable copying fee" is whatever you can justify on the
        basis of media cost, duplication charges, time of people involved,
        and so on.  (You will not be required to justify it to the
        Copyright Holder, but only to the computing community at large
        as a market that must bear the fee.)

        "Freely Available" means that no fee is charged for the item
        itself, though there may be fees involved in handling the item.
        It also means that recipients of the item may redistribute it
        under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the
Standard Version of this Package without restriction, provided that you
duplicate all of the original copyright notices and associated disclaimers.

2. You may apply bug fixes, portability fixes and other modifications
derived from the Public Domain or from the Copyright Holder.  A Package
modified in such a way shall still be considered the Standard Version.

3. You may otherwise modify your copy of this Package in any way, provided
that you insert a prominent notice in each changed file stating how and
when you changed that file, and provided that you do at least ONE of the
following:

    a) place your modifications in the Public Domain or otherwise make them
    Freely Available, such as by posting said modifications to Usenet or
    an equivalent medium, or placing the modifications on a major archive
    site such as uunet.uu.net, or by allowing the Copyright Holder to include
    your modifications in the Standard Version of the Package.

    b) use the modified Package only within your corporation or organization.

    c) rename any non-standard executables so the names do not conflict
    with standard executables, which must also be provided, and provide
    a separate manual page for each non-standard executable that clearly
```

(continues on next page)

(continued from previous page)

```
    documents how it differs from the Standard Version.

    d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or
executable form, provided that you do at least ONE of the following:

    a) distribute a Standard Version of the executables and library files,
    together with instructions (in the manual page or equivalent) on where
    to get the Standard Version.

    b) accompany the distribution with the machine-readable source of
    the Package with your modifications.

    c) give non-standard executables non-standard names, and clearly
    document the differences in manual pages (or equivalent), together
    with instructions on where to get the Standard Version.

    d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this
Package.  You may charge any fee you choose for support of this
Package.  You may not charge a fee for this Package itself.  However,
you may distribute this Package in aggregate with other (possibly
commercial) programs as part of a larger (possibly commercial) software
distribution provided that you do not advertise this Package as a
product of your own.  You may embed this Package's interpreter within
an executable of yours (by linking); this shall be construed as a mere
form of aggregation, provided that the complete Standard Version of the
interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as
output from the programs of this Package do not automatically fall
under the copyright of this Package, but belong to whoever generated
them, and may be sold commercially, and may be aggregated with this
Package.  If such scripts or library files are aggregated with this
Package via the so-called "undump" or "unexec" methods of producing a
binary executable image, then distribution of such an image shall
neither be construed as a distribution of this Package nor shall it
fall under the restrictions of Paragraphs 3 and 4, provided that you do
not represent such an executable image as a Standard Version of this
Package.

7. C subroutines (or comparably compiled subroutines in other
languages) supplied by you and linked into this Package in order to
emulate subroutines and variables of the language defined by this
Package shall not be considered part of this Package, but are the
equivalent of input as in Paragraph 6, provided these subroutines do
not change the language in any way that would cause it to fail the
regression tests for the language.

8. Aggregation of this Package with a commercial distribution is always
```

(continues on next page)

(continued from previous page)

```
permitted provided that the use of this Package is embedded; that is,
when no overt attempt is made to make this Package's interfaces visible
to the end user of the commercial distribution.  Such use shall not be
construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote
products derived from this software without specific prior written permission.

10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR
IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

                                The End
```

40.79 GPLv1+ License Text for perl5

The following license text applies to perl5:

Listing 79: Download: GPLv1+ for perl5

```
                    GNU GENERAL PUBLIC LICENSE
                     Version 1, February 1989

 Copyright (C) 1989 Free Software Foundation, Inc.
                    51 Franklin St, Fifth Floor, Boston, MA  02110-1301  USA

 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The license agreements of most software companies try to keep users
at the mercy of those companies.  By contrast, our General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  The
General Public License applies to the Free Software Foundation's
software and to any other program whose authors commit to using it.
You can use it for your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Specifically, the General Public License is designed to make
sure that you have the freedom to give away or sell copies of free
software, that you receive source code or can get it if you want it,
that you can change the software or use pieces of it in new free
programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.
```

(continues on next page)

(continued from previous page)

```
  For example, if you distribute copies of a such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must tell them their rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  The precise terms and conditions for copying, distribution and

modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any program or other work which
contains a notice placed by the copyright holder saying it may be
distributed under the terms of this General Public License.  The
"Program", below, refers to any such program or work, and a "work based
on the Program" means either the Program or any work containing the
Program or a portion of it, either verbatim or with modifications.  Each
licensee is addressed as "you".

  1. You may copy and distribute verbatim copies of the Program's source
code as you receive it, in any medium, provided that you conspicuously and
appropriately publish on each copy an appropriate copyright notice and
disclaimer of warranty; keep intact all the notices that refer to this
General Public License and to the absence of any warranty; and give any
other recipients of the Program a copy of this General Public License
along with the Program.  You may charge a fee for the physical act of
transferring a copy.

  2. You may modify your copy or copies of the Program or any portion of
it, and copy and distribute such modifications under the terms of Paragraph
1 above, provided that you also do the following:

    a) cause the modified files to carry prominent notices stating that
    you changed the files and the date of any change; and

    b) cause the whole of any work that you distribute or publish, that
    in whole or in part contains the Program or any part thereof, either
    with or without modifications, to be licensed at no charge to all
    third parties under the terms of this General Public License (except
```

(continues on next page)

(continued from previous page)

```
    that you may choose to grant warranty protection to some or all
    third parties, at your option).

    c) If the modified program normally reads commands interactively when
    run, you must cause it, when started running for such interactive use
    in the simplest and most usual way, to print or display an
    announcement including an appropriate copyright notice and a notice
    that there is no warranty (or else, saying that you provide a
    warranty) and that users may redistribute the program under these
    conditions, and telling the user how to view a copy of this General
    Public License.

    d) You may charge a fee for the physical act of transferring a
    copy, and you may at your option offer warranty protection in
    exchange for a fee.

Mere aggregation of another independent work with the Program (or its
derivative) on a volume of a storage or distribution medium does not bring

the other work under the scope of these terms.

  3. You may copy and distribute the Program (or a portion or derivative of
it, under Paragraph 2) in object code or executable form under the terms of
Paragraphs 1 and 2 above provided that you also do one of the following:

    a) accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of
    Paragraphs 1 and 2 above; or,

    b) accompany it with a written offer, valid for at least three
    years, to give any third party free (except for a nominal charge
    for the cost of distribution) a complete machine-readable copy of the
    corresponding source code, to be distributed under the terms of
    Paragraphs 1 and 2 above; or,

    c) accompany it with the information you received as to where the
    corresponding source code may be obtained.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form alone.)

Source code for a work means the preferred form of the work for making
modifications to it.  For an executable file, complete source code means
all the source code for all modules it contains; but, as a special
exception, it need not include source code for modules which are standard
libraries that accompany the operating system on which the executable
file runs, or for standard header files or definitions files that
accompany that operating system.

  4. You may not copy, modify, sublicense, distribute or transfer the
Program except as expressly provided under this General Public License.
Any attempt otherwise to copy, modify, sublicense, distribute or transfer
the Program is void, and will automatically terminate your rights to use
```

(continues on next page)

(continued from previous page)

```
the Program under this License.  However, parties who have received
copies, or rights to use copies, from you under this General Public
License will not have their licenses terminated so long as such parties
remain in full compliance.

  5. By copying, distributing or modifying the Program (or any work based
on the Program) you indicate your acceptance of this license to do so,
and all its terms and conditions.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the original
licensor to copy, distribute or modify the Program subject to these
terms and conditions.  You may not impose any further restrictions on the

recipients' exercise of the rights granted herein.

  7. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of the license which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
the license, you may choose any version ever published by the Free Software
Foundation.

  8. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  9. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  10. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
```

(continues on next page)

(continued from previous page)

```
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.


                     END OF TERMS AND CONDITIONS

        Appendix: How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to humanity, the best way to achieve this is to make it
free software which everyone can redistribute and change under these
terms.

  To do so, attach the following notices to the program.  It is safest to
attach them to the start of each source file to most effectively convey
the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) 19yy  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 1, or (at your option)
    any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston MA  02110-1301 USA


Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) 19xx name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the
appropriate parts of the General Public License.  Of course, the
```

(continues on next page)

(continued from previous page)

```
commands you use may be called something other than `show w' and `show
c'; they could even be mouse-clicks or menu items--whatever suits your
program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  program `Gnomovision' (a program to direct compilers to make passes
  at assemblers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

That's all there is to it!
```

40.80 PHP301 License Text for php74-bcmath

The following license text applies to php74-bcmath:

Listing 80: Download: PHP301 for php74-bcmath

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"
```

(continues on next page)

(continued from previous page)

```
  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.81 PHP301 License Text for php74-bz2

The following license text applies to php74-bz2:

Listing 81: Download: PHP301 for php74-bz2

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.82 PHP301 License Text for php74-ctype

The following license text applies to php74-ctype:

Listing 82: Download: PHP301 for php74-ctype

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.83 PHP301 License Text for php74-curl

The following license text applies to php74-curl:

Listing 83: Download: PHP301 for php74-curl

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.84 PHP301 License Text for php74-dom

The following license text applies to php74-dom:

Listing 84: Download: PHP301 for php74-dom

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.85 PHP301 License Text for php74-filter

The following license text applies to php74-filter:

Listing 85: Download: PHP301 for php74-filter

```
--------------------------------------------------------------------
                   The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.86 PHP301 License Text for php74-gettext

The following license text applies to php74-gettext:

Listing 86: Download: PHP301 for php74-gettext

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.87 PHP301 License Text for php74-intl

The following license text applies to php74-intl:

Listing 87: Download: PHP301 for php74-intl

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.88 PHP301 License Text for php74-json

The following license text applies to php74-json:

Listing 88: Download: PHP301 for php74-json

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.89 PHP301 License Text for php74-ldap

The following license text applies to php74-ldap:

Listing 89: Download: PHP301 for php74-ldap

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.90 PHP301 License Text for php74-mbstring

The following license text applies to php74-mbstring:

Listing 90: Download: PHP301 for php74-mbstring

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.91 PHP301 License Text for php74-opcache

The following license text applies to php74-opcache:

Listing 91: Download: PHP301 for php74-opcache

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.92 PHP301 License Text for php74-openssl

The following license text applies to php74-openssl:

Listing 92: Download: PHP301 for php74-openssl

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.93 MIT License Text for php74-openssl_x509_crl

The following license text applies to php74-openssl_x509_crl:

Listing 93: Download: MIT for php74-openssl_x509_crl

```
The MIT License (MIT)

Copyright (c) 2015 ukrbublik

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to deal
in the Software without restriction, including without limitation the rights
to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all
copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE
SOFTWARE.
```

40.94 PHP301 License Text for php74-pcntl

The following license text applies to php74-pcntl:

Listing 94: Download: PHP301 for php74-pcntl

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.
```

(continues on next page)

(continued from previous page)

```
  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.95 PHP301 License Text for php74-pdo

The following license text applies to php74-pdo:

Listing 95: Download: PHP301 for php74-pdo

```
--------------------------------------------------------------------
                   The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.96 PHP301 License Text for php74-pdo_sqlite

The following license text applies to php74-pdo_sqlite:

Listing 96: Download: PHP301 for php74-pdo_sqlite

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.97 LGPL21 License Text for php74-pear-Cache_Lite

The following license text applies to php74-pear-Cache_Lite:

Listing 97: Download: LGPL21 for php74-pear-Cache_Lite

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.
```

(continues on next page)

(continued from previous page)

```
  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be
introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it becomes
a de-facto standard.  To achieve this, non-free programs must be
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
```

(continues on next page)

(continued from previous page)

```
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must
be combined with the library in order to run.

                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.

  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
```

(continues on next page)

(continued from previous page)

```
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a
fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
```

(continues on next page)

(continued from previous page)

```
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in
these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.

  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
```

(continues on next page)

(continued from previous page)

```
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,
whether or not they are linked directly with the Library itself.

  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
```

(continues on next page)

(continued from previous page)

```
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you
distribute.

  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.

  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
```

(continues on next page)

(continued from previous page)

```
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with
this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded.  In such case, this License incorporates the limitation as if
written in the body of this License.
```

(continues on next page)

(continued from previous page)

```
  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by
the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.

                     END OF TERMS AND CONDITIONS
```

40.98 BSD3CLAUSE License Text for php74-pear-HTTP_Request2

The following license text applies to php74-pear-HTTP_Request2:

Listing 98: Download: BSD3CLAUSE for php74-pear-HTTP_Request2

```
HTTP_Request2

Copyright (c) 2008-2021, Alexey Borzov <avb@php.net>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in the
     documentation and/or other materials provided with the distribution.

  3. Neither the name of Alexey Borzov nor the names of his contributors
     may be used to endorse or promote products derived from this software
     without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF
THE POSSIBILITY OF SUCH DAMAGE.
```

40.99 BSD2CLAUSE License Text for php74-pear-Net_SMTP

The following license text applies to php74-pear-Net_SMTP:

Listing 99: Download: BSD2CLAUSE for php74-pear-Net_SMTP

```
Copyright 2002-2017 Jon Parise and Chuck Hagenbuch.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this
```

(continues on next page)

(continued from previous page)

```
   list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice,
   this list of conditions and the following disclaimer in the documentation
   and/or other materials provided with the distribution..

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.100 BSD2CLAUSE License Text for php74-pear-Net_Socket

The following license text applies to php74-pear-Net_Socket:

Listing 100: Download: BSD2CLAUSE for php74-pear-Net_Socket

```
Copyright 1997-2017 The PHP Group

Redistribution and use in source and binary forms, with or without modification, are␣
↪permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of␣
↪conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list␣
↪of conditions and the following disclaimer in the documentation and/or other materials␣
↪provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY␣
↪EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES␣
↪OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT␣
↪SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,␣
↪INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED␣
↪TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR␣
↪BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN␣
↪CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY␣
↪WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.101 BSD3CLAUSE License Text for php74-pear-Net_URL2

The following license text applies to php74-pear-Net_URL2:

Listing 101: Download: BSD3CLAUSE for php74-pear-Net_URL2

```
Copyright (c) 2002-2003, Richard Heyes
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

1) Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.
2) Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
3) Neither the name of the Richard Heyes nor the names of his
   contributors may be used to endorse or promote products derived
   from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.102 PHP301 License Text for php74-pear-XML_RPC2

The following license text applies to php74-pear-XML_RPC2:

Listing 102: Download: PHP301 for php74-pear-XML_RPC2

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.
```

(continues on next page)

(continued from previous page)

```
  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.
```

(continues on next page)

(continued from previous page)

```
For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.103 PHP301 License Text for php74-pear

The following license text applies to php74-pear:

Listing 103: Download: PHP301 for php74-pear

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.
```

(continues on next page)

(continued from previous page)

```
  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.104 PHP301 License Text for php74-pecl-mcrypt

The following license text applies to php74-pecl-mcrypt:

Listing 104: Download: PHP301 for php74-pecl-mcrypt

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.105 PHP301 License Text for php74-pecl-rrd

The following license text applies to php74-pecl-rrd:

Listing 105: Download: PHP301 for php74-pecl-rrd

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.106 MIT License Text for php74-phpseclib

The following license text applies to php74-phpseclib:

Listing 106: Download: MIT for php74-phpseclib

```
Copyright 2007-2016 TerraFrost and other contributors
http://phpseclib.sourceforge.net/

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
"Software"), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject to
the following conditions:

The above copyright notice and this permission notice shall be
included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND
NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE
LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION
OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION
WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

40.107 PHP301 License Text for php74-posix

The following license text applies to php74-posix:

Listing 107: Download: PHP301 for php74-posix

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.
```

(continues on next page)

(continued from previous page)

```
  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.108 PHP301 License Text for php74-readline

The following license text applies to php74-readline:

Listing 108: Download: PHP301 for php74-readline

```
--------------------------------------------------------------------
                   The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.109 PHP301 License Text for php74-session

The following license text applies to php74-session:

Listing 109: Download: PHP301 for php74-session

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
```

(continues on next page)

(continued from previous page)

```
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.110 PHP301 License Text for php74-shmop

The following license text applies to php74-shmop:

Listing 110: Download: PHP301 for php74-shmop

```
--------------------------------------------------------------------
                   The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
```

(continues on next page)

(continued from previous page)

```
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.111 BSD3CLAUSE License Text for php74-simplepie

The following license text applies to php74-simplepie:

Listing 111: Download: BSD3CLAUSE for php74-simplepie

```
Copyright (c) 2004-2007, Ryan Parman and Geoffrey Sneddon.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are
permitted provided that the following conditions are met:

        * Redistributions of source code must retain the above copyright notice, this␣
↪list of
          conditions and the following disclaimer.

        * Redistributions in binary form must reproduce the above copyright notice, this␣
↪list
          of conditions and the following disclaimer in the documentation and/or other␣
↪materials
          provided with the distribution.

        * Neither the name of the SimplePie Team nor the names of its contributors may␣
↪be used
          to endorse or promote products derived from this software without specific␣
↪prior
          written permission.
```

(continues on next page)

(continued from previous page)

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY␣
↪EXPRESS
OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF␣
↪MERCHANTABILITY
AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT␣
↪HOLDERS
AND CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON␣
↪ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING␣
↪NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.
```

40.112 PHP301 License Text for php74-simplexml

The following license text applies to php74-simplexml:

Listing 112: Download: PHP301 for php74-simplexml

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
```

(continues on next page)

(continued from previous page)

```
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.113 PHP301 License Text for php74-sockets

The following license text applies to php74-sockets:

Listing 113: Download: PHP301 for php74-sockets

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------
```

(continues on next page)

(continued from previous page)

```
Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.
```

(continues on next page)

(continued from previous page)

```
--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.114 PHP301 License Text for php74-sqlite3

The following license text applies to php74-sqlite3:

Listing 114: Download: PHP301 for php74-sqlite3

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
```

(continues on next page)

(continued from previous page)

```
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.115 PHP301 License Text for php74-sysvmsg

The following license text applies to php74-sysvmsg:

Listing 115: Download: PHP301 for php74-sysvmsg

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
```

(continues on next page)

(continued from previous page)

```
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------
```

(continues on next page)

(continued from previous page)

```
This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.116 PHP301 License Text for php74-sysvsem

The following license text applies to php74-sysvsem:

Listing 116: Download: PHP301 for php74-sysvsem

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
```

(continues on next page)

(continued from previous page)

```
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.117 PHP301 License Text for php74-sysvshm

The following license text applies to php74-sysvshm:

Listing 117: Download: PHP301 for php74-sysvshm

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:
```

(continues on next page)

(continued from previous page)

```
  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.
```

(continues on next page)

(continued from previous page)

```
The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.118 PHP301 License Text for php74-tokenizer

The following license text applies to php74-tokenizer:

Listing 118: Download: PHP301 for php74-tokenizer

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
```

(continues on next page)

(continued from previous page)

```
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.119 PHP301 License Text for php74-xml

The following license text applies to php74-xml:

Listing 119: Download: PHP301 for php74-xml

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
```

(continues on next page)

(continued from previous page)

```
    notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.
```

(continues on next page)

(continued from previous page)

```
For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.120 PHP301 License Text for php74-xmlreader

The following license text applies to php74-xmlreader:

Listing 120: Download: PHP301 for php74-xmlreader

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.
```

(continues on next page)

(continued from previous page)

```
  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.121 PHP301 License Text for php74-xmlwriter

The following license text applies to php74-xmlwriter:

Listing 121: Download: PHP301 for php74-xmlwriter

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.
```

(continues on next page)

(continued from previous page)

```
  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
```

(continues on next page)

(continued from previous page)

```
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.122 PHP301 License Text for php74-zlib

The following license text applies to php74-zlib:

Listing 122: Download: PHP301 for php74-zlib

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
```

(continues on next page)

(continued from previous page)

```
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.123 PHP301 License Text for php74

The following license text applies to php74:

Listing 123: Download: PHP301 for php74

```
--------------------------------------------------------------------
                  The PHP License, version 3.01
Copyright (c) 1999 - 2010 The PHP Group. All rights reserved.
--------------------------------------------------------------------

Redistribution and use in source and binary forms, with or without
modification, is permitted provided that the following conditions
are met:

  1. Redistributions of source code must retain the above copyright
     notice, this list of conditions and the following disclaimer.

  2. Redistributions in binary form must reproduce the above copyright
     notice, this list of conditions and the following disclaimer in
```

(continues on next page)

(continued from previous page)

```
     the documentation and/or other materials provided with the
     distribution.

  3. The name "PHP" must not be used to endorse or promote products
     derived from this software without prior written permission. For
     written permission, please contact group@php.net.

  4. Products derived from this software may not be called "PHP", nor
     may "PHP" appear in their name, without prior written permission
     from group@php.net.  You may indicate that your software works in
     conjunction with PHP by saying "Foo for PHP" instead of calling
     it "PHP Foo" or "phpfoo"

  5. The PHP Group may publish revised and/or new versions of the
     license from time to time. Each version will be given a
     distinguishing version number.
     Once covered code has been published under a particular version
     of the license, you may always continue to use it under the terms
     of that version. You may also choose to use such covered code
     under the terms of any subsequent version of the license
     published by the PHP Group. No one other than the PHP Group has
     the right to modify the terms applicable to covered code created
     under this License.

  6. Redistributions of any form whatsoever must retain the following
     acknowledgment:
     "This product includes PHP software, freely available from
     <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND
ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE PHP
DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT,
INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.

--------------------------------------------------------------------

This software consists of voluntary contributions made by many
individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project,
please see <http://www.php.net>.
```

(continues on next page)

(continued from previous page)

```
PHP includes the Zend Engine, freely available at
<http://www.zend.com>.
```

40.124 MIT License Text for py38-setuptools

The following license text applies to py38-setuptools:

Listing 124: Download: MIT for py38-setuptools

```
Copyright Jason R. Coombs

Permission is hereby granted, free of charge, to any person obtaining a copy
of this software and associated documentation files (the "Software"), to
deal in the Software without restriction, including without limitation the
rights to use, copy, modify, merge, publish, distribute, sublicense, and/or
sell copies of the Software, and to permit persons to whom the Software is
furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS
IN THE SOFTWARE.
```

40.125 RADVD License Text for radvd

The following license text applies to radvd:

Listing 125: Download: RADVD for radvd

```
  The author(s) grant permission for redistribution and use in source and
binary forms, with or without modification, of the software and documentation
provided that the following conditions are met:

0. If you receive a version of the software that is specifically labelled
   as not being for redistribution (check the version message and/or README),
   you are not permitted to redistribute that version of the software in any
   way or form.
1. All terms of all other applicable copyrights and licenses must be
   followed.
2. Redistributions of source code must retain the authors' copyright
   notice(s), this list of conditions, and the following disclaimer.
3. Redistributions in binary form must reproduce the authors' copyright
```

(continues on next page)

(continued from previous page)

```
   notice(s), this list of conditions, and the following disclaimer in the
   documentation and/or other materials provided with the distribution.
4. All advertising materials mentioning features or use of this software
   must display the following acknowledgement with the name(s) of the
   authors as specified in the copyright notice(s) substituted where
   indicated:

        This product includes software developed by the authors which are
        mentioned at the start of the source files and other contributors.

5. Neither the name(s) of the author(s) nor the names of its contributors
   may be used to endorse or promote products derived from this software
   without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY ITS AUTHORS AND CONTRIBUTORS ``AS IS'' AND ANY
EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY
DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON
ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.126 GPLv2 License Text for rate

The following license text applies to rate:

Listing 126: Download: GPLv2 for rate

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.
                          675 Mass Ave, Cambridge, MA 02139, USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Library General Public License instead.)  You can apply it to
your programs, too.
```

(continues on next page)

(continued from previous page)

```
  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and

modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".
```

(continues on next page)

(continued from previous page)

```
Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on

    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
```

(continues on next page)

(continued from previous page)

```
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not

compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
```

(continues on next page)

(continued from previous page)

```
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.
```

(continues on next page)

(continued from previous page)

```
This section is intended to make thoroughly clear what is believed to

be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
```

(continues on next page)

(continued from previous page)

```
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.
```

40.127 GPLv3 License Text for readline

The following license text applies to readline:

Listing 127: Download: GPLv3 for readline

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 3, 29 June 2007

 Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The GNU General Public License is a free, copyleft license for
software and other kinds of works.

  The licenses for most software and other practical works are designed
to take away your freedom to share and change the works.  By contrast,
the GNU General Public License is intended to guarantee your freedom to
share and change all versions of a program--to make sure it remains free
software for all its users.  We, the Free Software Foundation, use the
GNU General Public License for most of our software; it applies also to
any other work released this way by its authors.  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
them if you wish), that you receive source code or can get it if you
want it, that you can change the software or use pieces of it in new
free programs, and that you know you can do these things.

  To protect your rights, we need to prevent others from denying you
these rights or asking you to surrender the rights.  Therefore, you have
certain responsibilities if you distribute copies of the software, or if
you modify it: responsibilities to respect the freedom of others.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must pass on to the recipients the same
freedoms that you received.  You must make sure that they, too, receive
or can get the source code.  And you must show them these terms so they
know their rights.

  Developers that use the GNU GPL protect your rights with two steps:
```

(continues on next page)

(continued from previous page)

```
(1) assert copyright on the software, and (2) offer you this License
giving you legal permission to copy, distribute and/or modify it.

  For the developers' and authors' protection, the GPL clearly explains
that there is no warranty for this free software.  For both users' and
authors' sake, the GPL requires that modified versions be marked as
changed, so that their problems will not be attributed erroneously to
authors of previous versions.

  Some devices are designed to deny users access to install or run
modified versions of the software inside them, although the manufacturer
can do so.  This is fundamentally incompatible with the aim of
protecting users' freedom to change the software.  The systematic
pattern of such abuse occurs in the area of products for individuals to
use, which is precisely where it is most unacceptable.  Therefore, we
have designed this version of the GPL to prohibit the practice for those
products.  If such problems arise substantially in other domains, we
stand ready to extend this provision to those domains in future versions
of the GPL, as needed to protect the freedom of users.

  Finally, every program is threatened constantly by software patents.
States should not allow patents to restrict development and use of
software on general-purpose computers, but in those that do, we wish to
avoid the special danger that patents applied to a free program could
make it effectively proprietary.  To prevent this, the GPL assures that
patents cannot be used to render the program non-free.

  The precise terms and conditions for copying, distribution and
modification follow.

                       TERMS AND CONDITIONS

  0. Definitions.

  "This License" refers to version 3 of the GNU General Public License.

  "Copyright" also means copyright-like laws that apply to other kinds of
works, such as semiconductor masks.

  "The Program" refers to any copyrightable work licensed under this
License.  Each licensee is addressed as "you".  "Licensees" and
"recipients" may be individuals or organizations.

  To "modify" a work means to copy from or adapt all or part of the work
in a fashion requiring copyright permission, other than the making of an
exact copy.  The resulting work is called a "modified version" of the
earlier work or a work "based on" the earlier work.

  A "covered work" means either the unmodified Program or a work based
on the Program.

  To "propagate" a work means to do anything with it that, without
```

(continues on next page)

(continued from previous page)

```
permission, would make you directly or secondarily liable for
infringement under applicable copyright law, except executing it on a
computer or modifying a private copy.  Propagation includes copying,
distribution (with or without modification), making available to the
public, and in some countries other activities as well.

  To "convey" a work means any kind of propagation that enables other
parties to make or receive copies.  Mere interaction with a user through
a computer network, with no transfer of a copy, is not conveying.

  An interactive user interface displays "Appropriate Legal Notices"
to the extent that it includes a convenient and prominently visible
feature that (1) displays an appropriate copyright notice, and (2)
tells the user that there is no warranty for the work (except to the
extent that warranties are provided), that licensees may convey the
work under this License, and how to view a copy of this License.  If
the interface presents a list of user commands or options, such as a
menu, a prominent item in the list meets this criterion.

  1. Source Code.

  The "source code" for a work means the preferred form of the work
for making modifications to it.  "Object code" means any non-source
form of a work.

  A "Standard Interface" means an interface that either is an official
standard defined by a recognized standards body, or, in the case of
interfaces specified for a particular programming language, one that
is widely used among developers working in that language.

  The "System Libraries" of an executable work include anything, other
than the work as a whole, that (a) is included in the normal form of
packaging a Major Component, but which is not part of that Major
Component, and (b) serves only to enable use of the work with that
Major Component, or to implement a Standard Interface for which an
implementation is available to the public in source code form.  A
"Major Component", in this context, means a major essential component
(kernel, window system, and so on) of the specific operating system
(if any) on which the executable work runs, or a compiler used to
produce the work, or an object code interpreter used to run it.

  The "Corresponding Source" for a work in object code form means all
the source code needed to generate, install, and (for an executable
work) run the object code and to modify the work, including scripts to
control those activities.  However, it does not include the work's
System Libraries, or general-purpose tools or generally available free
programs which are used unmodified in performing those activities but
which are not part of the work.  For example, Corresponding Source
includes interface definition files associated with source files for
the work, and the source code for shared libraries and dynamically
linked subprograms that the work is specifically designed to require,
such as by intimate data communication or control flow between those
```

(continues on next page)

(continued from previous page)

```
subprograms and other parts of the work.

  The Corresponding Source need not include anything that users
can regenerate automatically from other parts of the Corresponding
Source.

  The Corresponding Source for a work in source code form is that
same work.

  2. Basic Permissions.

  All rights granted under this License are granted for the term of
copyright on the Program, and are irrevocable provided the stated
conditions are met.  This License explicitly affirms your unlimited
permission to run the unmodified Program.  The output from running a
covered work is covered by this License only if the output, given its
content, constitutes a covered work.  This License acknowledges your
rights of fair use or other equivalent, as provided by copyright law.

  You may make, run and propagate covered works that you do not
convey, without conditions so long as your license otherwise remains
in force.  You may convey covered works to others for the sole purpose
of having them make modifications exclusively for you, or provide you
with facilities for running those works, provided that you comply with
the terms of this License in conveying all material for which you do
not control copyright.  Those thus making or running the covered works
for you must do so exclusively on your behalf, under your direction
and control, on terms that prohibit them from making any copies of
your copyrighted material outside their relationship with you.

  Conveying under any other circumstances is permitted solely under
the conditions stated below.  Sublicensing is not allowed; section 10
makes it unnecessary.

  3. Protecting Users' Legal Rights From Anti-Circumvention Law.

  No covered work shall be deemed part of an effective technological
measure under any applicable law fulfilling obligations under article
11 of the WIPO copyright treaty adopted on 20 December 1996, or
similar laws prohibiting or restricting circumvention of such
measures.

  When you convey a covered work, you waive any legal power to forbid
circumvention of technological measures to the extent such circumvention
is effected by exercising rights under this License with respect to
the covered work, and you disclaim any intention to limit operation or
modification of the work as a means of enforcing, against the work's
users, your or third parties' legal rights to forbid circumvention of
technological measures.

  4. Conveying Verbatim Copies.
```

(continues on next page)

(continued from previous page)

```
  You may convey verbatim copies of the Program's source code as you
receive it, in any medium, provided that you conspicuously and
appropriately publish on each copy an appropriate copyright notice;
keep intact all notices stating that this License and any
non-permissive terms added in accord with section 7 apply to the code;
keep intact all notices of the absence of any warranty; and give all
recipients a copy of this License along with the Program.

  You may charge any price or no price for each copy that you convey,
and you may offer support or warranty protection for a fee.

  5. Conveying Modified Source Versions.

  You may convey a work based on the Program, or the modifications to
produce it from the Program, in the form of source code under the
terms of section 4, provided that you also meet all of these conditions:

    a) The work must carry prominent notices stating that you modified
    it, and giving a relevant date.

    b) The work must carry prominent notices stating that it is
    released under this License and any conditions added under section
    7.  This requirement modifies the requirement in section 4 to
    "keep intact all notices".

    c) You must license the entire work, as a whole, under this
    License to anyone who comes into possession of a copy.  This
    License will therefore apply, along with any applicable section 7
    additional terms, to the whole of the work, and all its parts,
    regardless of how they are packaged.  This License gives no
    permission to license the work in any other way, but it does not
    invalidate such permission if you have separately received it.

    d) If the work has interactive user interfaces, each must display
    Appropriate Legal Notices; however, if the Program has interactive
    interfaces that do not display Appropriate Legal Notices, your
    work need not make them do so.

  A compilation of a covered work with other separate and independent
works, which are not by their nature extensions of the covered work,
and which are not combined with it such as to form a larger program,
in or on a volume of a storage or distribution medium, is called an
"aggregate" if the compilation and its resulting copyright are not
used to limit the access or legal rights of the compilation's users
beyond what the individual works permit.  Inclusion of a covered work
in an aggregate does not cause this License to apply to the other
parts of the aggregate.

  6. Conveying Non-Source Forms.

  You may convey a covered work in object code form under the terms
of sections 4 and 5, provided that you also convey the
```

(continues on next page)

(continued from previous page)

```
machine-readable Corresponding Source under the terms of this License,
in one of these ways:

    a) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by the
    Corresponding Source fixed on a durable physical medium
    customarily used for software interchange.

    b) Convey the object code in, or embodied in, a physical product
    (including a physical distribution medium), accompanied by a
    written offer, valid for at least three years and valid for as
    long as you offer spare parts or customer support for that product
    model, to give anyone who possesses the object code either (1) a
    copy of the Corresponding Source for all the software in the
    product that is covered by this License, on a durable physical
    medium customarily used for software interchange, for a price no
    more than your reasonable cost of physically performing this
    conveying of source, or (2) access to copy the
    Corresponding Source from a network server at no charge.

    c) Convey individual copies of the object code with a copy of the
    written offer to provide the Corresponding Source.  This
    alternative is allowed only occasionally and noncommercially, and
    only if you received the object code with such an offer, in accord
    with subsection 6b.

    d) Convey the object code by offering access from a designated
    place (gratis or for a charge), and offer equivalent access to the
    Corresponding Source in the same way through the same place at no
    further charge.  You need not require recipients to copy the
    Corresponding Source along with the object code.  If the place to
    copy the object code is a network server, the Corresponding Source
    may be on a different server (operated by you or a third party)
    that supports equivalent copying facilities, provided you maintain
    clear directions next to the object code saying where to find the
    Corresponding Source.  Regardless of what server hosts the
    Corresponding Source, you remain obligated to ensure that it is
    available for as long as needed to satisfy these requirements.

    e) Convey the object code using peer-to-peer transmission, provided
    you inform other peers where the object code and Corresponding
    Source of the work are being offered to the general public at no
    charge under subsection 6d.

  A separable portion of the object code, whose source code is excluded
from the Corresponding Source as a System Library, need not be
included in conveying the object code work.

  A "User Product" is either (1) a "consumer product", which means any
tangible personal property which is normally used for personal, family,
or household purposes, or (2) anything designed or sold for incorporation
into a dwelling.  In determining whether a product is a consumer product,
```

(continues on next page)

(continued from previous page)

```
doubtful cases shall be resolved in favor of coverage.  For a particular
product received by a particular user, "normally used" refers to a
typical or common use of that class of product, regardless of the status
of the particular user or of the way in which the particular user
actually uses, or expects or is expected to use, the product.  A product
is a consumer product regardless of whether the product has substantial
commercial, industrial or non-consumer uses, unless such uses represent
the only significant mode of use of the product.

  "Installation Information" for a User Product means any methods,
procedures, authorization keys, or other information required to install
and execute modified versions of a covered work in that User Product from
a modified version of its Corresponding Source.  The information must
suffice to ensure that the continued functioning of the modified object
code is in no case prevented or interfered with solely because
modification has been made.

  If you convey an object code work under this section in, or with, or
specifically for use in, a User Product, and the conveying occurs as
part of a transaction in which the right of possession and use of the
User Product is transferred to the recipient in perpetuity or for a
fixed term (regardless of how the transaction is characterized), the
Corresponding Source conveyed under this section must be accompanied
by the Installation Information.  But this requirement does not apply
if neither you nor any third party retains the ability to install
modified object code on the User Product (for example, the work has
been installed in ROM).

  The requirement to provide Installation Information does not include a
requirement to continue to provide support service, warranty, or updates
for a work that has been modified or installed by the recipient, or for
the User Product in which it has been modified or installed.  Access to a
network may be denied when the modification itself materially and
adversely affects the operation of the network or violates the rules and
protocols for communication across the network.

  Corresponding Source conveyed, and Installation Information provided,
in accord with this section must be in a format that is publicly
documented (and with an implementation available to the public in
source code form), and must require no special password or key for
unpacking, reading or copying.

  7. Additional Terms.

  "Additional permissions" are terms that supplement the terms of this
License by making exceptions from one or more of its conditions.
Additional permissions that are applicable to the entire Program shall
be treated as though they were included in this License, to the extent
that they are valid under applicable law.  If additional permissions
apply only to part of the Program, that part may be used separately
under those permissions, but the entire Program remains governed by
this License without regard to the additional permissions.
```

(continues on next page)

(continued from previous page)

```
  When you convey a copy of a covered work, you may at your option
remove any additional permissions from that copy, or from any part of
it.  (Additional permissions may be written to require their own
removal in certain cases when you modify the work.)  You may place
additional permissions on material, added by you to a covered work,
for which you have or can give appropriate copyright permission.

  Notwithstanding any other provision of this License, for material you
add to a covered work, you may (if authorized by the copyright holders of
that material) supplement the terms of this License with terms:

    a) Disclaiming warranty or limiting liability differently from the
    terms of sections 15 and 16 of this License; or

    b) Requiring preservation of specified reasonable legal notices or
    author attributions in that material or in the Appropriate Legal
    Notices displayed by works containing it; or

    c) Prohibiting misrepresentation of the origin of that material, or
    requiring that modified versions of such material be marked in
    reasonable ways as different from the original version; or

    d) Limiting the use for publicity purposes of names of licensors or
    authors of the material; or

    e) Declining to grant rights under trademark law for use of some
    trade names, trademarks, or service marks; or

    f) Requiring indemnification of licensors and authors of that
    material by anyone who conveys the material (or modified versions of
    it) with contractual assumptions of liability to the recipient, for
    any liability that these contractual assumptions directly impose on
    those licensors and authors.

  All other non-permissive additional terms are considered "further
restrictions" within the meaning of section 10.  If the Program as you
received it, or any part of it, contains a notice stating that it is
governed by this License along with a term that is a further
restriction, you may remove that term.  If a license document contains
a further restriction but permits relicensing or conveying under this
License, you may add to a covered work material governed by the terms
of that license document, provided that the further restriction does
not survive such relicensing or conveying.

  If you add terms to a covered work in accord with this section, you
must place, in the relevant source files, a statement of the
additional terms that apply to those files, or a notice indicating
where to find the applicable terms.

  Additional terms, permissive or non-permissive, may be stated in the
form of a separately written license, or stated as exceptions;
```

(continues on next page)

(continued from previous page)

```
the above requirements apply either way.

  8. Termination.

  You may not propagate or modify a covered work except as expressly
provided under this License.  Any attempt otherwise to propagate or
modify it is void, and will automatically terminate your rights under
this License (including any patent licenses granted under the third
paragraph of section 11).

  However, if you cease all violation of this License, then your
license from a particular copyright holder is reinstated (a)
provisionally, unless and until the copyright holder explicitly and
finally terminates your license, and (b) permanently, if the copyright
holder fails to notify you of the violation by some reasonable means
prior to 60 days after the cessation.

  Moreover, your license from a particular copyright holder is
reinstated permanently if the copyright holder notifies you of the
violation by some reasonable means, this is the first time you have
received notice of violation of this License (for any work) from that
copyright holder, and you cure the violation prior to 30 days after
your receipt of the notice.

  Termination of your rights under this section does not terminate the
licenses of parties who have received copies or rights from you under
this License.  If your rights have been terminated and not permanently
reinstated, you do not qualify to receive new licenses for the same
material under section 10.

  9. Acceptance Not Required for Having Copies.

  You are not required to accept this License in order to receive or
run a copy of the Program.  Ancillary propagation of a covered work
occurring solely as a consequence of using peer-to-peer transmission
to receive a copy likewise does not require acceptance.  However,
nothing other than this License grants you permission to propagate or
modify any covered work.  These actions infringe copyright if you do
not accept this License.  Therefore, by modifying or propagating a
covered work, you indicate your acceptance of this License to do so.

  10. Automatic Licensing of Downstream Recipients.

  Each time you convey a covered work, the recipient automatically
receives a license from the original licensors, to run, modify and
propagate that work, subject to this License.  You are not responsible
for enforcing compliance by third parties with this License.

  An "entity transaction" is a transaction transferring control of an
organization, or substantially all assets of one, or subdividing an
organization, or merging organizations.  If propagation of a covered
work results from an entity transaction, each party to that
```

(continues on next page)

(continued from previous page)

```
transaction who receives a copy of the work also receives whatever
licenses to the work the party's predecessor in interest had or could
give under the previous paragraph, plus a right to possession of the
Corresponding Source of the work from the predecessor in interest, if
the predecessor has it or can get it with reasonable efforts.

  You may not impose any further restrictions on the exercise of the
rights granted or affirmed under this License.  For example, you may
not impose a license fee, royalty, or other charge for exercise of
rights granted under this License, and you may not initiate litigation
(including a cross-claim or counterclaim in a lawsuit) alleging that
any patent claim is infringed by making, using, selling, offering for
sale, or importing the Program or any portion of it.

  11. Patents.

  A "contributor" is a copyright holder who authorizes use under this
License of the Program or a work on which the Program is based.  The
work thus licensed is called the contributor's "contributor version".

  A contributor's "essential patent claims" are all patent claims
owned or controlled by the contributor, whether already acquired or
hereafter acquired, that would be infringed by some manner, permitted
by this License, of making, using, or selling its contributor version,
but do not include claims that would be infringed only as a
consequence of further modification of the contributor version.  For
purposes of this definition, "control" includes the right to grant
patent sublicenses in a manner consistent with the requirements of
this License.

  Each contributor grants you a non-exclusive, worldwide, royalty-free
patent license under the contributor's essential patent claims, to
make, use, sell, offer for sale, import and otherwise run, modify and
propagate the contents of its contributor version.

  In the following three paragraphs, a "patent license" is any express
agreement or commitment, however denominated, not to enforce a patent
(such as an express permission to practice a patent or covenant not to
sue for patent infringement).  To "grant" such a patent license to a
party means to make such an agreement or commitment not to enforce a
patent against the party.

  If you convey a covered work, knowingly relying on a patent license,
and the Corresponding Source of the work is not available for anyone
to copy, free of charge and under the terms of this License, through a
publicly available network server or other readily accessible means,
then you must either (1) cause the Corresponding Source to be so
available, or (2) arrange to deprive yourself of the benefit of the
patent license for this particular work, or (3) arrange, in a manner
consistent with the requirements of this License, to extend the patent
license to downstream recipients.  "Knowingly relying" means you have
actual knowledge that, but for the patent license, your conveying the
```

(continues on next page)

(continued from previous page)

```
covered work in a country, or your recipient's use of the covered work
in a country, would infringe one or more identifiable patents in that
country that you have reason to believe are valid.

  If, pursuant to or in connection with a single transaction or
arrangement, you convey, or propagate by procuring conveyance of, a
covered work, and grant a patent license to some of the parties
receiving the covered work authorizing them to use, propagate, modify
or convey a specific copy of the covered work, then the patent license
you grant is automatically extended to all recipients of the covered
work and works based on it.

  A patent license is "discriminatory" if it does not include within
the scope of its coverage, prohibits the exercise of, or is
conditioned on the non-exercise of one or more of the rights that are
specifically granted under this License.  You may not convey a covered
work if you are a party to an arrangement with a third party that is
in the business of distributing software, under which you make payment
to the third party based on the extent of your activity of conveying
the work, and under which the third party grants, to any of the
parties who would receive the covered work from you, a discriminatory
patent license (a) in connection with copies of the covered work
conveyed by you (or copies made from those copies), or (b) primarily
for and in connection with specific products or compilations that
contain the covered work, unless you entered into that arrangement,
or that patent license was granted, prior to 28 March 2007.

  Nothing in this License shall be construed as excluding or limiting
any implied license or other defenses to infringement that may
otherwise be available to you under applicable patent law.

  12. No Surrender of Others' Freedom.

  If conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot convey a
covered work so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you may
not convey it at all.  For example, if you agree to terms that obligate you
to collect a royalty for further conveying from those to whom you convey
the Program, the only way you could satisfy both those terms and this
License would be to refrain entirely from conveying the Program.

  13. Use with the GNU Affero General Public License.

  Notwithstanding any other provision of this License, you have
permission to link or combine any covered work with a work licensed
under version 3 of the GNU Affero General Public License into a single
combined work, and to convey the resulting work.  The terms of this
License will continue to apply to the part which is the covered work,
but the special requirements of the GNU Affero General Public License,
section 13, concerning interaction through a network will apply to the
```

(continues on next page)

(continued from previous page)

```
combination as such.

  14. Revised Versions of this License.

  The Free Software Foundation may publish revised and/or new versions of
the GNU General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

  Each version is given a distinguishing version number.  If the
Program specifies that a certain numbered version of the GNU General
Public License "or any later version" applies to it, you have the
option of following the terms and conditions either of that numbered
version or of any later version published by the Free Software
Foundation.  If the Program does not specify a version number of the
GNU General Public License, you may choose any version ever published
by the Free Software Foundation.

  If the Program specifies that a proxy can decide which future
versions of the GNU General Public License can be used, that proxy's
public statement of acceptance of a version permanently authorizes you
to choose that version for the Program.

  Later license versions may give you additional or different
permissions.  However, no additional obligations are imposed on any
author or copyright holder as a result of your choosing to follow a
later version.

  15. Disclaimer of Warranty.

  THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY
APPLICABLE LAW.  EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT
HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY
OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM
IS WITH YOU.  SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF
ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. Limitation of Liability.

  IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS
THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY
GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE
USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF
DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD
PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS),
EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF
SUCH DAMAGES.

  17. Interpretation of Sections 15 and 16.
```

(continues on next page)

(continued from previous page)

```
  If the disclaimer of warranty and limitation of liability provided
above cannot be given local legal effect according to their terms,
reviewing courts shall apply local law that most closely approximates
an absolute waiver of all civil liability in connection with the
Program, unless a warranty or assumption of liability accompanies a
copy of the Program in return for a fee.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
state the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software: you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation, either version 3 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program.  If not, see <http://www.gnu.org/licenses/>.

Also add information on how to contact you by electronic and paper mail.

  If the program does terminal interaction, make it output a short
notice like this when it starts in an interactive mode:

    <program>  Copyright (C) <year>  <name of author>
    This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, your program's commands
might be different; for a GUI interface, you would use an "about box".

  You should also get your employer (if you work as a programmer) or school,
```

(continues on next page)

(continued from previous page)

```
if any, to sign a "copyright disclaimer" for the program, if necessary.
For more information on this, and how to apply and follow the GNU GPL, see
<http://www.gnu.org/licenses/>.

  The GNU General Public License does not permit incorporating your program
into proprietary programs.  If your program is a subroutine library, you
may consider it more useful to permit linking proprietary applications with
the library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.  But first, please read
<http://www.gnu.org/philosophy/why-not-lgpl.html>.
```

40.128 GPLv2 License Text for rrdtool

The following license text applies to rrdtool:

Listing 128: Download: GPLv2 for rrdtool

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
```

(continues on next page)

(continued from previous page)

```
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.
```

(continues on next page)

(continued from previous page)

```
You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:
```

(continues on next page)

(continued from previous page)

```
    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
```

(continues on next page)

(continued from previous page)

```
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.
```

(continues on next page)

(continued from previous page)

```
Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS
```

40.129 GPLv2+ License Text for smartmontools

The following license text applies to smartmontools:

Listing 129: Download: GPLv2+ for smartmontools

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
```

(continues on next page)

(continued from previous page)

```
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
```

(continues on next page)

(continued from previous page)

```
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
```

(continues on next page)

(continued from previous page)

```
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
```

(continues on next page)

(continued from previous page)

```
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
```

(continues on next page)

(continued from previous page)

```
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
```

(continues on next page)

(continued from previous page)

```
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
```

(continues on next page)

(continued from previous page)

```
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

40.130 GPLv2 License Text for strongswan

The following license text applies to strongswan:

Listing 130: Download: GPLv2 for strongswan

```
Except for code in the blowfish, des, md4 and md5 plugins (see below) the
following terms apply:

For copyright information see the headers of individual source files.

This program is free software; you can redistribute it and/or modify it under
the terms of the GNU General Public License as published by the Free Software
Foundation; either version 2 of the License, or (at your option) any later
version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY
WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A
PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with
this program; if not, see <http://www.gnu.org/licenses>.

Linking strongSwan statically or dynamically with other modules is making a
combined work based on strongSwan. Thus, the terms and conditions of the GNU
General Public License cover the whole combination.
```

(continues on next page)

(continued from previous page)

```
In addition, as a special exception, the copyright holders of strongSwan give
you permission to combine strongSwan with free software programs or libraries
that are released under the GNU LGPL and with code included in the standard
release of the OpenSSL project's OpenSSL library under the OpenSSL or SSLeay
licenses (or modified versions of such code, with unchanged license). You may
copy and distribute such a system following the terms of the GNU GPL for
strongSwan and the licenses of the other code concerned, provided that you
include the source code of that other code when and as the GNU GPL requires
distribution of source code.

Note that people who make modified versions of strongSwan are not obligated to
grant this special exception for their modified versions; it is their choice
whether to do so. The GNU General Public License gives permission to release a
modified version without this exception; this exception also makes it possible
to release a modified version which carries forward this exception.


The DES implementation in the des plugin and the Blowfish implementation in the
blowfish plugin are under a BSD style license (see source files for details).
Note that these parts have an advertising clause in it.

The MD4 and MD5 implementations in the md4 and md5 plugins are from RSA Data
Security Inc., so this package must include the following phrase:
"derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm".
```

40.131 BSD2CLAUSE License Text for uclcmd

The following license text applies to uclcmd:

Listing 131: Download: BSD2CLAUSE for uclcmd

```
Copyright (c) 2014, Allan Jude
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this
  list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice,
  this list of conditions and the following disclaimer in the documentation
  and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE
FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
```

(continues on next page)

(continued from previous page)

```
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.132 BSD3CLAUSE License Text for unbound

The following license text applies to unbound:

Listing 132: Download: BSD3CLAUSE for unbound

```
Copyright (c) 2007, NLnet Labs. All rights reserved.

This software is open source.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice,
this list of conditions and the following disclaimer in the documentation
and/or other materials provided with the distribution.

Neither the name of the NLNET LABS nor the names of its contributors may
be used to endorse or promote products derived from this software without
specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED
TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR
PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.133 LGPL21+ License Text for vstr

The following license text applies to vstr:

Listing 133: Download: LGPL21+ for vstr

```
                  GNU LESSER GENERAL PUBLIC LICENSE
                       Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
     59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations
below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.
```

(continues on next page)

(continued from previous page)

```
  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be
introduced by others.
^L
  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it
becomes
a de-facto standard.  To achieve this, non-free programs must be
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
```

(continues on next page)

(continued from previous page)

```
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
that program using a modified version of the Library.

  The precise terms and conditions for copying, distribution and
modification follow.  Pay close attention to the difference between a
"work based on the library" and a "work that uses the library".  The
former contains code derived from the library, whereas the latter must
be combined with the library in order to run.
^L
                  GNU LESSER GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License Agreement applies to any software library or other
program which contains a notice placed by the copyright holder or
other authorized party saying it may be distributed under the terms of
this Lesser General Public License (also called "this License").
Each licensee is addressed as "you".

  A "library" means a collection of software functions and/or data
prepared so as to be conveniently linked with application programs
(which use some of those functions and data) to form executables.

  The "Library", below, refers to any such software library or work
which has been distributed under these terms.  A "work based on the
Library" means either the Library or any derivative work under
copyright law: that is to say, a work containing the Library or a
portion of it, either verbatim or with modifications and/or translated
straightforwardly into another language.  (Hereinafter, translation is
included without limitation in the term "modification".)

  "Source code" for a work means the preferred form of the work for
making modifications to it.  For a library, complete source code means
all the source code for all modules it contains, plus any associated
interface definition files, plus the scripts used to control
compilation
and installation of the library.

  Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running a program using the Library is not restricted, and output from
such a program is covered only if its contents constitute a work based
on the Library (independent of the use of the Library in a tool for
writing it).  Whether that is true depends on what the Library does
and what the program that uses the Library does.
```

(continues on next page)

(continued from previous page)

```
  1. You may copy and distribute verbatim copies of the Library's
complete source code as you receive it, in any medium, provided that
you conspicuously and appropriately publish on each copy an
appropriate copyright notice and disclaimer of warranty; keep intact
all the notices that refer to this License and to the absence of any
warranty; and distribute a copy of this License along with the
Library.

  You may charge a fee for the physical act of transferring a copy,
and you may at your option offer warranty protection in exchange for a

fee.

  2. You may modify your copy or copies of the Library or any portion
of it, thus forming a work based on the Library, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) The modified work must itself be a software library.

    b) You must cause the files modified to carry prominent notices
    stating that you changed the files and the date of any change.

    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.
```

(continues on next page)

(continued from previous page)

```
Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in
these notices.
^L
  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
from a designated place, then offering equivalent access to copy the
source code from the same place satisfies the requirement to
distribute the source code, even though third parties are not
compelled to copy the source along with the object code.

  5. A program that contains no derivative of any portion of the
Library, but is designed to work with the Library by being compiled or
linked with it, is called a "work that uses the Library".  Such a
work, in isolation, is not a derivative work of the Library, and
therefore falls outside the scope of this License.

  However, linking a "work that uses the Library" with the Library
creates an executable that is a derivative of the Library (because it
contains portions of the Library), rather than a "work that uses the
library".  The executable is therefore covered by this License.
Section 6 states terms for distribution of such executables.
```

(continues on next page)

(continued from previous page)

```
  When a "work that uses the Library" uses material from a header file
that is part of the Library, the object code for the work may be a
derivative work of the Library even though the source code is not.
Whether this is true is especially significant if the work can be
linked without the Library, or if the work is itself a library.  The
threshold for this to be true is not precisely defined by law.

  If such an object file uses only numerical parameters, data
structure layouts and accessors, and small macros and small inline
functions (ten lines or less in length), then the use of the object
file is unrestricted, regardless of whether it is legally a derivative
work.  (Executables containing this object code plus portions of the
Library will still fall under Section 6.)

  Otherwise, if the work is a derivative of the Library, you may
distribute the object code for the work under the terms of Section 6.
Any executables containing that work also fall under Section 6,
whether or not they are linked directly with the Library itself.
^L
  6. As an exception to the Sections above, you may also combine or
link a "work that uses the Library" with the Library to produce a
work containing portions of the Library, and distribute that work
under terms of your choice, provided that the terms permit
modification of the work for the customer's own use and reverse
engineering for debugging such modifications.

  You must give prominent notice with each copy of the work that the
Library is used in it and that the Library and its use are covered by
this License.  You must supply a copy of this License.  If the work
during execution displays copyright notices, you must include the
copyright notice for the Library among them, as well as a reference
directing the user to the copy of this License.  Also, you must do one
of these things:

    a) Accompany the work with the complete corresponding
    machine-readable source code for the Library including whatever
    changes were used in the work (which must be distributed under
    Sections 1 and 2 above); and, if the work is an executable linked
    with the Library, with the complete machine-readable "work that
    uses the Library", as object code and/or source code, so that the
    user can modify the Library and then relink to produce a modified
    executable containing the modified Library.  (It is understood
    that the user who changes the contents of definitions files in the
    Library will not necessarily be able to recompile the application
    to use the modified definitions.)

    b) Use a suitable shared library mechanism for linking with the
    Library.  A suitable mechanism is one that (1) uses at run time a
    copy of the library already present on the user's computer system,
    rather than copying library functions into the executable, and (2)
    will operate properly with a modified version of the library, if
    the user installs one, as long as the modified version is
```

(continues on next page)

(continued from previous page)

```
    interface-compatible with the version that the work was made with.

    c) Accompany the work with a written offer, valid for at
    least three years, to give the same user the materials
    specified in Subsection 6a, above, for a charge no more
    than the cost of performing this distribution.

    d) If distribution of the work is made by offering access to copy
    from a designated place, offer equivalent access to copy the above
    specified materials from the same place.

    e) Verify that the user has already received a copy of these
    materials or that you have already sent this user a copy.

  For an executable, the required form of the "work that uses the
Library" must include any data and utility programs needed for
reproducing the executable from it.  However, as a special exception,
the materials to be distributed need not include anything that is
normally distributed (in either source or binary form) with the major
components (compiler, kernel, and so on) of the operating system on
which the executable runs, unless that component itself accompanies
the executable.

  It may happen that this requirement contradicts the license
restrictions of other proprietary libraries that do not normally
accompany the operating system.  Such a contradiction means you cannot
use both them and the Library together in an executable that you
distribute.
^L
  7. You may place library facilities that are a work based on the
Library side-by-side in a single library together with other library
facilities not covered by this License, and distribute such a combined
library, provided that the separate distribution of the work based on
the Library and of the other library facilities is otherwise
permitted, and provided that you do these two things:

    a) Accompany the combined library with a copy of the same work
    based on the Library, uncombined with any other library
    facilities.  This must be distributed under the terms of the
    Sections above.

    b) Give prominent notice with the combined library of the fact
    that part of it is a work based on the Library, and explaining
    where to find the accompanying uncombined form of the same work.

  8. You may not copy, modify, sublicense, link with, or distribute
the Library except as expressly provided under this License.  Any
attempt otherwise to copy, modify, sublicense, link with, or
distribute the Library is void, and will automatically terminate your
rights under this License.  However, parties who have received copies,
or rights, from you under this License will not have their licenses
terminated so long as such parties remain in full compliance.
```

(continues on next page)

(continued from previous page)

```
  9. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Library or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Library (or any work based on the
Library), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with
this License.
^L
  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all.  For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply, and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
```

(continues on next page)

(continued from previous page)

```
original copyright holder who places the Library under this License
may add an explicit geographical distribution limitation excluding those
countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number.  If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation.  If the Library does not specify a
license version number, you may choose any version ever published by
the Free Software Foundation.
^L
  14. If you wish to incorporate parts of the Library into other free
programs whose distribution conditions are incompatible with these,
write to the author to ask for permission.  For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this.  Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

                            NO WARRANTY

  15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE.  THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU.  SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

  16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.

                     END OF TERMS AND CONDITIONS
```

(continues on next page)

(continued from previous page)

```
^L
           How to Apply These Terms to Your New Libraries

  If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change.  You can do so by permitting
redistribution under these terms (or, alternatively, under the terms
of the ordinary General Public License).

  To apply these terms, attach the following notices to the library.
It is safest to attach them to the start of each source file to most
effectively convey the exclusion of warranty; and each file should
have at least the "copyright" line and a pointer to where the full
notice is found.


    <one line to give the library's name and a brief idea of what it
does.>
    Copyright (C) <year>  <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307  USA

Also add information on how to contact you by electronic and paper
mail.

You should also get your employer (if you work as a programmer) or
your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library `Frob' (a library for tweaking knobs) written by James
Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That's all there is to it!
```

40.134 GPLv2+ License Text for wol

The following license text applies to wol:

Listing 134: Download: GPLv2+ for wol

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.
     59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Library General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.

  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
```

(continues on next page)

(continued from previous page)

```
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and

modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.

  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
```

(continues on next page)

(continued from previous page)

```
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on

    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
```

(continues on next page)

(continued from previous page)

```
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not

compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
```

(continues on next page)

(continued from previous page)

```
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to

be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
```

(continues on next page)

(continued from previous page)

```
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY

  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.


                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
```

(continues on next page)

(continued from previous page)

```
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License
    along with this program; if not, write to the Free Software
    Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA


Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year  name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Library General
Public License instead of this License.
```

40.135 BSD3CLAUSE License Text for wpa_supplicant

The following license text applies to wpa_supplicant:

Listing 135: Download: BSD3CLAUSE for wpa_supplicant

```
wpa_supplicant and hostapd
--------------------------

Copyright (c) 2002-2019, Jouni Malinen <j@w1.fi> and contributors
All Rights Reserved.
```

(continues on next page)

(continued from previous page)

```
These programs are licensed under the BSD license (the one with
advertisement clause removed).

If you are submitting changes to the project, please see CONTRIBUTIONS
file for more instructions.


This package may include either wpa_supplicant, hostapd, or both. See
README file respective subdirectories (wpa_supplicant/README or
hostapd/README) for more details.

Source code files were moved around in v0.6.x releases and compared to
earlier releases, the programs are now built by first going to a
subdirectory (wpa_supplicant or hostapd) and creating build
configuration (.config) and running 'make' there (for Linux/BSD/cygwin
builds).


License
-------

This software may be distributed, used, and modified under the terms of
BSD license:

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are
met:

1. Redistributions of source code must retain the above copyright
   notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright
   notice, this list of conditions and the following disclaimer in the
   documentation and/or other materials provided with the distribution.

3. Neither the name(s) of the above-listed copyright holder(s) nor the
   names of its contributors may be used to endorse or promote products
   derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

40.136 XINETD License Text for xinetd

The following license text applies to xinetd:

Listing 136: Download: XINETD for xinetd

```
ORIGINAL LICENSE:
This software is

(c) Copyright 1992 by Panagiotis Tsirigotis

The author (Panagiotis Tsirigotis) grants permission to use, copy,
and distribute this software and its documentation for any purpose
and without fee, provided that the above copyright notice extant in
files in this distribution is not removed from files included in any
redistribution and that this copyright notice is also included in any
redistribution.

Modifications to this software may be distributed, either by distributing
the modified software or by distributing patches to the original software,
under the following additional terms:

1. The version number will be modified as follows:
      a. The first 3 components of the version number
         (i.e <number>.<number>.<number>) will remain unchanged.
      b. A new component will be appended to the version number to indicate
         the modification level. The form of this component is up to the
         author of the modifications.

2. The author of the modifications will include his/her name by appending it
   along with the new version number to this file and will be responsible for
   any wrong behavior of the modified software.

The author makes no representations about the suitability of this
software for any purpose.  It is provided "as is" without any express
or implied warranty.


Modifications:
Version: 2.1.8.7-current
Copyright 1998-2001 by Rob Braun

Sensor Addition
Version: 2.1.8.9pre14a
Copyright 2001 by Steve Grubb

This is an exerpt from an email I recieved from the original author, allowing
xinetd as maintained by me, to use the higher version numbers:

I appreciate your maintaining the version string guidelines as specified
in the copyright. But I did not mean them to last as long as they did.

So, if you want, you may use any 2.N.* (N >= 3) version string for future
```

(continues on next page)

(continued from previous page)

```
xinetd versions that you release. Note that I am excluding the 2.2.* line;
using that would only create confusion. Naming the next release 2.3.0
would put to rest the confusion about 2.2.1 and 2.1.8.*.
```

CHAPTER

FORTYONE

CONFIGURATION RECIPES

CHAPTER

FORTYTWO

ADDITIONAL COMMERCIAL RESOURCES

- Netgate TAC
- Netgate Professional Services
- pfSense Training

INDEX

D

DNS, **1929**

H

HTTP, **1929**

I

ICMP, **1929**
IP, **1929**

L

LAN, **1929**

N

NAT, **1929**

S

SSH, **1929**

T

TCP, **1929**

U

UDP, **1929**

V

VM, **1929**
VPN, **1929**

W

WAN, **1929**