

Documentation 2.0

ZABBIX

16.12.2025

Contents

Zabbix Manual	4
Copyright notice	4
1. Introduction	4
1 Frontend improvements	4
1 Manual structure	4
2 What is Zabbix	5
3 Zabbix features	5
4 Zabbix overview	6
5 What's new in Zabbix 2.0.0	7
6 What's new in Zabbix 2.0.1	21
7 What's new in Zabbix 2.0.2	22
8 What's new in Zabbix 2.0.3	24
9 What's new in Zabbix 2.0.4	26
10 What's new in Zabbix 2.0.5	27
11 What's new in Zabbix 2.0.6	28
12 What's new in Zabbix 2.0.7	29
13 What's new in Zabbix 2.0.8	30
14 What's new in Zabbix 2.0.9	31
15 What's new in Zabbix 2.0.10	32
16 What's new in Zabbix 2.0.11	32
17 What's new in Zabbix 2.0.12	33
18 What's new in Zabbix 2.0.13	33
19 What's new in Zabbix 2.0.14	34
20 What's new in Zabbix 2.0.15	34
21 What's new in Zabbix 2.0.16	34
22 What's new in Zabbix 2.0.17	34
2. Zabbix concepts	35
1 Zabbix definitions	35
2 Server	36
3 Agent	38
4 Proxy	39
5 Java gateway	41
6 Sender	43
7 Get	44
3. Installation	44
1 Getting Zabbix	44
2 Requirements	44
3 Installation from packages	49
4 Installation from sources	51
5 Known issues	59
6 Upgrade procedure	60
7 Upgrade notes for 2.0.0	61
8 Upgrade notes for 2.0.1	63
9 Upgrade notes for 2.0.2	63
10 Upgrade notes for 2.0.3	63
11 Upgrade notes for 2.0.4	63
12 Upgrade notes for 2.0.5	63
13 Upgrade notes for 2.0.6	63
14 Upgrade notes for 2.0.7	64
15 Upgrade notes for 2.0.8	64

16 Upgrade notes for 2.0.9	64
17 Upgrade notes for 2.0.10	64
18 Upgrade notes for 2.0.11	64
19 Upgrade notes for 2.0.12	64
20 Upgrade notes for 2.0.13	65
21 Upgrade notes for 2.0.14	65
22 Upgrade notes for 2.0.15	65
23 Upgrade notes for 2.0.16	65
24 Upgrade notes for 2.0.17	65
4. Quickstart	65
1 Login and configuring user	66
2 New host	69
3 New item	70
4 New trigger	72
5 Receiving problem notification	73
6 New template	76
5. Zabbix appliance	78
6. Configuration	82
1 Hosts and host groups	86
2 Items	90
3 Triggers	180
4 Events	187
5 Visualisation	188
6 Templates	203
7 Notifications upon events	204
8 Macros	226
9 Users and user groups	227
7. IT services	230
8. Web monitoring	233
1 Web monitoring items	236
2 Real life scenario	238
9. Maintenance	241
10. Regular expressions	243
11. Event acknowledgment	244
12. Configuration export/import	244
Groups	246
Hosts	246
13. Discovery	252
1 Network discovery	252
2 Active agent auto-registration	258
3 Low-level discovery	259
14. Distributed monitoring	272
1 Proxies	273
2 Nodes	275
15. Web interface	278
1 Frontend sections	279
2 User profile	331
3 Global search	335
4 Frontend maintenance mode	336
5 Definitions	337
6 Creating your own theme	338
7 Debug mode	339
16. Appendixes	339
Zabbix API	339
1 Frequently asked questions / Troubleshooting	693
2 Installation	694
3 Daemon configuration	697
4 Items	716
5 Triggers	727
6 Macros	743
7 Setting time periods	780
8 Command execution	780
9 Recipes for monitoring	781

10 Performance tuning	783
11 Version compatibility	784
Zabbix manpages	784
zabbix_agentd	784
NAME	785
SYNOPSIS	785
DESCRIPTION	785
FILES	785
SEE ALSO	785
AUTHOR	785
Index	785
zabbix_get	786
NAME	786
SYNOPSIS	786
DESCRIPTION	786
EXAMPLES	786
SEE ALSO	786
AUTHOR	786
Index	786
zabbix_proxy	787
NAME	787
SYNOPSIS	787
DESCRIPTION	787
FILES	787
SEE ALSO	787
AUTHOR	788
Index	788
zabbix_sender	788
NAME	788
SYNOPSIS	788
DESCRIPTION	788
EXAMPLES	789
SEE ALSO	789
AUTHOR	789
Index	789
zabbix_server	790
NAME	790
SYNOPSIS	790
DESCRIPTION	790
FILES	790
SEE ALSO	790
AUTHOR	791
Index	791
Zabbix API	791

Zabbix Manual

Welcome to the user manual for Zabbix 2.0 software. These pages are created to help our users successfully manage their monitoring tasks with Zabbix, from the simple to the more complex.

Copyright notice

Zabbix documentation is NOT distributed under a GPL license. Use of Zabbix documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Zabbix disseminates it (that is, electronically for download on a Zabbix web site) or on a USB or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Zabbix. Zabbix reserves any and all rights to this documentation not expressly granted above.

1. Introduction

Please use the sidebar to access content in the Introduction section.

1 Frontend improvements

1.1.1 Updated translations

- * French

1.2 Daemon improvements

1.3 Miscellaneous improvements Fixed issues with mysql user parameter configuration script mysql.size parameter. It contained a complex bash expression and was failing if the default shell was not bash (CVE-2016-4338).

1 Manual structure

Structure

The content of this Zabbix 2.0 manual is divided into sections and subsections to provide easy access to particular subjects of interest.

When you navigate to respective sections, make sure that you expand section folders to reveal full content of what is included in subsections and individual pages.

Cross-linking between pages of related content is provided as much as possible to make sure that relevant information is not missed by the users.

Sections

Introduction provides general information about current Zabbix software. Reading this section should equip you with some good reasons to choose Zabbix.

Zabbix concepts explain the terminology used in Zabbix and provides details on Zabbix components.

Installation and **Quickstart** sections should help you to get started with Zabbix. **Zabbix appliance** is an alternative for getting a quick taster of what it is like to use Zabbix.

Configuration is one of the largest and more important sections in this manual. It contains loads of essential advice about how to set up Zabbix to monitor your environment, from setting up hosts to getting essential data to viewing data to configuring notifications and remote commands to be executed in case of problems.

IT services section details how to use Zabbix for a high-level overview of your monitoring environment.

Web monitoring should help you learn how to monitor the availability of web sites.

[Maintenance](#), [Regular expressions](#), [Event acknowledgment](#) and [XML export/import](#) are further sections that reveal how to use these various aspects of Zabbix software.

[Discovery](#) contains instructions for setting up automatic discovery of network devices, active agents, file systems, network interfaces, etc.

[Distributed monitoring](#) deals with the possibilities of using Zabbix in larger and more complex environments.

[Web interface](#) contains information specific for using the web interface of Zabbix.

Detailed lists of technical information are included in [Appendixes](#). This is where you will also find FAQ and Zabbix API sections.

2 What is Zabbix

Overview

Zabbix was created by Alexei Vladishev, and currently is actively developed and supported by Zabbix SIA.

Zabbix is an enterprise-class open source distributed monitoring solution.

Zabbix is software that monitors numerous parameters of a network and the health and integrity of servers. Zabbix uses a flexible notification mechanism that allows users to configure e-mail based alerts for virtually any event. This allows a fast reaction to server problems. Zabbix offers excellent reporting and data visualisation features based on the stored data. This makes Zabbix ideal for capacity planning.

Zabbix supports both polling and trapping. All Zabbix reports and statistics, as well as configuration parameters, are accessed through a web-based frontend. A web-based frontend ensures that the status of your network and the health of your servers can be assessed from any location. Properly configured, Zabbix can play an important role in monitoring IT infrastructure. This is equally true for small organisations with a few servers and for large companies with a multitude of servers.

Zabbix is free of cost. Zabbix is written and distributed under the GPL General Public License version 2. It means that its source code is freely distributed and available for the general public.

[Commercial support](#) is available and provided by Zabbix Company.

Learn more about [Zabbix features](#).

Users of Zabbix

Many organisations of different size around the world rely on Zabbix as a primary monitoring platform.

3 Zabbix features

Overview

Zabbix is a highly integrated network monitoring solution, offering a multiplicity of features in a single package.

Data gathering

- availability and performance checks
- support for SNMP (both trapping and polling), IPMI, JMX monitoring
- custom checks
- gathering desired data at custom intervals
- performed by server/proxy and by agents

Flexible threshold definitions

- you can define very flexible problem thresholds, called triggers, referencing values from the backend database

Highly configurable alerting

- sending notifications can be customized for the escalation schedule, recipient, media type
- notifications can be made meaningful and helpful using macro variables
- automatic actions include remote commands

Real-time graphing

- monitored items are immediately graphed using the built-in graphing functionality

Web monitoring capabilities

- Zabbix can follow a path of simulated mouse clicks on a web site and check for functionality and response time

Extensive visualisation options

- ability to create custom graphs that can combine multiple items into a single view
- network maps
- custom screens and slide shows for a dashboard-style overview
- reports
- high-level (business) view of monitored resources

Historical data storage

- data stored in a database
- configurable history
- built-in housekeeping procedure

Easy configuration

- add monitored devices as hosts
- hosts are picked up for monitoring, once in the database
- apply templates to monitored devices

Use of templates

- grouping checks in templates
- templates can inherit other templates

Network discovery

- automatic discovery of network devices
- agent auto registration
- discovery of file systems, network interfaces and SNMP OIDs

Fast web interface

- a web-based frontend in PHP
- accessible from anywhere
- you can click your way through
- audit log

Zabbix API

- Zabbix API provides programmable interface to Zabbix for mass manipulations, 3rd party software integration and other purposes.

Permissions system

- secure user authentication
- certain users can be limited to certain views

Full featured and easily extensible agent

- deployed on monitoring targets
- can be deployed on both Linux and Windows

Binary daemons

- written in C, for performance and small memory footprint
- easily portable

Ready for complex environments

- remote monitoring made easy by using a Zabbix proxy

4 Zabbix overview

Architecture

Zabbix consists of several major software components, the responsibilities of which are outlined below.

Server

Zabbix server is the central component to which agents report availability and integrity information and statistics. The server is the central repository in which all configuration, statistical and operational data are stored.

Database storage

All configuration information as well as the data gathered by Zabbix is stored in a database.

Web interface

For an easy access to Zabbix from anywhere and from any platform, the web-based interface is provided. The interface is part of Zabbix server, and usually (but not necessarily) runs on the same physical machine as the one running the server.

Note:

Zabbix web interface must run on the same physical machine if SQLite is used.

Proxy

Zabbix proxy can collect performance and availability data on behalf of Zabbix server. A proxy is an optional part of Zabbix deployment; however, it may be very beneficial to distribute the load of a single Zabbix server.

Agent

Zabbix agents are deployed on monitoring targets to actively monitor local resources and applications and report the gathered data to Zabbix server.

Data flow

In addition it is important to take a step back and have a look at the overall data flow within Zabbix. In order to create an item that gathers data you must first create a host. Moving to the other end of the Zabbix spectrum you must first have an item to create a trigger. You must have a trigger to create an action. Thus if you want to receive an alert that your CPU load is too high on *Server X* you must first create a host entry for *Server X* followed by an item for monitoring its CPU, then a trigger which activates if the CPU is too high, followed by an action which sends you an email. While that may seem like a lot of steps, with the use of templating it really isn't. However, due to this design it is possible to create a very flexible setup.

5 What's new in Zabbix 2.0.0

5.1 Major new features 5.1.1 Low level discovery

Low level discovery documentation

5.1.1.1 Network interface discovery

Network interface discovery support has been added for the following platforms:

- Linux
- Windows
- Solaris
- HP-UX
- AIX

5.1.1.2 Mounted filesystem discovery

Mounted filesystem discovery returns filesystem names along with their types.

Network interface discovery support has been added for the following platforms:

- Linux
- Windows
- Solaris
- HP-UX
- AIX
- FreeBSD
- OpenBSD
- NetBSD
- MacOS X

It is possible to filter by regular expression on the filesystem name or type.

5.1.2 Automatic inventory data collection

Host inventory now can be populated **automatically**. Items have a property that allows to store their value in any of the inventory fields if inventory is set to *Automatic* mode.

5.1.3 WEB monitoring by Zabbix proxies

WEB scenarios could be monitored by remote proxies. It allows monitoring of WEB applications from various geographical locations. Also new items "*Last error message of scenario '\$1'*" with key *web.test.error[<a scenario name>]* will be added to all existing scenarios during database upgrade.

5.1.4 Multi-homed host support

Zabbix now supports multiple interfaces per host. You can configure several interface-host pairs in the host properties, and specify which should be used as the default for several item types (Zabbix agent, SNMP and IPMI). Additionally, you can override interface on the item level. When a template is linked to a host that has multiple interfaces defined, default interface for that type used (which currently is the first interface in the list). A single item can only be monitored on a single interface.

Interfaces that are used in any items can not be removed and link *Remove* is greyed out for them.

Together with the ability to use user macros in any of the port fields this should provide a flexible and powerful way to monitor multi-homed hosts.

For SNMP items, port on the item level, if specified, overrides the interface port.

Interface for an item can be selected in item properties. Both in item properties and item mass update form interfaces are grouped by their type.

5.1.5 Direct JMX monitoring support

Ability to monitor JMX without deploying additional software has been introduced in the form of a **Java gateway**.

5.1.6 Visible host name

Visible host name has been implemented; UTF8 support; new macros ({HOST.HOST}, {HOST.NAME}, {HOST.IP}); {IPADDRESS}, {HOSTNAME} deprecated

5.1.7 SNMP trap monitoring

SNMP trap item documentation.

5.2 Frontend improvements 5.2.1 A new look

Zabbix frontend has received a redesign. The login page, table view, configuration wizard, default trigger severity colours all have a fresh new look.

5.2.2 Reworked frontend

The frontend has been changed to be more user friendly, especially on smaller screens. Many pages that previously required horizontal scrolling now fit in the screen thanks to improved layout and tabbed interface.

The following pages have been redone:

- Host and template properties
 - Host properties now also include host navigation interface, making it easy to jump from host properties to editing items, triggers, graphs or other entities for that host (screenshots to compare 1.8 vs 2.0)
- Proxy properties
- Item properties
- Trigger properties
- Action properties
- Graph properties
 - Graph item reordering now isn't done by repeated clicks on *Up* and *Down* links. Instead, drag'n'drop reordering is possible.

5.2.3 Dropping IE6 support

Support for Microsoft Internet Explorer 6 will not be provided anymore. While we were at that, we also dropped IE7 support.

5.2.3.1 Outdated browser warning

If an outdated browser will be detected by Zabbix frontend, a warning page will be displayed. It will link to download pages for multiple more recent browsers. User would not be blocked from accessing Zabbix frontend, though - that would still be possible via the link in the lower right corner.

Zabbix frontend is built on advanced, modern technologies and does not support old browsers. It is highly recommended that you choose and install a modern browser. It is free of charge and only takes a couple of minutes.



[Google Chrome
Download page](#)



[Mozilla Firefox
Download page](#)



[Internet Explorer
Download page](#)



[Opera browser
Download page](#)



[Apple Safari
Download page](#)

Why is it recommended to upgrade the web browser?

New browsers usually come with support for new technologies, increasing web page speed, better privacy settings and so on. They also resolve security and functional issues.

[Continue despite this warning »](#)

Zabbix 1.9.7 Copyright 2001-2011 by Zabbix SIA

5.2.4 SQLite supported by frontend

5.2.5 Improved frontend performance

Multiple sections of frontend have their performance improved, including:

- Dashboard performance has been improved for systems with lots of events.
- Network maps for non-superusers
- Screens for non-superusers
- IT services monitoring and configuration sections
- Availability report

5.2.6 New network map icons

For Zabbix 2.0, a completely new set of network icons is provided, and several new icons have been introduced.

<list of icons>

5.2.7 More intuitive range selecting

Previously, selecting a range of entities in Zabbix frontend was not very intuitive - one had to hold down Ctrl both for the range start and range end. Starting with 2.0, range selection works as in most other software. One just clicks on a checkbox that should start range selection, then, holding down shift, clicks on the checkbox that should be the end for the range. It works the same way for selecting and unselecting (taking the choice from the very first click).

5.2.8 Customisable dashboard

Dashboard elements can now be arranged in a custom layout by dragging and dropping them in the desired location. They can be placed in up to 3 columns.

5.2.9 Translations managed by gettext

Frontend translations are now managed by gettext. This means that standard tools like [Pootle](#), [Poedit](#) or [Virtaal](#) can be used to translate Zabbix frontend. That should also allow working on smaller sections of the translation and improve collaboration possibilities.

5.2.10 Network map improvements

Network maps have received a couple of improvements regarding URL handling.

**** Multiple URLs for map elements ****

In previous Zabbix versions, it was possible to add single URL for each network map element. Now it is possible to add multiple, while also giving a label for each. This can be done on two levels:

- Map element level
- Map level

For example, here two URLs have been defined on the element (host) level:

Edit map element

Type

Host

Label

Riga

Label location

Bottom

Host

A Test Host

Select

Icon (default)

Server

Use advanced icons

☐

Coordinate X

195

Coordinate Y

175

Links

Name	URL	
Wiki page	http://wiki.internal/Riga	Remove
Web management interface	http://10.10.10.10/manage	Remove

Add

Apply

Remove

Close

On the map level, URLs can be defined for all element types. Here, URL is defined for each available element type:

System map: "Eurasia-UTF8" ?

Name: Eurasia-UTF8

Width: 932

Height: 573

Background image: eurasia

Icon highlighting: ☒

Mark elements on triggers status change: ☐

Expand single problem: ☒

Icon label type: Label

Icon label location: Left

Problem display: All

Links


Name	URL	Element	
Group graphs	charts.php?&hostid=	Host group	Remove
Host graphs	charts.php?&hostid=	Host	Remove
Some link for all images	http://wiki/about_imag	Image	Remove
Some link for all maps	http://wiki/about_map	Map	Remove
Trigger events	events.php?triggerid=	Trigger	Remove

[Add](#)

Save Delete Cancel

Of course, it is also possible to add multiple URLs for the same element type.

When these URLs have been configured, in the monitoring section clicking on a map element will list them (by name) in the popup menu. If the element has its own URLs and there are map level URLs for its type defined, they will be combined in the same menu:



From the already existing macros, in map URLs macro {TRIGGER.ID} can be used.

Additionally, 3 new macros have been added that also can be used in element URL fields:

- {HOST.ID}
- {HOSTGROUP.ID}
- {MAP.ID}

**** Grid selection saving ****

In previous versions, map grid options like whether to display grid or align to it were discarded after the map configuration form

was closed. Starting with 2.0.0, selected grid options are saved for each map individually.

** Setting label type per map element type **

In previous versions, one could select map label type for all map elements in the map. Starting with 2.0.0, it is possible to enable *Advanced labels* feature per map and choose map element label per element type - image, trigger, host, host group and map.

Advanced labels	☒
Host Group label type	Label
Host label type	Label
Trigger label type	Element name
Map label type	Status only
Image label type	Nothing
	Custom label
	Label

** Displaying hosts from a host group **

For host group element, it is now possible to specify that it should display all hosts in a group. In that case, area can be specified where hosts will be placed.

** Icon mapping **

Automatic icon selection for host type map elements has been added. It utilises values from host inventory fields to select icon for map elements.

Icon maps are configured in *Administration* → *General* → *Icon maps*. Each icon map has a unique name, default icon and a list of mappings. Each mapping has a match order, expression for matching (PCRE regular expression or Zabbix expression), icon and reference to an inventory field. Icon map can be assigned to a map in the map configuration form, and overridden on map element level.

When deciding on which icon should be used, regular expression will be matched against specified inventory field. Order in which mappings are tried depends on match order.

** Select elements by dragging **

You can now drag a rectangle in map configuration view and select all elements in it.

** Highlighting selected link **

Selected map link now is highlighted in the list of links.

** Changed element highlighting **

Instead of filled rectangle, to indicate selection elements are now highlighted with a thick rectangular border.

** Element properties window changes **

Map element properties window has received multiple modifications.

- Links are only shown when one element is selected
- List of selected elements is displayed above properties themselves (was displayed below before)
- When editing multiple elements, properties that are not marked for updating using the checkbox on the left hand side are not editable
- Only one endpoint of a link may be changed when editing a link. Endpoint that is currently selected can not be changed in this view anymore. Additionally, endpoints are sorted by type in the dropdown

Host	SNMP Device
	Another Host
	A Test Host
	A Test Host 2
	autoreg_2
	HP laserjet
	IPMI Host
	A Test Host
	ackmacros
Map	booklet
	Drag and drop map
Trigger	A Test Host:CPU Load too high on A Test Host for last 3 minutes & <> \$#@!%^*()00;.:',./
	A Test Host:Severity 1 missing on A Test Host
	A Test Host:Severity 2 missing on A Test Host
	A Test Host:Severity 3 missing on A Test Host
Host group	Server

5.2.11 Customisable trigger severity names and colours

Trigger severity names and colours can now be easily customised in *Administration* → *General* → *Trigger severities*.

5.2.12 Customisable problem displaying

The way triggers are displayed in the frontend can now be customised in *Administration* → *General* → *Trigger displaying options*. That includes both OK/PROBLEM and acknowledged/not acknowledged.

Additionally, time for how long OK triggers are shown after changing state and for how long trigger is blinking after changing state can now be configured in the frontend, instead of modifying `include/defines.inc.php` file.

5.2.13 Inventory report

Ability to report on hosts by their gathered inventory data has been added to *Monitoring* → *Inventory* view.

5.2.14 Event export to CSV

Events from *Monitoring* → *Events* page can now be exported to a CSV file by using a button in the upper right corner.

Displaying 1 to 7 of 7 found

Export to CSV

Filter

Time	Host	Description	Status	Severity	Duration	Actions
24 Jan 2011 18:28:19	A Test Host	Apache error "1"	PROBLEM	Not classified	19h 24m 31s	-
24 Jan 2011 18:11:18	A Test Host	Apache error "1"	PROBLEM	Not classified	17m 1s	-
24 Jan 2011 15:18:35	A Test Host	Apache error "1"	PROBLEM	Not classified	2h 52m 43s	-
23 Jan 2011 21:25:27	A Test Host	Apache error "1"	PROBLEM	Not classified	17h 53m 8s	-
20 Jan 2011 22:59:13	A Test Host	Apache error "1"	PROBLEM	Not classified	2d 22h 26m	-
20 Jan 2011 19:29:06	A Test Host	Apache error "1"	PROBLEM	Not classified	3h 30m 7s	-
20 Jan 2011 15:40:17	A Test Host	Apache error "1"	PROBLEM	Not classified	3h 48m 49s	-

Zabbix 1.9.2 Copyright 2001-2010 by SIA Zabbix

Connected as 'Admin'

Only currently visible events are exported.

5.2.15 User executable script improvements

Confirmation for user executable scripts

User scripts that can be configured on the frontend in *Administration* → *Scripts* can now display a confirmation message before being executed. A custom message confirmation message can be entered by the administrator. This feature should be especially useful with potentially dangerous operations (like a reboot script) or ones that might take a long time.

Enable confirmation ☒

Confirmation text [Test confirmation](#)

Confirmation can be tested in the configuration form. In that case, *Execute* button is disabled.

Execution confirmation

Remote system will be rebooted. Are you sure?

Execute

Cancel

Ability to specify script type

Before, if a user script would have to be an IPMI command, it had to follow a syntax of `IPMI <command>`. Starting with Zabbix 2.0, it is an option in the script editing form.

Type

Script

IPMI

Script

Ability to execute scripts on Zabbix agent

Before Zabbix 2.0, user scripts could be executed on the Zabbix server only. Starting with 2.0, a script can be configured to be executed on the agent instead. For this feature to work, agent must be **configured to allow remote commands**.

Execute on

☐ Zabbix agent
☒ Zabbix server

5.2.16 Media type improvements

Displaying of actions that use media types


In media type configuration screen it's now possible to view and access actions that directly refer to this media type (it's selected in the *Send only to* dropdown). When sorting by that column, entries are only sorted by whether they are referenced, actions referencing media types do not influence the sort order. It is possible to go directly to action configuration from this page as well.

MEDIA TYPES			
Displaying 1 to 5 of 5 found			
☐	Description	Type	Used in actions
☐	SMS	SMS	-
☐	ccc	Email	action_no_user , integrity
☐	Jabber	Jabber	a
☐	Email	Email	some action , Test Action
☐	Test Script	Script	SNMP Action
Details			
GSM modem: "/dev/ttyS0"			
SMTP server: "localhost", SMTP helo: "localhost", SMTP email: "zabbix@localhost"			
Jabber identifier: "jabber@company.com"			
SMTP server: "mail.zabbix.com", SMTP helo: "zabbix.com", SMTP email: "zabbix@company.com"			
Script name: "testscript"			

Enabling and disabling media types

5.2.16.1 Disabling graph legend

For graphs with many items added, legend could grow pretty huge, and there might be other reasons why somebody might want to hide the legend. Starting with Zabbix 2.0.0, that is possible for custom graphs with a simple checkbox:



5.2.17 Improved popup placement

Popups with additional information that appear on mouse roll-over are designed to better fit within smaller screens, such as in mobile devices.

5.2.18 Improved IT services view and configuration

monitoring - shows 80-100% upon mouseover (and a tooltip), reordered sla values, aligned by decimal point, removed graph "show" link (now it's on the bar, link to report is on the service name), expected sla shown in black not green, more...


multiple service dependencies can be selected now

5.2.19 Templated host screens

Screens can be defined in templates. They are available for all hosts that are linked to the template. They can be accessed from the host popup menu everywhere, and from the search results.

5.2.20 Ability to sort triggers in screens

Triggers element (by host or by host group) in screens now can be sorted by severity, host or date.



<https://support.zabbix.com/browse/ZBXNEXT-746>

In the monitoring section, this sorting can only be viewed but not changed. Additionally, in the monitoring section this element shows how many triggers are visible, if not all of them can be shown.

5.2.21 Warning when a too large file is attempted to be uploaded

Previously it was not easy to debug why some upload failed if it exceeded PHP option max_upload_filesize limit. Now Zabbix shows a warning when uploading a too large XML or image file.

5.2.22 Improved configuration wizard

For configuration wizard requirements step, *Recommended* column was removed, as it did not serve any real purpose. Additionally, for failed PHP configuration requirements exact configuration parameters are now listed in the error messages.

Better text, visuals, removed 2nd and 8th steps

5.2.23 Trigger config helper

Trigger configuration helper does not allow setting parameter in seconds for trigger function last() anymore

5.2.24 Improved search results page

In the search results page, individual blocks have been placed below each other. They can also be collapsed and expanded now.

5.2.25 Host popup menu added to more pages

Host popup menu has been added to the *Monitoring* → *Events*, "Event details", and *Monitoring* → *Overview* pages. In *Monitoring* → *Overview* this menu is available only when hosts are displayed on the left hand side.

5.2.26 New option "Expand macros" for maps configuration screen

In 2.0 macros expanding is turned off by default and new option is introduced for map configuration screen that allows to enable or disable it. Maps that utilize a lot of macros should benefit greatly from turning macro expanding off performance wise.

5.2.27 Improved debug layout

In the frontend debug functionality, SQL query formatting has been improved for better readability.

5.2.28 Allow deleting trigger dependencies during mass update

If no trigger dependencies are added and their controlling checkbox is marked in the trigger mass update, all dependencies for selected triggers will be removed now.

5.2.29 New and updated translations

5.2.29.1 New translations

- Turkish

5.2.29.2 Updated translations

- Czech
- Dutch
- French
- Greek
- German
- Japanese
- Latvian
- Polish
- Ukrainian
- Russian
- Spanish

5.3 New and improved items 5.3.1 Item descriptions added

...and 'description' is now 'name'

Added {ITEM.DESCRPTION} and {ITEM.KEY} macros.

5.3.2 Improved items

system.localtime key now supports parameters:

- `utc` (default) - returns Unix timestamp like before;
- `local` - returns time in the 'yyyy-mm-dd, hh:mm:ss.nn, +hh:mm' format

net.tcp.dns and **net.tcp.dns.query** have been renamed to **net.dns** and **net.dns.record**, respectively. Old syntax still supported as an alias. They are also now supported on Windows.

Additional functionality:

- IP as the first parameter is not ignored anymore, thus it is possible to specify different DNS server than localhost
- Both items gained ability to specify timeout and retry count (not supported on Windows)
- **net.dns** gained support for specifying DNS record type like `A`, `CNAME` and `MX` (**net.dns.record** already supported that). See [item documentation](#) for more information

net.tcp.service and **net.tcp.service.perf** now support https and telnet.

system.cpu.util on Solaris was unified with other platforms - *kernel* parameter was changed to *system* and *wait* parameter was changed to *ioawait*.

system.cpu.num is now supported for OSX

system.cpu.load now accepts *percpu* as the first parameter - if specified, total load divided by online CPU count is returned (default is still *all* and returns normal load value)

vm.memory.size now supports more modes for various platforms. See [supported by platform](#) and [vm.memory.size parameters](#) pages for more information.

5.3.2.1 Skipping older data in log monitoring

added optional suppressing of processing of older data in `log[]`, `logrt[]` and `eventlog[]`

additional parameter **mode** controls this:

- `all` (default)
- `skip` (skipping processing of older data)

5.3.2.2 Simple check format changed

The format of **simple checks** has been changed. The possibility of using IP addresses in simple checks has been added. The upgrade process converts all existing simple checks.

5.3.2.3 Ability to send empty strings

User parameters, external checks and other item types that return text (character, log, text types of information) now may send an empty string which will be accepted as a valid value. Before this, an empty value returned by an user parameter or external check would be treated as an unsupported item, and ignored for other item types by the Zabbix server.

5.3.2.4 External check parameter changes

External check parameter handling was changed. Previously, only one parameter was accepted. Starting with Zabbix 2.0, the key syntax conforms to other types of items and multiple comma-separated parameters may be passed.

Additionally, the hardcoded first parameter {HOST.CONNN} for external checks has been removed and now external check parameters may be fully customised. Upgrade process adds this macro as a first parameter to all existing external checks.

5.3.2.5 Multi-line value support for external checks

Multi-line values can now be stored for the output of external scripts.

5.3.2.6 Time postfix support in aggregate items

Aggregate items now support time postfixes.

5.3.3 New items

- **vfs.file.contents[]** allows to retrieve full file contents with each check (limited to 64KB, will turn unsupported if file is bigger)

Items that are especially useful for automated inventory data collection:

- system.hw.chassis[full|type|vendor|model|serial] - default is [full], root permissions needed
- system.hw.cpu[all|cpunum,full|maxfreq|vendor|model|curfreq] - default is [all,full]
- system.hw.devices[pci|usb] - default is [pci]
- system.hw.macaddr[interface,short|full] - default is [all,full], interface is regex
- system.sw.arch
- system.sw.os[name|short|full] - default is [name]
- system.sw.packages[package,manager,short|full] - default is [all,all,full], package is regex

A **status** item has been removed as it was not easy to use and understand. Instead, a new **internal item** - **zabbix[host,<type>,available]** has been added. It allows to properly monitor host availability by each of states - Zabbix agent, SNMP, IPMI, JMX - and returns values for available, unavailable and unknown. The database upgrade process **doesn't convert** all existing 'status' keys and associated triggers to the new syntax, so they have to be converted manually.

5.4 Daemon performance improvements 5.4.1 Improved housekeeper performance on PostgreSQL

5.4.2 Caching of general configuration

Multiple general configuration parameters that previously were only read upon server startup are now also stored in the configuration cache. They are updated together with common configuration cache update.

- *Refresh unsupported items*
- *Discovery group*
- Custom severity names
- For how long alerts are stored
- For how long events are stored

Refresh unsupported items is now also synchronised to the Zabbix proxy.

5.4.3 Trigger cache

Trigger cache has been introduced. Currently it holds the following information about triggers:

- description
- expression
- error
- severity
- type
- value
- value_flags

Trigger cache is part of the configuration cache and is updated the same way and at the same time as other parts of it. Note that upgrading to 2.0.0 may increase configuration cache usage, especially if Zabbix instance has lots of triggers.

This should improve trigger related operation performance and reduce the load on the database.

5.4.4 User macro caching improvements

In Zabbix 1.8, individual daemon processes (for example, pollers) cache user macros and update this information every 5 minutes. In Zabbix 2.0, user macros are included in the general configuration cache and updated every **CacheUpdateFrequency** seconds.

5.4.5 Improved history (DB) syncer and escalator performance

The amount of work history syncer and escalator processes have to do for escalations has been reduced, which resulted in improved performance when processing large amount of events.

5.5 Other improvements

5.5.1 Nanosecond support

5.5.2 Database integrity

5.5.2.1 Database creation

Attention:

Order of **database file import** has been changed - images must be imported before **data.sql** now.

5.5.3 Improved active agent auto-registration

Active agent auto-registration now supports monitoring added hosts with passive checks. When active agent asks for checks, if it has configuration parameters ListenIP or ListenPort defined in the configuration file, these are sent along to the server. If multiple IP addresses are specified, the first one is sent to the server.

Server, when adding a new auto-registered host, uses received IP address and port to configure the agent. If no port value is received, 10050 is used. If no IP address value is received, the one used for incoming connection is used.

For active agent auto-registration notifications there are also two new macros supported:

- {HOST.IP} (or its deprecated version {IPADDRESS})
- {HOST.PORT}

5.5.4 Automatic filling of DNS names

For network discovery and active agent autoregistration, host DNS entry is now populated by doing a reverse lookup from Zabbix server or Zabbix proxy, if discovery is performed by a proxy.

5.5.5 Event redesign

The way Zabbix handles unknown events has been redesigned. Before, in *Monitoring* → *Events* section, if event flow was OK → UNKNOWN → OK, hiding unknown events would display two successive OK events. Starting with 2.0.0, only the oldest OK event would be display in such a case. An exception is triggers with "Multiple PROBLEM events" - for these, multiple successive events would be always shown.

(should document how acknowledges work from server & frontend side)

Trigger status UNKNOWN has been removed - triggers will always be in OK or PROBLEM state, although they might get additional unknown status. This status will only be used for informational purposes.

5.5.6 Supported macro changes

5.5.6.1 Common macros

In interface IP and DNS fields the following macros are supported:

- {HOST.HOST} (or its deprecated version {HOSTNAME})
- {HOST.NAME}
- {HOST.CONN}
- {HOST.DNS}
- {HOST.IP} (or its deprecated version {IPADDRESS})

In trigger names, additional following macros are supported:

- {HOST.HOST} (or its deprecated version {HOSTNAME})
- {HOST.NAME}
- {HOST.CONN}
- {HOST.DNS}
- {HOST.IP} (or its deprecated version {IPADDRESS})

5.5.6.2 User macro support

For 2.0, user macros can now be used in these additional locations:

- Hosts:
 - Interface IP/DNS
 - Interface port
- Passive proxy
 - Interface port
- Items
 - SNMPv3 security name
 - SNMPv3 auth pass
 - SNMPv3 priv pass
 - SNMPv1/v2 community
 - SNMPv1/v2 OID
 - DB monitor additional parameters
 - SSH username
 - SSH public key
 - SSH private key
 - SSH password
 - SSH script
 - Telnet username
 - Telnet password
 - Telnet script
 - Calculated item expression
- Discovery
 - SNMPv3 security name
 - SNMPv3 auth pass
 - SNMPv3 priv pass
 - SNMPv1/v2 community
 - SNMPv1/v2 OID

Additionally, macro **{TRIGGER.KEY<1-9>}** has been renamed to **{ITEM.KEY<1-9>}**.

5.5.7 Boolean item type

New item type has been added - boolean. It translates textual representation into either 0 or 1. All values are matched in a case-insensitive way. **TRUE** is stored as **1** and **FALSE** is stored as **0**. Currently recognised values are:

TRUE	FALSE
true	false
t	f
yes	no
y	n
up	down
running	unused
enabled	disabled
available	unavailable

Additionally, any non-zero numeric value is considered to be **TRUE** and zero is considered to be **FALSE**.

5.5.8 API access changes

User group access property "API access" has been removed. Starting with Zabbix 2.0, all users have full access to the API.

5.5.9 Unified time format

In previous versions of Zabbix, specifying time till midnight differed depending on the location - in some locations it was "23:59", in others - "24:00". Starting with Zabbix 2.0 this has been unified so midnight is always specified as "24:00".

5.5.10 "Host profiles" renamed to "host inventory" and merged

Before, Zabbix hosts had "normal" and "extended" profile. In Zabbix 2.0 both profiles were merged into single "host inventory".

5.5.11 Escalations enabled by default

Starting with Zabbix 2.0, escalations are always enabled and can not be disabled. By default a condition "Trigger value = PROBLEM" is added to new actions, so only problem notifications are sent. This makes recovery message checkbox operate in a more intuitive way. This action condition can be removed by the user.

5.5.12 Remote command improvements

Specify IPMI as a type instead of obscure syntax

Execute custom global scripts

Execute SSH and telnet commands

For custom scripts, ability to execute on the Zabbix server or agent was added

5.5.13 Changed default locations

5.5.13.1 Configuration files location

Before 2.0 this was hard-coded to `/etc/zabbix`. Since 2.0 the default changed to `/usr/local/etc` and depends on compile-time installation variable `sysconfdir` (which was ignored before Zabbix 2.0).

5.5.13.2 External scripts location

Before 2.0 the default location for external scripts was `/etc/zabbix/externalscripts`. Since 2.0 this changed to `/usr/local/share/zabbix/externalscripts` and depends on compile-time installation variable `datadir`. This location can be changed using configuration parameter `ExternalScripts` (server and proxy).

5.5.13.3 Alert scripts location

Before 2.0 the default location for alert scripts that are used for actions was `/home/zabbix/bin`. Since 2.0 this changed to `/usr/local/share/zabbix/alertscripts` and depends on compile-time installation variable `datadir`. This location can be changed using configuration parameter `AlertScriptsPath` (server).

5.5.14 Configuration files and directories added during installation

Attention:

If you are using a packaged version of Zabbix the following information can be irrelevant because the locations might depend on a distribution policy.

When installing Zabbix the following configuration files are created if missing:

- `/usr/local/etc/zabbix_server.conf`
- `/usr/local/etc/zabbix_proxy.conf`
- `/usr/local/etc/zabbix_agent.conf`
- `/usr/local/etc/zabbix_agentd.conf`

Note:

`/usr/local/etc` is default value of compile-time installation variable `sysconfdir`.

The following directories are created for includes:

- `/usr/local/etc/zabbix_server.conf.d`
- `/usr/local/etc/zabbix_proxy.conf.d`
- `/usr/local/etc/zabbix_agent.conf.d`
- `/usr/local/etc/zabbix_agentd.conf.d`

Note:

These directories are not automatically handled as "Include" directories, you would have to set them in appropriate configuration file.

The following directory is created when installing server or proxy:

- `/usr/local/share/zabbix/externalscripts`

The following directory is created when installing server:

- `/usr/local/share/zabbix/alertscripts`

Note:

`/usr/local/share` is default value of compile-time installation variable `datadir`.

5.5.15 Unsupported item information sent by proxies

Information about unsupported items is now passed from Zabbix proxies to the server.

5.5.16 Simplified ServerActive logic

Zabbix agent daemon parameter **ServerActive** has been simplified. Instead of overriding **Server** parameter, **ServerActive** now is the only parameter that's used for active checks. Additionally, **ServerPort** parameter has been removed - host ports now should be specified in the **ServerActive** parameter.

Note:

See the "[See also](#)" section at the bottom of this page to read more details about these changes.

5.5.17 Simplified Zabbix agent daemon configuration

Zabbix agent daemon configuration has been simplified by removing **DisableActive** and **DisablePassive** parameters. Instead, not specifying **ServerActive** will not do any active item processing, and specifying **StartAgents=0** will disable processing of incoming connections for passive items.

5.5.18 Monitoring of files larger than 2GB

While Zabbix did offer monitoring of files larger than 2GB before, this support was not consistent - it varied by platform and by item. Since 2.0, support for files exceeding 2GB is available also on 32-bit systems in all components, including tasks like logfile monitoring, file size monitoring etc. See [large file support](#) section for more information.

See also

1. [Differences in the Zabbix agent configuration for active and passive checks starting from version 2.0.0](#)


6 What's new in Zabbix 2.0.1

6.1 Frontend improvements 6.1.1 Warning message about server being down

A warning message will be shown in the frontend if Zabbix server is not reachable (for example, not running). This message can be configured on the Zabbix installation level in *Administration* → *General* → *GUI*.

Show warning if Zabbix server is down ☒

When server has been detected to be down, a message will be displayed at the top of the browser window. This message will be visible even if user has scrolled down in the page.



If mouse cursor is moved over the message, it is hidden to reveal contents below.

6.1.2 Improved availability displaying in host properties

Displaying of host availability in host properties was improved. Instead of single word describing only agent availability, 4 icons were added, same as in the list of hosts. Added bonus is that they now take less space than previous textual representation.

|<|<|<|<|

|<|<|<|<|

6.1.3 Miscellaneous frontend improvements

Action properties (conditions and operations) have been updated to use direct *Remove* links instead of checkboxes and buttons.

XML import now accepts all valid empty tags in all locations, previously it required specific empty tag syntax depending on the location. Since 2.0.1 both `<tag></tag>` and `<tag/>` may be used.

6.1.4 Updated translations

- Brazilian Portuguese
- French
- Greek
- Japanese

- Ukrainian
- Russian

6.2 Daemon improvements

- Zabbix Windows agent previously included the trailing backslash in disk drive discovery - for example, the returned values would include `C:\`. Since 2.0.1, the returned values have the trailing backslash stripped (for example, `C:`). This allows to use these values in Windows performance counters as well.
- Previously, on FreeBSD a monitored disk device had to be specified without `/dev` prefix. Now the prefix may be optionally used.
- Previously, Zabbix agents were limited to monitoring 8 disks or partitions on a system. This limit has been increased to 1024 now. Note that monitoring a large amount of disks will increase the memory usage.
- Previously, Zabbix agents allocated a static amount of memory for monitoring 8 disks or partitions. This amount is now dynamic and is allocated only if monitoring of a specific disk or a partition is requested, thus reducing the memory usage on systems where less than 8 disks are being monitored. If no disks are being monitored, Zabbix agent memory usage will decrease by approximately 400 KB. Memory is allocated incrementally as more disk monitoring is being requested. The steps for this increase and approximate memory usage (for disk monitoring only, in MB) are as follows:

Disks	0	1	2	3	4	8	16	32	64	128	256	512	768	1024
Memory	0	0.05	0.1	0.15	0.2	0.4	0.8	1.5	3.1	6.2	12.4	24.8	37.2	49.6

7 What's new in Zabbix 2.0.2

7.1 Frontend improvements

- Filters in host, host inventory and audit pages were improved to fit in screens, less than 1024 pixels wide
- Performance in graph related pages was improved
- Performance in the dashboard's "Host status" widget was improved, especially on setups with many groups
- Global messages previously appeared in all pages, including popups (for example, for item or host selection). Now they will only appear in pages that would also display the main menu
- In general action properties, "Default escalation period" was renamed to "Default operation step duration"
- In action operations, "Escalation period" was renamed to "Step duration". "Period (sec)" and "Delay" columns were renamed to "Duration (sec)" and "Start in" respectively, and also swapped places.
- Frontend will now validate time periods entered in item flexible intervals, working time definition and other locations
- More XML import status messages can now be translated
- More specific reasons will be given when template linking fails in some cases
- Item value macros in map labels now support multiplier suffixes (for example, "{host:item.function(5m)}")
- In the global search results, host and template configuration links for host group results are now separated. Previously, only link to host editing would be displayed, and no indication of how many hosts are there would be visible. Since 2.0.2, configuration section of search results has been expanded with templates entry, and both host and template columns have object count as well. This is also more consistent with other blocks in the global search results

Before 2.0.2

Host groups				
Host group	Latest data	Triggers	Events	Edit hosts
Zabbix servers	Latest data	Triggers	Events	Edit hosts
Displaying 1 of 1 found				

Since 2.0.2

Host groups					
Host group	Latest data	Triggers	Events	Hosts	Templates
Zabbix servers	Latest data	Triggers	Events	Hosts (2)	Templates (0)
Displaying 1 of 1 found					

- *Change in Monitoring* → *Latest data* for **unixtime** units has been improved. Instead of interpreting change using the same unit as for the item, it will now use **uptime** unit

Before 2.0.2

Since 2.0.2


Last value	Change	Last value	Change
2012.07.04 14:12:56	+1970.01.01 03:00:30	2012.07.04 14:13:56	+00:00:30

7.1.1 Updated translations


- French
- German
- Japanese
- Russian

7.2 Daemon improvements

- Previously, HOST.CONN, HOST.IP and HOST.DNS macros only resolved to the main agent interface. This was not convenient for devices that only had SNMP interface. Since 2.0.2 these macros will use the main agent interface, however, if it is not present, main SNMP interface will be used. If SNMP is also not present, main JMX interface will be used. If JMX is not present either, main IPMI interface will be used.
- Windows agent binaries will now include version information resource. For example, on Windows 2008 the improvement will look like this:



File information in 2.0.1



File information in 2.0.2

7.3 Miscellaneous improvements

- Zabbix sender manpage now explains which parameters are supported from the agent configuration file
- Item reference for frontend now includes all items and its syntax matches the syntax used in the manual

8 What's new in Zabbix 2.0.3

8.1 Frontend improvements 8.1.1 Flicker-free screens

Previously an automatic page refresh or change in the timeline scrollbar in the monitoring section reloaded the whole page. This was especially visible in **screens**, where it may take some time for all screen elements to load properly, and the result often was a page where elements jumped around for a while.

Since Zabbix 2.0.3, screens in monitoring do not refresh the whole page, but all the individual elements are reloaded in the background and replaced.

Additionally, image background reloading was implemented in *Latest data*, graphs and slideshows. This should result in a better user experience and reduced network traffic for all of these pages.

8.1.2 Displaying of multiple map links

In previous 2.0 versions, map links were not available for editing when multiple elements were selected. Since 2.0.3, it is possible as follows:

- If a single map element is selected, all links that are connected to it are listed
- When multiple map elements are selected, links that connect any two of them are listed

Links between the selected elements			
	From	To	Link indicators
Edit	Server	Hub (small)	
Edit	Server	SNMP Device	A Test Host:clamd is not running A Test Host:WEB service is down
Edit	SNMP Device	Hub (small)	

8.1.3 Filtering triggers by maintenance status

It is possible to filter *Monitoring* → *Triggers* page by maintenance status now.

Filter

Triggers status

Problem

Acknowledge status

Any

Events

Hide all

Min severity

All

Age less than

☐ 14 days

Show details

☐

Filter by name


Show hosts in maintenance

☒


FilterReset

8.1.4 Improved classic theme tab colours


Previously, classic theme tab colours were slightly confusing. The active tab had grey background, which made it look inactive, while other tabs had light background. This was changed so that the active tab now has a light background, while the inactive ones - blue. Active tab also lost the lower border, so that it indeed looks like one with the visible tab.



Host properties before 2.0.3



Previous theming was even more confusing when there were two tabs only, and it was not that clear which is the active one.



8.1.5 Maintenance indicated by an icon

Previously, maintenance being active was indicated:

- in *Monitoring* → *Triggers* - with a string, printed next to the host
- in *Monitoring* → *Dashboard* - by colouring host name in orange

In the triggers page this was increasing the cell width considerably, and in the dashboard it did not allow to nicely provide access both to host context menu and maintenance details. This was changed in both locations to use a new icon - orange, round background with white wrench on it.

|<|<|<|<|

|<|<|<|<|

8.1.6 Miscellaneous frontend improvements

- In item mass update, field length limits have been extended.
- In availability report, dropdown labels have been renamed from *Template* and *Template group* to *Host* and *Host group*
- Map links in the editor will now be sorted first by *From*, then by *To* element. If a single element is selected, it will be used as *From* for all links. If multiple elements are selected, for each link the element with the lower **selementid** will be used for *From*.
- Performance for the screen, host screen and graph sections was improved.
- All frontend problems visible in Internet Explorer with "Compatibility mode" enabled have been solved

8.1.7 New translations

- American English
- Chinese (Taiwan)

8.1.8 Updated translations

- Brazilian Portuguese
- Chinese (China)
- French
- German
- Japanese
- Latvian
- Spanish
- Russian
- Ukrainian

8.2 API improvements

- **trigger.get** and **triggerprototype.get** methods have a new parameter **expandExpression**
- **maintenance.get** method has a new parameter **selectTimeperiods**

8.3 Daemon improvements

- Notification macro **{ESC.HISTORY}** will include remote commands now
- Suffixes "KMGtSmhdw" are now supported in the second parameter of count() trigger function
- Previously it was possible to send large amount of data to the Zabbix server, potentially exhausting memory. This is now limited to accept only 128MB when using Zabbix protocol. Any other data (including older Zabbix protocols) stays limited at 16MB
- Zabbix server and proxy will now retry connections to MySQL if connecting to a DNS name and resolution fails
- When encoding email subject from UTF-8 to Base64, lines longer than 75 bytes will be split up to conform with RFC-2047

8.4 Miscellaneous improvements Database patch performance for upgrading 1.8 to 2.0 was significantly improved, especially for systems with lots of events.

9 What's new in Zabbix 2.0.4

9.1 Frontend improvements 9.1.1 Shading of screen elements that failed to update

Since 2.0.3, screen elements are updated in background, without reloading the whole page. Zabbix 2.0.4 will darken elements that have failed to update for some reason.

9.1.2 Miscellaneous frontend improvements

- Full cloning a template will now clone templated screens as well.
- It is now possible to have empty recovery message body.
- Previously, if a trigger had no events, it was shown as acknowledged in *Monitoring* → *Triggers*. Now a string **No events** will be shown instead.
- Page refresh should now reload less files, thus improving overall frontend performance and decreasing the network traffic.
- In network maps, **HOST.HOST** macro will now work in trigger elements.
- Screen editing will not allow to set row or column span that would exceed row or column count.
- In *Monitoring* → *Overview* vertically displayed text that used to be made up by images, now is actually vertical text.

9.1.3 New translations

- Finnish
- Indonesian
- Persian

9.1.4 Updated translations

- Brazilian Portuguese
- Chinese (China)
- Chinese (Taiwan)
- Czech
- French
- German
- Hungarian
- Italian
- Japanese
- Latvian
- Polish
- Slovak
- Russian
- Ukrainian

9.2 API improvements Method **templatescreen.copy** has been added.

9.3 Daemon improvements Macro **{TRIGGER.COMMENT}** has been renamed to **{TRIGGER.DESCRPTION}**, to match the field name. Old macro is still supported.

Default disks checked have been improved for vfs.dev.read and vfs.dev.write items when using Zabbix agent daemon with -t or -p flags. On Linux, FreeBSD, OpenBSD and Solaris disks that have higher probability of existing will be used.

Macros **{HOST.HOST}**, **{HOST.NAME}**, **{HOST.IP}**, **{HOST.DNS}** and **{HOST.CONN}** are now supported in the parameters field for SSH, telnet and DB monitor items.

9.3.1 Improved SNMPv3 failure handling

Previously, Zabbix server and proxy would handle all SNMPv3 querying failures as a failure of the whole host. Since 2.0.4, if SNMPv3 *security name* or *security auth passphrase* are wrong, this will be detected and such items will turn unsupported. Item information will be updated with a proper error message, which will be visible in the item configuration list.

9.3.2 Java gateway improvements

Zabbix Java gateway now supports **dots in attribute name**. To use such an attribute name, dots must be escaped with backslashes.

10 What's new in Zabbix 2.0.5

10.1 Frontend improvements

- Frontend overall performance was improved by optimising performance checks
- SQL query count was reduced and performance improved in *IT services* and most pages that access trigger information
- Graphs in web monitoring will use new flicker-free refresh method now
- Error messages for screen row and column span validation were improved
- Reduced the amount of persistent connections and to Oracle backend and changed to suggested Oracle access functions

10.1.1 New translations

- Bulgarian

10.1.2 Completed translations

The following translations are 100% completed for this release:

- American English
- Brazilian Portuguese
- French
- German
- Greek
- Japanese
- Portuguese (Portugal)
- Russian
- Ukrainian

10.1.3 Updated translations

- American English
- Brazilian Portuguese
- Dutch
- Finnish
- French
- German
- Greek
- Hungarian
- Japanese
- Latvian
- Polish
- Portuguese (Portugal)
- Russian
- Slovak
- Spanish
- Swedish
- Turkish
- Ukrainian

10.2 Daemon improvements

- Source IP setting for SNMP checks was not supported with all Net-SNMP versions. Since Zabbix 2.0.5 this setting should work in all cases.
- Performance for Oracle as the Zabbix backend database was improved by prefetching 2MB of data in SELECT statements.
- Zabbix Windows agent now supports Windows 8 and Windows Server 2012 in the **system.uname** item.
- Zabbix Windows agent now supports non-ASCII interface names on Windows Vista/Server 2008 and later.

- **User macros** and **global regular expressions** now are supported in SNMP trap item key parameters.
- Zabbix agent **user parameter** output previously was limited to 64KB. It can now return up to 512KB of data. Additionally, peak memory usage of agent in cases when the output size is smaller may be reduced by up to 60KB.

10.2.1 Improved SNMP performance

- Performance of the SNMP normalize function was improved.
- Previously, each dynamic index was obtained individually when requested. On systems with thousands of dynamic indexes Zabbix pollers were 100% busy for more than 10 minutes until the cache was built up. Now the first dynamic index request will cause all indexes in the same SNMP table to be looked up and processed in one go. This greatly speeds up the cache building and reduces the SNMP query count. Note that all entries are looked up even if only some would be used later.

10.2.2 Increased cache size limits

The limit for Zabbix server and proxy configuration parameters *CacheSize*, *HistoryCacheSize* and *HistoryTextCacheSize*, as well as server parameter *TrendCacheSize*, was increased from 1GB to 2GB.

10.3 Miscellaneous improvements Zabbix sender now supports *ServerActive* parameter from the agent daemon configuration file. The first IP address is used.

11 What's new in Zabbix 2.0.6

11.1 New notification macros

- Macros **{TRIGGER.TEMPLATE.NAME}** and **{TRIGGER.HOSTGROUP.NAME}** will return the list of templates and host groups, respectively, in which the trigger is defined.
- Macros **{TRIGGER.NAME.ORIG}**, **{ITEM.NAME.ORIG}** and **{ITEM.KEY.ORIG}** will return the original name/key with macros not expanded.

See **macros supported by location** for details.

11.2 Frontend improvements 11.2.1 New translations

- Lithuanian

11.2.2 Completed translations

The following translations are 100% completed for this release:

- American English
- Brazilian Portuguese
- French
- German
- Greek
- Japanese
- Russian
- Slovak
- Ukrainian

11.2.3 Updated translations

- American English
- Brazilian Portuguese
- Dutch
- French
- German
- Greek
- Hungarian
- Japanese
- Latvian
- Russian
- Slovak
- Spanish
- Turkish
- Ukrainian

11.2.4 Miscellaneous frontend improvements

- Dashboard performance was improved by reducing the amount of SQL statements.
- Action conditions with deleted objects are handled differently. Starting with Zabbix 2.0.6 if any object (host, template, trigger, etc) that is used in the action condition is deleted, the condition is deleted and the action is disabled to avoid incorrect execution of the action. Before 2.0.6 the missing object was displayed as *unknown* and the condition remained in place.
- The value part of an action condition is now displayed in italics to make action conditions easier to read; like this: Host=*New host*

11.3 Daemon improvements 11.3.1 Improved LLD performance

Previously, low level discovery could generate a large amount of database queries, and even result in deadlocks. Starting with Zabbix 2.0.6, each low-level discovery prototype will be processed in its own transaction, avoiding the deadlocks.

Item prototype processing has also been significantly improved:

- Already discovered items will be updated only if some of their properties have changed.
- Only changed fields will be updated.
- The amount and size of the SQL statements has been reduced.
- Field validation LLD macro resolution has been improved:
 - validity of UTF-8 sequences
 - validity of field length

In a test run, the amount of SQL queries in the initial LLD run was reduced 64 times. In further runs when there are no changes to the discovered entities the amount of SQL queries was reduced 250 times. In further runs when there were changes the amount of SQL queries was reduced 210 times. Note that the total size of SQL queries was reduced even more.

11.3.2 Miscellaneous daemon improvements

The amount of traffic that is exchanged between Zabbix server and Java gateway was reduced.

The Zabbix server escalator process generates alerts while processing escalations. Previously these alerts were not available to the alerter process immediately, but only after the escalator finished its job. This has been fixed by splitting one transaction to handle all escalations into a separate transaction for each, which results in alerts becoming available to the alerter while the escalator is busy.

Logging failed active check connection

Zabbix agent now generates log warnings when an active check connection to server/proxy fails:

```
active check data upload to [127.0.0.1:10051] started to fail ([connect] cannot connect to [[127.0.0.1]:10051])
active check configuration update from [127.0.0.1:10051] started to fail (cannot connect to [[127.0.0.1]:10051])
```

and when it is restored later:

```
active check data upload to [127.0.0.1:10051] is working again
active check configuration update from [127.0.0.1:10051] is working again
```

12 What's new in Zabbix 2.0.7

12.1 Frontend improvements 12.1.1 Completed translations

The following translations are 100% completed for this release:

- American English
- Brazilian Portuguese
- French
- German
- Greek
- Japanese
- Russian
- Slovak
- Ukrainian

12.1.2 Updated translations

- Brazilian Portuguese

- Latvian

12.1.3 Value mapping in map labels

Value mapping support has been added in map labels so that raw values can be represented in a human-readable way.

12.1.4 Improved performance of System status widget

The performance of the dashboard *System status* widget has been improved by 3-7% on average for both memory usage and execution time.

12.1.5 Improved performance of SQL query condition calculation

Time required on 100 000 unsorted values has been reduced from 133 seconds to 2.5 seconds. This low-level improvement should help with frontend performance on large installations.

12.2 Daemon improvements 12.2.1 Improved performance with Oracle

Zabbix 2.0.5 improved performance by adding a 2MB prefetch. In some cases, this could lead to worse performance, thus Zabbix 2.0.7 changes the prefetch to be row based, which would improve the performance in those edge cases.

12.2.2 Improved performance of configuration cache update

Previously, Zabbix used a configuration cache mutex and string pool mutex. Starting with Zabbix 2.0.7, only one mutex is used. This change reduces the amount of time needed for the initial updating of the configuration cache by 20% on average.

12.2.3 Improved proxy performance

Zabbix proxy performance has been improved by reducing the amount of database queries on the proxy side during configuration updates. Previously, configuration updates changed all fields in the records that were to be modified. This has been modified to only update the values that have changed. The improvement is important for Zabbix proxies with a large number of items (hundreds of thousands and more).

12.2.4 Zabbix agent on Solaris

Getting detailed statistics about each CPU on Solaris systems is more efficient in Zabbix 2.0.7. This reduces CPU utilization by Zabbix agent, especially on systems with many cores.

system.swap.size calculation algorithm is changed to imitate "swap -s".

Warning:

This change is reverted in Zabbix 2.0.9 where **system.swap.size** is calculated again as in Zabbix 2.0.6 (imitates "swap -l").

12.2.5 No invalid recovery escalations

Escalation handling has been improved in case of trigger events. Before there was a possibility that for a trigger event invalid records could be added to the escalation table, resulting in extra work on adding/deleting records, which in case of big installations could affect performance of the database.

12.2.6 Zabbix agent on Solaris 64-bit platforms will be compiled as 64-bit application

By default now in case of a 64-bit *Solaris* the agent, sender and get will be compiled as 64-bit applications (previously they were compiled as 32-bit). The reason is that otherwise some functionality like `proc.mem` reports wrong values for 64-bit processes.

Note:

If you don't need support for large numbers on a 64-bit platform you will have to recompile Zabbix agent and make sure a compiler produces 32-bit application. Also note that some older compilers produce only 32-bit executables.

12.2.7 Improved performance of zabbix[queue] internal check

Before during queue calculations a configuration lock was made for every active item. Now the queue is calculated without acquiring configuration locks, improving performance on systems with a large number of items.

13 What's new in Zabbix 2.0.8

13.1 Frontend improvements 13.1.1 Completed translations

The following translations are 100% completed for this release:

- American English
- Brazilian Portuguese
- French
- German
- Greek
- Japanese
- Russian
- Slovak
- Ukrainian

13.1.2 Updated translations

- Dutch

13.1.3 Suppressed warning messages about deprecated mysql extension

Before Zabbix 2.0.8, PHP version 5.5+ users would get a warning message about the mysql PHP extension being deprecated. Starting with Zabbix 2.0.8, warning messages about the deprecated mysql PHP extension are suppressed. *Note:* using the mysqli PHP extension becomes an official requirement in the next major Zabbix 2.2.0 version.

13.2 Daemon improvements 13.2.1 ODBC monitoring improvements

- Before Zabbix 2.0.8 the login timeout for ODBC *Database monitor* items had a fixed value, 5 seconds. Now it is configurable and equal to "Timeout" parameter in server/proxy configuration file (3 seconds by default).
- More detailed ODBC **error messages**

13.2.2 Improved warnings for messages exceeding size limits

Previously, when a message exceeded size limits Zabbix ignored it by writing a simple warning to the log file:

```
19561:20130806:165515.375 Message size 99999999 exceeds the maximum size 67108864 bytes. Message ignored.
```

In Zabbix 2.0.8, the IP address is added to the warning to identify the message source:

```
19561:20130806:165515.375 Message size 99999999 from IP 10.0.3.1 exceeds the maximum size 67108864 bytes.
```

13.2.3 Security fixes

Zabbix server now correctly enables SSL host verification when using Ez Texting service to send alerts.

14 What's new in Zabbix 2.0.9

14.1 Frontend improvements 14.1.1 Faster loading graphs

The initial loading of graphs could be noticeably slow, especially with PostgreSQL. This was much improved by faster working SQL statements, significantly (up to 100-1000x) improving the loading times with all databases.

14.1.2 Completed translations

The following translations are 100% completed for this release:

- American English
- Brazilian Portuguese
- French
- German
- Greek
- Japanese
- Russian
- Slovak
- Ukrainian

14.1.3 Updated translations

- Brazilian Portuguese
- Chinese (China)
- Italian

- Lithuanian
- Romanian
- Slovak

14.2 Daemon improvements 14.2.1 Zabbix agent on Solaris

system.swap.size item calculation algorithm is changed back to imitate "swap -l" (as in Zabbix 2.0.6). Zabbix versions 2.0.7 and 2.0.8 imitated "swap -s" calculation, which was later recognized as a mistake.

15 What's new in Zabbix 2.0.10

15.1 Frontend improvements 15.1.1 Improved global script API calls

The loading speed of several Monitoring section screens (Last 20 issues, Latest data, Status of triggers, Latest events, ...) with a large number of global scripts was improved.

15.1.2 Faster loading of event details

The loading speed of the event detail page (*Monitoring* → *Events* → *Event details*) has been improved.

15.1.3 Completed translations

The following translations are 100% completed for this release:

15.1.4 Updated translations

- American English
- Brazilian Portuguese
- Czech
- Dutch
- French
- German
- Greek
- Japanese
- Italian
- Romanian
- Russian
- Slovak
- Spanish
- Ukrainian

15.2 Daemon improvements Sometimes, connection problems could have been interpreted as erroneous data package and thus misleading "Message size exceeds the maximum size bytes" would be reported instead. Since 2.0.10, the correct error message will be reported.

16 What's new in Zabbix 2.0.11

16.1 Frontend improvements 16.1.1 Updated translations

- French
- Italian
- Japanese
- Slovak
- Russian
- Ukrainian

16.1.2 LDAP bind password no longer viewable in clear text

- LDAP authentication bind password, once stored in the database, was accessible to Zabbix Super Admin level users in clear text in HTML source code. This has been fixed, by hiding the password from clear view.

16.2 Daemon improvements

- Zabbix application names in syslog fixed to meet RFC 5424 for APP-NAME. See [Syslog application names change](#)
- A trigger can now only be processed by one main, history syncer or timer process at a time, which should eliminate problems like multiple successive OK events and might lead to a performance improvement for timer processes on large systems, because they will not do duplicate work by processing triggers already being processed by history syncers.
- Trigger processing performance during low level discovery has been improved.
- Low level discovered triggers won't be deleted and will still work if relevant items are not discovered anymore (until those items get deleted).

17 What's new in Zabbix 2.0.12

17.1 Frontend improvements 17.1.1 Updated translations

- Brazilian Portuguese
- Italian
- Japanese
- Persian
- Russian
- Slovak
- Ukrainian

17.2 Daemon improvements

- Graph processing performance during low level discovery has been significantly improved. Testing with 2048 graphs showed a 600 times smaller amount of SQL requests during the initial discovery. Further runs without changes showed a 2500 times smaller amount of SQL requests, and if a change to graph name was required, the SQL request count was 1500 times lower. The total size of SQL statements was 3.7 times lower for the initial discovery, 3000 times lower for further runs without changes and 1500 times lower when a change to graph name was required.
- Graphs created by low-level discovery from now on will not be deleted and will still work if relevant items are not discovered anymore (until those items get deleted).
- Batch processing of IT services has been added. It resolves possible deadlocks and improves performance when processing large IT service trees. Testing with 800 IT services and having a tree depth of 4 levels showed a 300% performance improvement.
- Significantly improved log file monitoring (`log[]` and `logrt[]` item keys):
 - more efficient log file reading and matching of records against regular expression.
 - more efficient selecting of log files when checking `logrt[]` items.
 - for log file records longer than 256 kB only the first 256 kB are matched against the regular expression and the rest of the record is ignored. However, if Zabbix agent is stopped while it is dealing with a long record the agent internal state is lost and the long record may be analysed again and differently after the agent is started again.
 - changes for `logrt[]` items from 2.0.11 to 2.0.12:
 - * On UNIX platforms a `'logrt[]'` item becomes NOTSUPPORTED if a directory where the log files are e
 - * Unfortunately, on Microsoft Windows if a directory does not exist the item will not become NOTSUPPORT
 - * An absence of log files for `'logrt[]'` item does not make it NOTSUPPORTED (on 2.0.11 it caused NO
 - * Errors of reading log files for `'logrt[]'` item are logged as warnings into Zabbix agent log file
- * Zabbix agent log file can be helpful to find out why a `'log[]'` or `'logrt[]'` item became NOTSUPPORT
- * Please note that even though performance of `'log[]'` and `'logrt[]'` item checks has been improved th
- * Startup and shutdown scripts for Java gateway no longer hide error messages on startup. They now also de
- * No more partial server/proxy configuration. Error checks have been added to the process of loading confi

17.3 Miscellaneous improvements

- Zabbix agent daemon manpage now describes the meaning of value types in **-p** or **-t** output.
- An example robots.txt file has been added in the frontend.

18 What's new in Zabbix 2.0.13

18.1 Frontend improvements 18.1.1 Updated translations

- Slovak
- Ukrainian

18.2 Template changes

- **Template JMX Generic:** typo in item name "mpTenured" has been fixed to be "mp Tenured". If re-importing the template, this change requires no manual updates.

18.3 Daemon improvements

- Improved handling of log file rotation/truncation for `logrt[]` and `log[]` items. Special attention is paid to cases when several log files have the same last modification time. For more details see [Important notes](#) in the Log file monitoring section.
- Java gateway now uses Android JSON library instead of JSON.org library. When upgrading, apart from the gateway itself, it is necessary to replace the JSON library file and update `startup.sh` script. See [Java gateway file overview](#) for details.

19 What's new in Zabbix 2.0.14

19.1 Frontend improvements 19.1.1 Updated translations

- Italian
- Polish
- Spanish

20 What's new in Zabbix 2.0.15

20.1 Frontend improvements 20.1.1 Updated translations

- Czech
- Japanese
- Polish

20.2 Daemon improvements

- While item `net.tcp.service[ntp]` has existed for a long time, it almost never worked, because it tried to probe NTP protocol over TCP. It was rewritten to work over UDP and it now works.
- For Java gateway, it is now possible to specify timeout for JMX network operations using **TIMEOUT configuration option** in `startup.sh`.
- In actions, it is now possible to execute a custom script on the server if trigger expression contains multiple hosts.

20.3 Miscellaneous improvements 20.3.1 Validation of global regular expressions in LLD rules

A check for valid reference has been added for global regular expressions in LLD rules. If entered reference is not valid, due to misspelling or missing referenced global regular expression, the respective LLD rule will become unsupported and appropriate error message will be displayed.

21 What's new in Zabbix 2.0.16

21.1 Frontend improvements 21.1.1 Updated translations

- Czech
- Russian

22 What's new in Zabbix 2.0.17

22.1 Frontend improvements 22.1.1 Updated translations

- * English (United States)
- * French
- * Japanese
- * Korean
- * Polish
- * Ukrainian

22.2 Daemon improvements

22.3 Miscellaneous improvements

2. Zabbix concepts

Please use the sidebar to access content in the Zabbix concepts section.

1 Zabbix definitions

Overview

In this section you can learn the meaning of some terms commonly used in Zabbix.

Definitions

host

- a networked device that you want to monitor, with IP/DNS.

host group

- a logical grouping of hosts; it may contain hosts and templates. Hosts and templates within a host group are not in any way linked to each other. Host groups are used when assigning access rights to hosts for different user groups.

item

- a particular piece of data that you want to receive off of a host, a metric of data.

trigger

- a logical expression that defines a problem threshold and is used to "evaluate" data received in items

When received data are above the threshold, triggers go from 'Ok' into a 'Problem' state. When received data are below the threshold, triggers stay in/return to an 'Ok' state.

event

- a single occurrence of something that deserves attention such as a trigger changing state or a discovery/agent auto-registration taking place

action

- a predefined means of reacting to an event.

An action consists of operations (e.g. sending a notification) and conditions (when the operation is carried out)

escalation

- a custom scenario for executing operations within an action; a sequence of sending notifications/executing remote commands

media

- a means of delivering notifications; delivery channel

notification

- a message about some event sent to a user via the chosen media channel

remote command

- a pre-defined command that is automatically executed on a monitored host upon some condition

template

- a set of entities (items, triggers, graphs, screens, applications, low-level discovery rules) ready to be applied to one or several hosts

The job of templates is to speed up the deployment of monitoring tasks on a host; also to make it easier to apply mass changes to monitoring tasks. Templates are linked directly to individual hosts.

application

- a grouping of items in a logical group

web scenario

- one or several HTTP requests to check the availability of a web site

frontend

- the web interface provided with Zabbix

Zabbix API

- Zabbix API allows you to use the JSON RPC protocol to create, update and fetch Zabbix objects (like hosts, items, graphs and others) or perform any other custom tasks

Zabbix server

- a central process of Zabbix software that performs monitoring, interacts with Zabbix proxies and agents, calculates triggers, sends notifications; a central repository of data

Zabbix agent

- a process deployed on monitoring targets to actively monitor local resources and applications

Zabbix proxy

- a process that may collect data on behalf of Zabbix server, taking some processing load off of the server

node

- a full Zabbix server configured as an element within a hierarchy of distributed monitoring; it is responsible for monitoring its own location

2 Server

Overview

Zabbix server is the central process of Zabbix software.

The server performs the polling and trapping of data, it calculates triggers, sends notifications to users. It is the central component to which Zabbix agents and proxies report data on availability and integrity of systems. The server can itself remotely check networked services (such as web servers and mail servers) using simple service checks.

The server is the central repository in which all configuration, statistical and operational data is stored, and it is the entity in Zabbix that will actively alert administrators when problems arise in any of the monitored systems.

The functioning of a basic Zabbix server is broken into three distinct components; they are: Zabbix server, web frontend and database storage.

All of the configuration information for Zabbix is stored in the database, which both the server and the web frontend interact with. For example, when you create a new item using the web frontend (or API) it is added to the items table in the database. Then, about once a minute Zabbix server will query the items table for a list of the items which are active that is then stored in a cache within the Zabbix server. This is why it can take up to two minutes for any changes made in Zabbix frontend to show up in the latest data section.

Server process

Zabbix server runs as a daemon process. The server can be started by executing:

```
shell> cd sbin
shell> ./zabbix_server
```

You can use the following command line parameters with Zabbix server :

-c --config <file>	absolute path to the configuration file (default is /usr/local/etc/zabbix_
-n --new-nodeid <nodeid>	convert database data to new nodeid
-R --runtime-control <option>	perform administrative functions
-h --help	give this help
-V --version	display version number

Note:

Runtime control is not supported on OpenBSD and NetBSD.

Examples of command line parameters:

```
shell> zabbix_server -c /usr/local/etc/zabbix_server.conf
shell> zabbix_server --help
shell> zabbix_server -V
```

Runtime control

Runtime control options:

Option	Description
config_cache_reload	Reload configuration cache. Ignored if cache is being currently loaded.

Example of using runtime control to reload the server configuration cache:

```
shell> zabbix_server -c /usr/local/etc/zabbix_server.conf -R config_cache_reload
```

Process user

Zabbix server is designed to run as a non-root user. It will run as whatever non-root user it is started as. So you can run server as any non-root user without any issues.

If you will try to run it as 'root', it will switch to a hardcoded 'zabbix' user, which **must be present** on your system.

These settings currently cannot be user configured, neither during compilation nor in the configuration file.

If Zabbix server and **agent** are run on the same machine it is recommended to use a different user for running the server than for running the agent. Otherwise, if both are run as the same user, the agent can access the server configuration file and any Admin level user in Zabbix can quite easily retrieve, for example, the database password.

Configuration file

See the **configuration file** options for details on configuring zabbix_server.

Start-up scripts

The scripts are used to automatically start/stop Zabbix processes during system's start-up/shutdown. The scripts are located under directory misc/init.d.

Supported platforms

Due to the security requirements and mission-critical nature of server operation, UNIX is the only operating system that can consistently deliver the necessary performance, fault tolerance and resilience. Zabbix operates on market leading versions.

Zabbix server is tested on the following platforms:

- Linux
- Solaris
- AIX
- HP-UX
- Mac OS X
- FreeBSD
- OpenBSD
- NetBSD
- SCO Open Server
- Tru64/OSF1

Note:

Zabbix may work on other Unix-like operating systems as well.

3 Agent

Overview

Zabbix agent is deployed on a monitoring target to actively monitor local resources and applications (hard drives, memory, processor statistics etc).

The agent gathers operational information locally and reports data to Zabbix server for further processing. In case of failures (such as a hard disk running full or a crashed service process), Zabbix server can actively alert the administrators of the particular machine that reported the failure.

Zabbix agents are extremely efficient because of use of native system calls for gathering statistical information.

Passive and active checks

Zabbix agents can perform passive and active checks.

In a **passive check** the agent responds to a data request. Zabbix server (or proxy) asks for data, for example, CPU load, and Zabbix agent sends back the result.

Active checks require more complex processing. The agent must first retrieve a list of items from Zabbix server for independent processing. Then it will periodically send new values to the server.

Whether to perform passive or active checks is configured by selecting the respective monitoring **item type**. Zabbix agent processes items of type 'Zabbix agent' or 'Zabbix agent (active)'.

Supported platforms

Zabbix agent is supported for:

- Linux
- IBM AIX
- FreeBSD
- NetBSD
- OpenBSD
- HP-UX
- Mac OS X
- Solaris: 8, 9, 10, 11
- Windows: all desktop and server versions since 2000

Installation

See **installation instructions** for Zabbix agent.

Attention:

In general, 32bit Zabbix agents will work on 64bit systems, but may fail in some cases.

Running agent processes

Zabbix agent is run on the host being monitored.

Zabbix agent on UNIX runs as a daemon process. To start the agent (standalone daemon), execute:

```
shell> cd sbin
shell> ./zabbix_agentd
```

Zabbix agent on Windows runs as a Windows service. You can run a single instance of Zabbix agent or multiple instances of the agent on a host. A single instance can use the default configuration file or a configuration file specified in the command line. In case of multiple instances each agent instance must have its own configuration file (one of the instances can use the default configuration file).

The following command line parameters can be used with Zabbix agent:

Parameter	Description
UNIX and Windows agent	

Parameter	Description
-c --config <config-file>	Absolute path to the configuration file. You may use this option to specify a configuration file that is not the default one. On UNIX, default is /usr/local/etc/zabbix_agentd.conf or as set by compile-time variables --sysconfdir or --prefix On Windows, default is c:\zabbix_agentd.conf
-p --print	Print known items and exit. <i>Note:</i> To return user parameter results as well, you must specify the configuration file (if it is not in the default location).
-t --test <item key>	Test specified item and exit. <i>Note:</i> To return user parameter results as well, you must specify the configuration file (if it is not in the default location).
-h --help	Display help information
-V --version	Display version number
Windows agent only	
-m --multiple-agents	Use multiple agent instances (with -i,-d,-s,-x functions). To distinguish service names of instances, each service name will include the Hostname value from the specified configuration file.
Windows agent only (functions)	
-i --install	Install Zabbix Windows agent as service
-d --uninstall	Uninstall Zabbix Windows agent service
-s --start	Start Zabbix Windows agent service
-x --stop	Stop Zabbix Windows agent service

See **more details** on installing and running Zabbix agent on Windows.

Specific **examples** of using command line parameters:

- printing all built-in agent items with values
- testing a user parameter with "mysql.ping" key defined in the specified configuration file
- installing a "Zabbix Agent" service for Windows using the default path to configuration file c:\zabbix_agentd.conf
- installing a "Zabbix Agent [Hostname]" service for Windows using the configuration file zabbix_agentd.conf located in the same folder as agent executable and make the service name unique by extending it by Hostname value from the config file

```
shell> zabbix_agentd --print
shell> zabbix_agentd -t "mysql.ping" -c /etc/zabbix/zabbix_agentd.conf
shell> zabbix_agentd.exe -i
shell> zabbix_agentd.exe -i -m -c zabbix_agentd.conf
```

Process user

Zabbix agent on UNIX is designed to run as a non-root user. It will run as whatever non-root user it is started as. So you can run agent as any non-root user without any issues.

If you will try to run it as 'root', it will switch to a hardcoded 'zabbix' user, which must be present on your system. You can only run agent as 'root' if you modify the 'AllowRoot' parameter in the agent configuration file accordingly.

Locale

Note that the agent requires a UTF-8 locale so that some textual agent items can return the expected content. Most modern Unix-like systems have a UTF-8 locale as default, however, there are some systems where that may need to be set specifically.

Configuration file

For details on configuring Zabbix agent see the configuration file options for **zabbix_agentd** or **Windows agent**.

4 Proxy

Overview

Zabbix proxy is a process that may collect monitoring data from one or more monitored devices and send the information to the Zabbix server, essentially working on behalf of the server. All collected data is buffered locally and then transferred to the Zabbix server the proxy belongs to.

Deploying a proxy is optional, but may be very beneficial to distribute the load of a single Zabbix server. If only proxies collect data, processing on the server becomes less CPU and disk I/O hungry.

A Zabbix proxy is the ideal solution for centralized monitoring of remote locations, branches and networks with no local administrators.

Zabbix proxy requires a separate database.

Attention:

Note that databases supported with Zabbix proxy are SQLite, MySQL and PostgreSQL. Using Oracle or IBM DB2 is at your own risk and may contain some limitations as, for example, in **return values** of low-level discovery rules.

See also: [Using proxies in a distributed environment](#)

Proxy process

Zabbix proxy runs as a daemon process. The proxy can be started by executing:

```
shell> cd sbin
shell> ./zabbix_proxy
```

You can use the following command line parameters with Zabbix proxy:

<code>-c --config <file></code>	absolute path to the configuration file
<code>-R --runtime-control <option></code>	perform administrative functions
<code>-h --help</code>	give this help
<code>-V --version</code>	display version number

Examples of command line parameters:

```
shell> zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf
shell> zabbix_proxy --help
shell> zabbix_proxy -V
```

Runtime control

Runtime control options:

Option	Description
config_cache_reload	Reload configuration cache. Ignored if cache is being currently loaded. Active Zabbix proxy will connect to the Zabbix server and request configuration data.

Example of using runtime control to reload the proxy configuration cache:

```
shell> zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R config_cache_reload
```

Note:

Runtime control is not supported on OpenBSD and NetBSD.

Process user

Zabbix proxy is designed to run as a non-root user. It will run as whatever non-root user it is started as. So you can run proxy as any non-root user without any issues.

If you will try to run it as 'root', it will switch to a hardcoded 'zabbix' user, which must be present on your system.

These settings currently cannot be user configured, neither during compilation nor in the configuration file.

Configuration file

See the [configuration file](#) options for details on configuring zabbix_proxy.

Supported platforms

Zabbix proxy runs on the same list of [server#supported platforms](#) as Zabbix server.

5 Java gateway

Overview

Zabbix 2.0 adds native support for monitoring JMX applications by introducing a new Zabbix daemon called "Zabbix Java gateway". Zabbix Java gateway is a daemon written in Java. When Zabbix server wants to know the value of a particular JMX counter on a host, it asks Zabbix Java gateway, which uses the [JMX management API](#) to query the application of interest remotely. The application does not need any additional software installed, it just has to be started with `-Dcom.sun.management.jmxremote` option on the command line.

Java gateway accepts incoming connection from Zabbix server or proxy and can only be used as a "passive proxy". As opposed to Zabbix proxy, it may also be used from Zabbix proxy (Zabbix proxies cannot be chained). Access to each Java gateway is configured directly in Zabbix server or proxy configuration file, thus only one Java gateway may be configured per Zabbix server or Zabbix proxy. If a host will have items of type **JMX agent** and items of other type, only the **JMX agent** items will be passed to Java gateway for retrieval.

When an item has to be updated over Java gateway, Zabbix server or proxy will connect to the Java gateway and request the value, which Java gateway in turn retrieves and passes back to the server or proxy. As such, Java gateway does not cache any values.

Zabbix server or proxy has a specific type of processes that connect to Java gateway, controlled by the option **StartJavaPollers**. Internally, Java gateway starts multiple threads, controlled by the **START_POLLERS** option. On the server side, if a connection takes more than **Timeout** seconds, it will be terminated, but Java gateway might still be busy retrieving value from the JMX counter. To solve this, since Zabbix 2.0.15 there is the **TIMEOUT** option in Java gateway that allows to set timeout for JMX network operations.

Zabbix server or proxy will try to pool requests to a single JMX target together as much as possible (affected by item intervals) and send them to the Java Gateway in a single connection for better performance.

It is suggested to have **StartJavaPollers** less than or equal to **START_POLLERS**, otherwise there might be situations when no threads are available in the Java gateway to service incoming requests.

Sections below describe how to get and run Zabbix Java gateway, how to configure Zabbix server (or Zabbix proxy) to use Zabbix Java gateway for JMX monitoring, and how to configure Zabbix items in Zabbix GUI that correspond to particular JMX counters.

5.1 Getting Java gateway

There are two ways to get Java gateway. One is to download Java gateway package from Zabbix website and the other is to compile Java gateway from source.

5.1.1 Downloading from Zabbix website

Zabbix Java gateway packages (RHEL, Debian, Ubuntu) are available for download at <http://www.zabbix.com/download.php>.

5.1.2 Compiling from source

In order to compile Java gateway, you first run `./configure` script with `--enable-java` option. It is advisable that you specify `--prefix` option to request installation path other than the default `/usr/local`, because installing Java gateway will create a whole directory tree, not just a single executable.

```
$ ./configure --enable-java --prefix=$PREFIX
```

To compile and package Java gateway into a JAR file, run `make`. Note that for this step you will need `javac` and `jar` executables in your path.

```
$ make
```

Now you have `zabbix-java-gateway-$VERSION.jar` file in `src/zabbix_java/bin`. If you are comfortable with running Java gateway from `src/zabbix_java` in the distribution directory, then you can proceed to instructions for configuring and running Java gateway. Otherwise, make sure you have enough privileges and run `make install`.

```
$ make install
```

5.2 Overview of files in Java gateway distribution

Regardless of how you obtained Java gateway, you should have ended up with a collection of shell scripts, JAR and configuration files under `$PREFIX/sbin/zabbix_java`. The role of these files is summarized below.

`bin/zabbix-java-gateway-$VERSION.jar`

Java gateway JAR file itself.

```
lib/logback-core-0.9.27.jar
lib/logback-classic-0.9.27.jar
lib/slf4j-api-1.6.1.jar
lib/android-json-4.3_r3.1.jar
```

Dependencies of Java gateway: [Logback](#), [SLF4J](#), and [Android JSON](#) library (note that up to Zabbix 2.0.12 [JSON.org](#) library was used).

```
lib/logback.xml
lib/logback-console.xml
```

Configuration files for Logback.

```
shutdown.sh
startup.sh
```

Convenience scripts for starting and stopping Java gateway.

```
settings.sh
```

Configuration file that is sourced by startup and shutdown scripts above.

5.3 Configuring and running Java gateway

By default, Java gateway listens on port 10052. If you plan on running Java gateway on a different port, you can specify that in settings.sh script. See the description of [Java gateway configuration file](#) for how to specify this and other options.

Warning:

Port 10052 is not [IANA registered](#).

Once you are comfortable with the settings, you can start Java gateway by running the startup script:

```
$ ./startup.sh
```

Likewise, once you no longer need Java gateway, run the shutdown script to stop it:

```
$ ./shutdown.sh
```

Note that unlike server or proxy, Java gateway is lightweight and does not need a database.

5.4 Configuring server for use with Java gateway

Now that Java gateway is running, you have to tell Zabbix server where to find Zabbix Java gateway. This is done by specifying JavaGateway and JavaGatewayPort parameters in [server configuration file](#). If the host on which JMX application is running is monitored by Zabbix proxy, then you specify the connection parameters in [proxy configuration file](#) instead.

```
JavaGateway=192.168.3.14
JavaGatewayPort=10052
```

By default, server does not start any processes related to JMX monitoring. If you wish to use it, however, you have to specify the number of pre-forked instances of Java pollers. You do this in the same way you specify regular pollers and trappers.

```
StartJavaPollers=5
```

Do not forget to restart server or proxy, once you are done with configuring them.

5.5 Debugging Java gateway

In case there are any problems with Java gateway or an error message that you see about an item in the frontend is not descriptive enough, you might wish to take a look at Java gateway log file.

By default, Java gateway logs its activities into /tmp/zabbix_java.log file with log level "info". Sometimes that information is not enough and there is a need for information at log level "debug". In order to increase logging level, modify file lib/logback.xml and change the level attribute of <root> tag to "debug":

```
<root level="debug">
  <appender-ref ref="FILE" />
</root>
```

Note that unlike Zabbix server or Zabbix proxy, there is no need to restart Zabbix Java gateway after changing logback.xml file - changes in logback.xml will be picked up automatically. When you are done with debugging, you can return the logging level to "info".

If you wish to log to a different file or a completely different medium like database, adjust logback.xml file to meet your needs. See [Logback Manual](#) for more details.

Sometimes for debugging purposes it is useful to start Java gateway as a console application rather than a daemon. To do that, comment out `PID_FILE` variable in `settings.sh`. If `PID_FILE` is omitted, `startup.sh` script starts Java gateway as a console application and makes Logback use `lib/logback-console.xml` file instead, which not only logs to console, but has logging level "debug" enabled as well.

Finally, note that since Java gateway uses SLF4J for logging, you can replace Logback with the framework of your choice by placing an appropriate JAR file in `lib` directory. See [SLF4J Manual](#) for more details.

6 Sender

Overview

Zabbix sender is a command line utility that may be used to send performance data to Zabbix server for processing.

The utility is usually used in long running user scripts for periodical sending of availability and performance data.

For sending results directly to Zabbix server or proxy, a **trapper item** type must be configured.

Running Zabbix sender

An example of running Zabbix UNIX sender:

```
shell> cd bin
shell> ./zabbix_sender -z zabbix -s "Linux DB3" -k db.connections -o 43
```

where:

- `z` - Zabbix server host (IP address can be used as well)
- `s` - technical name of monitored host (as registered in Zabbix frontend)
- `k` - item key
- `o` - value to send

Attention:

If objects have whitespaces, these objects must be quoted using double quotes.

See the [Zabbix sender manpage](#) for more information.

Zabbix sender accepts strings in UTF-8 encoding (for both UNIX-like systems and Windows) without byte order mark (BOM) first in the file.

Zabbix sender on Windows can be run similarly:

```
zabbix_sender.exe [options]
```

Since Zabbix 1.8.4, `zabbix_sender` realtime sending scenarios have been improved to gather multiple values passed to it in close succession and send them to the server in a single connection. A value that is not further apart from the previous value than 0.2 seconds can be put in the same stack, but maximum pooling time still is 1 second.

If sending many values from an input file, Zabbix sender will batch them at 250 values in one go (all values will be processed), for example:

```
# zabbix_sender -z 127.0.0.1 -i /tmp/traptest.txt
Info from server: "Processed 250 Failed 0 Total 250 Seconds spent 0.002668"
Info from server: "Processed 50 Failed 0 Total 50 Seconds spent 0.000540"
sent: 300; skipped: 0; total: 300
```

All entries from an input file are sent in a sequential order top-down.

If the target item has triggers referencing it, all timestamps in an input file must be in an increasing order, otherwise event calculation will not be correct.

Note:

Zabbix sender will terminate if invalid (not following *parameter=value* notation) parameter entry is present in specified configuration file.

7 Get

Overview

Zabbix get is a process which can be used to communicate with Zabbix agent and retrieve required information from the agent.

The utility is usually used for the troubleshooting of Zabbix agents.

Running Zabbix get

An example of running Zabbix get under UNIX to get the processor load value from the agent:

```
shell> cd bin
shell> ./zabbix_get -s 127.0.0.1 -p 10050 -k system.cpu.load[all,avg1]
```

Another example of running Zabbix get for capturing a string from a website:

```
shell> cd bin
shell> ./zabbix_get -s 192.168.1.1 -p 10050 -k "web.page.regex[www.zabbix.com,,,\"USA: ([a-zA-Z0-9.-]+)\\""
```

Note that the item key here contains a space so quotes are used to mark the item key to the shell. The quotes are not part of the item key; they will be trimmed by the shell and will not be passed to Zabbix agent.

Zabbix get accepts the following command line parameters:

-s --host <host name or IP>	Specify host name or IP address of a host.
-p --port <port number>	Specify port number of agent running on the host. Default is 10050.
-I --source-address <IP address>	Specify source IP address.
-k --key <item key>	Specify key of item to retrieve value of.
-h --help	Give this help.
-V --version	Display version number.

Zabbix get on Windows can be run similarly:

```
zabbix_get.exe [options]
```

3. Installation

Please use the sidebar to access content in the Installation section.

1 Getting Zabbix

Overview

There are three ways of getting Zabbix:

- Install it from the [packages](#)
- Download the latest source archive and [compile it yourself](#)
- Download the [virtual appliance](#)

To download the latest sources or the virtual appliance, go to the [Zabbix download page](#), where direct links to latest versions are provided. To download older versions, see the link below stable version downloads.

2 Requirements

Hardware

Memory

Zabbix requires both physical and disk memory. 128 MB of physical memory and 256 MB of free disk space could be a good starting point. However, the amount of required disk memory obviously depends on the number of hosts and parameters that are being monitored. If you're planning to keep a long history of monitored parameters, you should be thinking of at least a couple of

gigabytes to have enough space to store the history in the database. Each Zabbix daemon process requires several connections to a database server. Amount of memory allocated for the connection depends on configuration of the database engine.

Note:

The more physical memory you have, the faster the database (and therefore Zabbix) works!

CPU

Zabbix and especially Zabbix database may require significant CPU resources depending on number of monitored parameters and chosen database engine.

Other hardware

A serial communication port and a serial GSM modem are required for using SMS notification support in Zabbix. USB-to-serial converter will also work.

Examples of hardware configuration

The table provides several examples of hardware configurations:

Name	Platform	CPU/Memory	Database	Monitored hosts
<i>Small</i>	CentOS	Virtual Appliance	SQLite	100
<i>Medium</i>	CentOS	2 CPU cores/2GB	MySQL InnoDB	500
<i>Large</i>	RedHat Enterprise Linux	4 CPU cores/8GB	RAID10 MySQL InnoDB or PostgreSQL	>1000
<i>Very large</i>	RedHat Enterprise Linux	8 CPU cores/16GB	Fast RAID10 MySQL InnoDB or PostgreSQL	>10000

Note:

Actual configuration depends on the number of active items and refresh rates very much. It is highly recommended to run the database on a separate box for large installations.

Supported platforms

Due to security requirements and mission-critical nature of monitoring server, UNIX is the only operating system that can consistently deliver the necessary performance, fault tolerance and resilience. Zabbix operates on market leading versions.

Zabbix is tested on the following platforms:

- Linux
- IBM AIX
- FreeBSD
- NetBSD
- OpenBSD
- HP-UX
- Mac OS X
- Solaris
- Windows: all desktop and server versions since 2000 (Zabbix agent only)

Note:

Zabbix may work on other Unix-like operating systems as well.

Software

Zabbix is built around a modern Apache web server, leading database engines, and PHP scripting language.

Database management system

Software	Version	Comments
MySQL	5.0 or later	Required if MySQL is used as Zabbix backend database. InnoDB engine is required.

Software	Version	Comments
<i>Oracle</i>	10g or later	Required if Oracle is used as Zabbix backend database.
<i>PostgreSQL</i>	8.1 or later	Required if PostgreSQL is used as Zabbix backend database. It is suggested to use at least PostgreSQL 8.3, which introduced much better VACUUM performance .
<i>SQLite</i>	3.3.5 or later	Required if SQLite is used as Zabbix backend database.
<i>IBM DB2</i>	9.7 or later	Required if IBM DB2 is used as Zabbix backend database.

Attention:

IBM DB2 support is experimental!

Attention:

SQLite3 support for use with Zabbix server is experimental!

Frontend

The following software is required to run Zabbix frontend:

Software	Version	Comments
<i>Apache</i>	1.3.12 or later	
<i>PHP</i>	5.1.6 or later	PHP v7 is not supported.
PHP extensions: <i>gd</i>	2.0 or later	PHP GD extension must support PNG images (<i>--with-png-dir</i>), JPEG (<i>--with-jpeg-dir</i>) images and FreeType 2 (<i>--with-freetype-dir</i>).
<i>bcmath</i>		php-bcmath (<i>--enable-bcmath</i>)
<i>ctype</i>		php-ctype (<i>--enable-ctype</i>)
<i>libXML</i>	2.6.15 or later	php-xml or php5-dom, if provided as a separate package by the distributor.
<i>xmlreader</i>		php-xmlreader, if provided as a separate package by the distributor.
<i>xmlwriter</i>		php-xmlwriter, if provided as a separate package by the distributor.
<i>session</i>		php-session, if provided as a separate package by the distributor.
<i>sockets</i>		php-net-socket (<i>--enable-sockets</i>). Required for user script support.
<i>mbstring</i>		php-mbstring (<i>--enable-mbstring</i>)
<i>gettext</i>		php-gettext (<i>--with-gettext</i>)
<i>ldap</i>		php-ldap. Required only if LDAP authentication is used in the frontend.
<i>ibm_db2</i>		Required if IBM DB2 is used as Zabbix backend database.
<i>mysql</i>		Required if MySQL is used as Zabbix backend database.
<i>oci8</i>		Required if Oracle is used as Zabbix backend database.
<i>pgsql</i>		Required if PostgreSQL is used as Zabbix backend database.
<i>sqlite3</i>		Required if SQLite is used as Zabbix backend database.

Note:

Zabbix may work on previous versions of Apache, MySQL, Oracle, and PostgreSQL as well.

Attention:

For other fonts than the default DejaVu, PHP function `imagerotate` might be required. If it is missing, these fonts might be rendered incorrectly in Monitoring → Overview header and other locations. This function is only available if PHP is compiled with bundled GD, which is not the case in Debian and other distributions.

WEB browser on client side

Cookies and Java Script must be enabled.

Latest versions of Google Chrome, Mozilla Firefox, Microsoft Internet Explorer and Opera are supported. Other browsers (Apple Safari, Konqueror) may work with Zabbix as well.

Server

Requirement	Description
<i>OpenIPMI</i>	Required for IPMI support.
<i>libssh2</i>	Required for SSH support. Version 1.0 or higher.
<i>fping</i>	Required for ICMP ping items.
<i>libcurl</i>	Required for web monitoring.
<i>libiksemel</i>	Required for Jabber support.
<i>net-snmp</i>	Required for SNMP support.

Java gateway

If you obtained Zabbix from the source repository or an archive, then the necessary dependencies are already included in the source tree.

If you obtained Zabbix from your distribution's package, then the necessary dependencies are already provided by the packaging system.

In both cases above, the software is ready to be used and no additional downloads are necessary.

If, however, you wish to provide your versions of these dependencies (for instance, if you are preparing a package for some Linux distribution), below is the list of library versions that Java gateway is known to work with. Zabbix may work with other versions of these libraries, too.

The following table lists JAR files that are currently bundled with Java gateway in the original code:

Library	License	Website	Comments
<i>logback-core-0.9.27.jar</i>	EPL 1.0, LGPL 2.1	http://logback.qos.ch/	Tested with 0.9.27, 1.0.13, and 1.1.1.
<i>logback-classic-0.9.27.jar</i>	EPL 1.0, LGPL 2.1	http://logback.qos.ch/	Tested with 0.9.27, 1.0.13, and 1.1.1.
<i>slf4j-api-1.6.1.jar</i>	MIT License	http://www.slf4j.org/	Tested with 1.6.1, 1.6.6, and 1.7.6.
<i>android-json-4.3_r3.1.jar</i>	Apache License 2.0	https://android.googlesource.com/platform/libcore/+master/json	Tested with 2.3.3_r1.1 and 4.3_r3.1. See <code>src/zabbix_java/lib/README</code> for instructions on creating a JAR file.

Java gateway compiles and runs with Java 1.6 and above. It is recommended that those who provide a precompiled version of the gateway for others use Java 1.6 for compilation, so that it runs on all versions of Java up to the latest one.

Database size

Zabbix configuration data require a fixed amount of disk space and do not grow much.

Zabbix database size mainly depends on these variables, which define the amount of stored historical data:

- Number of processed values per second

This is the average number of new values Zabbix server receives every second. For example, if we have 3000 items for monitoring with refresh rate of 60 seconds, the number of values per second is calculated as $3000/60 = 50$.

It means that 50 new values are added to Zabbix database every second.

- Housekeeper settings for history

Zabbix keeps values for a fixed period of time, normally several weeks or months. Each new value requires a certain amount of disk space for data and index.

So, if we would like to keep 30 days of history and we receive 50 values per second, total number of values will be around $(30 \times 24 \times 3600) \times 50 = 129.600.000$, or about 130M of values.

Depending on the database engine used, type of received values (floats, integers, strings, log files, etc), the disk space for keeping a single value may vary from 40 bytes to hundreds of bytes. Normally it is around 90 bytes per value for numeric items. In our case, it means that 130M of values will require $130M \times 90 \text{ bytes} = 10.9GB$ of disk space.

Note:

The size of text/log item values is impossible to predict exactly, but you may expect around 500 bytes per value.

- Housekeeper setting for trends

Zabbix keeps a 1-hour max/min/avg/count set of values for each item in the table **trends**. The data is used for trending and long period graphs. The one hour period can not be customised.

Zabbix database, depending on database type, requires about 90 bytes per each total. Suppose we would like to keep trend data for 5 years. Values for 3000 items will require $3000 \times 24 \times 365 \times 90 = 2.2GB$ per year, or **11GB** for 5 years.

- Housekeeper settings for events

Each Zabbix event requires approximately 170 bytes of disk space. It is hard to estimate the number of events generated by Zabbix daily. In the worst case scenario, we may assume that Zabbix generates one event per second.

It means that if we want to keep 3 years of events, this would require $3 \times 365 \times 24 \times 3600 \times 170 = 15GB$

The table contains formulas that can be used to calculate the disk space required for Zabbix system:

Parameter	Formula for required disk space (in bytes)
<i>Zabbix configuration</i>	Fixed size. Normally 10MB or less.
<i>History</i>	$\text{days} \times (\text{items} / \text{refresh rate}) \times 24 \times 3600 \times \text{bytes}$ items : number of items days : number of days to keep history refresh rate : average refresh rate of items bytes : number of bytes required to keep single value, depends on database engine, normally ~90 bytes.
<i>Trends</i>	$\text{days} \times (\text{items} / 3600) \times 24 \times 3600 \times \text{bytes}$ items : number of items days : number of days to keep history bytes : number of bytes required to keep single trend, depends on database engine, normally ~90 bytes.
<i>Events</i>	$\text{days} \times \text{events} \times 24 \times 3600 \times \text{bytes}$ events : number of event per second. One (1) event per second in worst case scenario. days : number of days to keep history bytes : number of bytes required to keep single trend, depends on database engine, normally ~170 bytes.

Note:

Average values such as ~90 bytes for numeric items, ~170 bytes for events have been gathered from real-life statistics using a MySQL backend database.

So, the total required disk space can be calculated as:

Configuration + History + Trends + Events

The disk space will NOT be used immediately after Zabbix installation. Database size will grow then it will stop growing at some point, which depends on housekeeper settings.

Note:

Disk space requirements for nodes in distributed setup are calculated in a similar way, but this also depends on a total number of child nodes linked to a node.

Time synchronisation

It is very important to have precise system date on server with Zabbix running. [ntpd](#) is the most popular daemon that synchronizes the host's time with the time of other machines.

3 Installation from packages

Overview

Zabbix SIA provides official RPM and DEB packages for Red Hat Enterprise Linux, Debian and Ubuntu LTS.

Package files are available at repo.zabbix.com. yum and apt repositories are also available on the server. A step-by-step tutorial for installing Zabbix from packages is provided here.

Red Hat Enterprise Linux / CentOS Supported for versions: RHEL5 and RHEL6

Installing repository configuration package

Install the repository configuration package. This package contains yum configuration files.

Zabbix 2.0 for RHEL5:

```
# rpm -ivh http://repo.zabbix.com/zabbix/2.0/rhel/5/x86_64/zabbix-release-2.0-1.el5.noarch.rpm
```

Zabbix 2.0 for RHEL6:

```
# rpm -ivh http://repo.zabbix.com/zabbix/2.0/rhel/6/x86_64/zabbix-release-2.0-1.el6.noarch.rpm
```

Installing Zabbix packages

Install Zabbix packages. Example for Zabbix server and web frontend with mysql database.

Note:

Zabbix official repository provides fping, iksemel, libssh2 packages as well. These packages are located in the *non-supported* directory.

```
# yum install zabbix-server-mysql zabbix-web-mysql
```

Example for installing Zabbix agent only.

```
# yum install zabbix-agent
```

Creating initial database

Create zabbix database and user on MySQL.

```
# mysql -uroot
mysql> create database zabbix character set utf8 collate utf8_bin;
mysql> grant all privileges on zabbix.* to zabbix@localhost identified by 'zabbix';
mysql> exit
```

Import initial schema and data.

```
# cd /usr/share/doc/zabbix-server-mysql-2.0.4/create
# mysql -uroot zabbix < schema.sql
# mysql -uroot zabbix < images.sql
# mysql -uroot zabbix < data.sql
```

Starting Zabbix server process

Edit database configuration in zabbix_server.conf

```
# vi /etc/zabbix/zabbix_server.conf
DBHost=localhost
DBName=zabbix
DBUser=zabbix
```

```
DBPassword=zabbix
```

Start Zabbix server process.

```
# service zabbix-server start
```

Editing PHP configuration for Zabbix frontend

Apache configuration file for Zabbix frontend is located in /etc/httpd/conf.d/zabbix.conf. Some PHP settings are already configured.

```
php_value max_execution_time 300
php_value memory_limit 128M
php_value post_max_size 16M
php_value upload_max_filesize 2M
php_value max_input_time 300
# php_value date.timezone Europe/Riga
```

It's necessary to uncomment the "date.timezone" setting and set the right timezone for you. After changing the configuration file restart the apache web server.

```
# service httpd restart
```

Zabbix frontend is available at <http://zabbix-frontend-hostname/zabbix> in the browser. Default username/password is Admin/zabbix.

Debian / Ubuntu Supported for version: Debian 6 (Squeeze), Debian 7 (Wheezy), Ubuntu 10.04 LTS (Lucid Lynx), Ubuntu 12.04 LTS (Precise Pangolin)

Installing repository configuration package

Install the repository configuration package. This package contains apt configuration files.

Zabbix 2.0 for Debian 6:

```
# wget http://repo.zabbix.com/zabbix/2.0/debian/pool/main/z/zabbix-release/zabbix-release_2.0-1squeeze_all.d
# dpkg -i zabbix-release_2.0-1squeeze_all.deb
# apt-get update
```

Zabbix 2.0 for Debian 7:

```
# wget http://repo.zabbix.com/zabbix/2.0/debian/pool/main/z/zabbix-release/zabbix-release_2.0-1wheezy_all.d
# dpkg -i zabbix-release_2.0-1wheezy_all.deb
# apt-get update
```

Zabbix 2.0 for Ubuntu 10.04 LTS:

```
# wget http://repo.zabbix.com/zabbix/2.0/ubuntu/pool/main/z/zabbix-release/zabbix-release_2.0-1lucid_all.d
# dpkg -i zabbix-release_2.0-1lucid_all.deb
# apt-get update
```

Zabbix 2.0 for Ubuntu 12.04 LTS:

```
# wget http://repo.zabbix.com/zabbix/2.0/ubuntu/pool/main/z/zabbix-release/zabbix-release_2.0-1precise_all.d
# dpkg -i zabbix-release_2.0-1precise_all.deb
# apt-get update
```

Installing and configuring MySQL database

Note:

This step is for Debian 6 and Ubuntu 10.04 with MySQL database only.

dbconfig-common on Debian 6 and Ubuntu 10.04 cannot create database with utf-8 encoding. It is necessary to install mysql-server package and configure mysql configuration to use utf-8 in advance.

For Debian 7 and Ubuntu, this step is not necessary.

Install MySQL server package:

```
# apt-get install mysql-server
```

Set character encoding parameter in /etc/mysql/my.cnf

```
[mysqld]
default-character-set=utf8
```

Note:

The encoding setting above is for MySQL 5.1. If you use MySQL 5.5 or later, please use "character-set-server=utf-8" instead.

Then start mysqld

```
# service mysql start
```

Installing Zabbix packages

Install Zabbix packages. dbconfig-common will create the database and populate the initial schema and data automatically. If backend db is located on a different server, please set dbc_remote_questions_default='true' in /etc/dbconfig-common/config.

Example for Zabbix server and web frontend with mysql database.

```
# apt-get install zabbix-server-mysql zabbix-frontend-php
```

Note:

The zabbix-frontend-php package, during installation, will configure a font, which is used on generated images. If you updated the package from any other repository and text is empty on graphs or maps, please check if a "ttf-dejavu-core" package is installed and try to execute "dpkg-reconfigure zabbix-frontend-php" command.

Example for installing Zabbix agent only.

```
# apt-get install zabbix-agent
```

Editing PHP configuration for Zabbix frontend

Apache configuration file for Zabbix frontend is located in /etc/apache2/conf.d/zabbix. Some PHP settings are already configured.

```
php_value max_execution_time 300
php_value memory_limit 128M
php_value post_max_size 16M
php_value upload_max_filesize 2M
php_value max_input_time 300
# php_value date.timezone Europe/Riga
```

It's necessary to uncomment the "date.timezone" setting and set the correct timezone for you. After changing the configuration file restart the apache web server.

```
# service apache2 restart
```

Zabbix frontend is available at <http://zabbix-frontend-hostname/zabbix> in the browser. Default username/password is Admin/zabbix.

Troubleshooting See the section on [installation-specific issue troubleshooting](#).

4 Installation from sources

Overview

You can get the very latest version of Zabbix by compiling it from the sources.

A step-by-step tutorial for installing Zabbix from the sources is provided here.

1 Installing Zabbix daemons

1 Download the source archive

Go to the [Zabbix download page](#) and download the source archive. Once downloaded, extract the sources, by running:

```
$ tar -zxvf zabbix-2.0.0.tar.gz
```

Note:

Enter the correct Zabbix version in the command. It must match the name of the downloaded archive.

2 Create user account

For all of the Zabbix daemon processes, an unprivileged user is required. If a Zabbix daemon is started from an unprivileged user account, it will run as that user.

However, if a daemon is started from a 'root' account, it will switch to a 'zabbix' user account, which must be present. To create such a user account (in its own group, "zabbix") on Linux systems, run:

```
groupadd zabbix
useradd -g zabbix zabbix
```

A separate user account is not required for Zabbix frontend installation.

If Zabbix **server** and **agent** are run on the same machine it is recommended to use a different user for running the server than for running the agent. Otherwise, if both are run as the same user, the agent can access the server configuration file and any Admin level user in Zabbix can quite easily retrieve, for example, the database password.

Attention:

Running Zabbix as root, bin, or any other account with special rights is a security risk.

3 Create Zabbix database

For Zabbix **server** and **proxy** daemons, as well as Zabbix frontend, a database is required. It is not needed to run Zabbix **agent**.

SQL **scripts are provided** for creating database schema and inserting the dataset. Zabbix proxy database needs only the schema while Zabbix server database requires also the dataset on top of the schema.

Having created a Zabbix database, proceed to the following steps of compiling Zabbix.

4 Configure the sources

When configuring the sources for a Zabbix server or proxy, you must specify the database type to be used. Only one database type can be compiled with a server or proxy process at a time.

To see all of the supported configuration options, inside the extracted Zabbix source directory run:

```
./configure --help
```

To configure the sources for a Zabbix server and agent, you may run something like:

```
./configure --enable-server --enable-agent --with-mysql --enable-ipv6 --with-net-snmp --with-libcurl
```

To configure the sources for a Zabbix server (with PostgreSQL etc.), you may run:

```
./configure --enable-server --with-postgresql --with-net-snmp
```

To configure the sources for a Zabbix proxy (with SQLite etc.), you may run:

```
./configure --prefix=/usr --enable-proxy --with-net-snmp --with-sqlite3 --with-ssh2
```

To configure the sources for a Zabbix agent, you may run:

```
./configure --enable-agent
```

You may use the `--enable-static` flag to statically link libraries. If you plan to distribute compiled binaries among different servers, you must use this flag to make these binaries work without required libraries. Note that `--enable-static` does not work in [Solaris](#).

Attention:

Using `--enable-static` option is not recommended when building server.// //

In order to build the server statically you must have a static version of every external library needed. There is no strict check for that in configure script.

Note:

Command-line utilities `zabbix_get` and `zabbix_sender` are compiled if `--enable-agent` option is used.

Note:

Use `--with-ibm-db2` flag to specify location of the CLI API.

Use `--with-oracle` flag to specify location of the OCI API.

Note:

`--with-ucd-snmp` flag can be used instead of `--with-net-snmp`. If no SNMP support is required, both `--with-net-snmp` and `--with-ucd-snmp` may be skipped.

5 Make and install everything

Note:

If installing from SVN, it is required to run first:

```
$ make dbschema
```

```
make install
```

This step should be run as a user with sufficient permissions (commonly 'root', or by using sudo).

Running `make install` will by default install the daemon binaries (`zabbix_server`, `zabbix_agentd`, `zabbix_proxy`) in `/usr/local/sbin` and the client binaries (`zabbix_get`, `zabbix_sender`) in `/usr/local/bin`.

Note:

To specify a different location than `/usr/local`, use a `--prefix` key in the previous step of configuring sources, for example `--prefix=/home/zabbix`. In this case daemon binaries will be installed under `<prefix>/sbin`, while utilities under `<prefix>/bin`. Man pages will be installed under `<prefix>/share`.

6 Review and edit configuration files

- edit the Zabbix agent configuration file **`/usr/local/etc/zabbix_agentd.conf`**

You need to configure this file for every host with `zabbix_agentd` installed.

You must specify the Zabbix server **IP address** in the file. Connections from other hosts will be denied.

- edit the Zabbix server configuration file **`/usr/local/etc/zabbix_server.conf`**

You must specify the database name, user and password (if using any).

Note:

With SQLite the full path to database file must be specified; DB user and password are not required.

The rest of the parameters will suit you with their defaults if you have a small installation (up to ten monitored hosts). You should change the default parameters if you want to maximize the performance of Zabbix server (or proxy) though. See the [Performance tuning] section for more details.

- if you have installed a Zabbix proxy, edit the proxy configuration file **`/usr/local/etc/zabbix_proxy.conf`**

You must specify the server IP address and proxy hostname (must be known to the server), as well as the database name, user and password (if using any).

Note:

With SQLite the full path to database file must be specified; DB user and password are not required.

7 Start up the daemons

Run `zabbix_server` on the server side.

```
shell> zabbix_server
```

Note:

Make sure that your system allows allocation of 36MB (or a bit more) of shared memory, otherwise the server may not start and you will see "Cannot allocate shared memory for <type of cache>." in the server log file. This may happen on FreeBSD, Solaris 8.

See the "[See also](#)" section at the bottom of this page to find out how to configure shared memory.

Run `zabbix_agentd` on all the monitored machines.

```
shell> zabbix_agentd
```

Note:

Make sure that your system allows allocation of 2MB of shared memory, otherwise the agent may not start and you will see "Cannot allocate shared memory for collector." in the agent log file. This may happen on Solaris 8.

If you have installed Zabbix proxy, run `zabbix_proxy`.

```
shell> zabbix_proxy
```

2 Installing Zabbix web interface

Copying PHP files

Zabbix frontend is written in PHP, so to run it a PHP supported webserver is needed. Installation is done by simply copying the PHP files from frontends/php to the webserver HTML documents directory.

Common locations of HTML documents directories for Apache web servers include:

- /usr/local/apache2/htdocs (default directory when installing Apache from source)
- /srv/www/htdocs (OpenSUSE, SLES)
- /var/www/html (Fedora, RHEL, CentOS)
- /var/www (Debian, Ubuntu)

It is suggested to use a subdirectory instead of the HTML root. To create a subdirectory and copy Zabbix frontend files into it, execute the following commands, replacing the actual directory:

```
mkdir <htdocs>/zabbix
cd frontends/php
cp -a . <htdocs>/zabbix
```

If installing from SVN and planning to use any other language than English, you must generate translation files. To do so, run:

```
locale/make_mo.sh
```

msgfmt utility from gettext package is required.

Note:

Additionally, to use any other language than English, its locale should be installed on the web server. See the ["See also"](#) section in the "User profile" page to find out how to install it if required.

Installing frontend

Step 1


In your browser, open Zabbix URL: `http://<server_ip_or_name>/zabbix`

You should see the first screen of the frontend installation wizard.



Step 2

Make sure that all software prerequisites are met.



Pre-requisite	Minimum value	Description
<i>PHP version</i>	5.1.6	
<i>PHP memory_limit option</i>	128MB	In php.ini: memory_limit = 128M
<i>PHP post_max_size option</i>	16MB	In php.ini: post_max_size = 16M
<i>PHP upload_max_filesize option</i>	2MB	In php.ini: upload_max_filesize = 2M
<i>PHP max_execution_time option</i>	300 seconds	In php.ini: max_execution_time = 300
<i>PHP max_input_time option</i>	300 seconds	In php.ini: max_input_time = 300
<i>PHP session.auto_start option</i>	must be disabled	In php.ini: session.auto_start = 0. Required since Zabbix 2.0.4 version .
<i>Database support</i>	One of: IBM DB2, MySQL, Oracle, PostgreSQL, SQLite	One of the following modules must be installed: ibm_db2, mysql, oci8, pgsql, sqlite3
<i>bcmath</i>		php-bcmath
<i>mbstring</i>		php-mbstring
<i>sockets</i>		php-net-socket. Required for user script support.
<i>gd</i>	2.0 or higher	php-gd. PHP GD extension must support PNG images (<i>--with-png-dir</i>), JPEG (<i>--with-jpeg-dir</i>) images and FreeType 2 (<i>--with-freetype-dir</i>).
<i>libxml</i>	2.6.15	php-xml or php5-dom
<i>xmlwriter</i>		php-xmlwriter
<i>xmlreader</i>		php-xmlreader
<i>ctype</i>		php-ctype
<i>session</i>		php-session
<i>gettext</i>		php-gettext

Step 3

Enter details for connecting to the database. Zabbix database must already be created.

1. Welcome
2. Check of pre-requisites
3. Configure DB connection
4. Zabbix server details
5. Pre-Installation summary
6. Install

www.zabbix.com
Licensed under [GPL v2](#)

Cancel

3. Configure DB connection

Please create database manually,
and set the configuration parameters for connection to this database.

Press "Test connection" button when done.

Database type	MySQL
Database host	localhost
Database port	0 0 - use default port
Database name	zabbix
User	zabbix
Password	*****

OK

Test connection

<< Previous

Next >>

Step 4

Enter Zabbix server details.

1. Welcome
2. Check of pre-requisites
3. Configure DB connection
4. Zabbix server details
5. Pre-Installation summary
6. Install

www.zabbix.com
Licensed under [GPL v2](#)

Cancel

4. Zabbix server details

Please enter host name or host IP address
and port number of Zabbix server,
as well as the name of the installation (optional).

Host	localhost
Port	10051
Name	

<< Previous

Next >>

Step 5

Review a summary of settings.

1. Welcome
2. Check of pre-requisites
3. Configure DB connection
4. Zabbix server details
5. Pre-Installation summary
6. Install

www.zabbix.com
Licensed under [GPL v2](#)

Cancel

5. Pre-Installation summary

Please check configuration parameters.
If all is correct, press "Next" button, or "Previous" button to change configuration parameters.

Database type	MySQL
Database server	localhost
Database port	default
Database name	zabbix
Database user	zabbix
Database password	*****
Zabbix server	localhost
Zabbix server port	10051
Zabbix server name	

<< Previous
Next >>

Step 6

Download the configuration file and place it under conf/.

1. Welcome
2. Check of pre-requisites
3. Configure DB connection
4. Zabbix server details
5. Pre-Installation summary
6. Install

www.zabbix.com
Licensed under [GPL v2](#)

Cancel

6. Install

Configuration file
"/srv/www/htdocs/zabbix/conf/zabbix.conf.php"
created: **Fail**

Retry


Unable to create the configuration file.
Please install it manually, or fix permissions on the conf directory.

Press the "Download configuration file" button, download the configuration file and save it as
"/srv/www/htdocs/zabbix/conf/zabbix.conf.php"

Download configuration file

When done, press the "Retry" button

<< Previous
Finish




Note:

Providing the webserver user has write access to conf/ directory the configuration file would be saved automatically and it would be possible to proceed to the next step right away.

Step 7

Finish the installation.



Step 8

Zabbix frontend is ready! The default user name is **Admin**, password **zabbix**.



Proceed to [getting started with Zabbix](#).

Troubleshooting See the section on [installation-specific issue troubleshooting](#).

See also

1. [How to configure shared memory for Zabbix daemons](#)

5 Known issues

IPMI checks

IPMI checks will not work with the standard OpenIPMI library from Debian/Ubuntu package. To fix, recompile OpenIPMI library with OpenSSL enabled as discussed in [ZBX-6139](#).

SSH checks

Some Linux distributions like Debian, Ubuntu do not support encrypted private keys (with passphrase) if the libssh2 library is installed from packages. Please see [ZBX-4850](#) for more details.

ODBC checks

Zabbix server or proxy that uses MySQL as its database may or may not work correctly with MySQL ODBC library due to an [upstream bug](#). Please see [ZBX-7665](#) for more information and available workarounds.

Simple checks

There is a bug in **fping** versions earlier than v3.10 that mishandles duplicate echo replay packets. This may cause unexpected results for `icmping`, `icmpingloss`, `icmpingsec` items. It is recommended to use the latest version of **fping**. Please see [ZBX-11726](#) for more details.

Known issues for 2.0.0

- Reordering entries in lists by dragging them sometimes moves elements inside the entry to the left hand side - this is a JQuery issue

Known issues for 2.0.0 - 2.0.1

- In action operations, the "Delay" column was calculated incorrectly if custom operation step durations were used - fixed for 2.0.2 (and also renamed to "Start in")

Known issues for 2.0.0 - 2.0.2

- Some pages could be displayed incorrectly by Internet Explorer with "Compatibility mode" enabled - fixed for 2.0.3

Known issues for 2.0.0 - 2.0.4

- IT services custom service times representation in the frontend could get broken after upgrading from 1.8 to 2.0 - fixed in the upgrade patch of 2.0.5

Known issues for 2.0.0 - 2.0.8

- Log size was not stored on Zabbix proxy, making the agent re-read the whole file upon restart.

Known issues for 2.0.0 - 2.0.10

- LDAP authentication bind password, once stored in the database, was accessible to Zabbix Super Admin level users in clear text in HTML source code. Fixed for 2.0.11, by hiding the password from clear view.

Known issues for 2.0.1

- Event exporting to CSV also exports extra data - fixed for 2.0.2

Known issues for 2.0.4 - 2.0.6

- It is not possible to redirect line "info from server" from Zabbix sender output - fixed for 2.0.7

Known issues for 2.0.12

- `logrt []` item processing is broken in Zabbix agent (see [ZBX-8238](#) for more information). As a workaround use Zabbix agent 2.0.11 on hosts with `logrt []` items or apply a [patch](#). The problem does not affect 2.0.12 server/proxy.

Known issues for 2.0.13 and later

* `log []` and `logrt []` items repeatedly reread log file from the beginning if file system is 100% full and the log file is being appended (see [ZBX-10884](#) for more information).

6 Upgrade procedure

Overview

This section provides the steps required for a successful upgrade from Zabbix 1.8.x to 2.0.

Database upgrade to version 2.0 may take a long time, as all the historical data tables have to be upgraded.

Attention:

Make sure to read [upgrade notes](#) before proceeding with the upgrade.

1 Stop Zabbix server

Stop Zabbix server to make sure that no new data is inserted into database.

2 Back up the existing Zabbix database

This is a very important step. Make sure that you have a backup of your database. It will help if the upgrade procedure fails (lack of disk space, power off, any unexpected problem).

3 Back up configuration files, PHP files and Zabbix binaries

Make a backup copy of Zabbix binaries, configuration files and PHP files.

4 Install new server binaries

You may use pre-compiled binaries or [compile your own](#).

5 Review server configuration parameters

Some parameters of `zabbix_server.conf` might have changed from 1.8, new parameters added. You may want to review them.

6 Upgrade the database

Before you upgrade:

- Make sure the database user has enough permissions (create table, drop table, create index, drop index)
- Make sure you have enough free disk space.

Database upgrade scripts are located in the `upgrades/dbpatches/2.0/<db engine>` directory. Run the scripts:

- MySQL: `upgrades/dbpatches/2.0/mysql/patch.sql`
- PostgreSQL: `upgrades/dbpatches/2.0/postgresql/patch.sql`
- Oracle: `upgrades/dbpatches/2.0/oracle/patch.sql`
- IBM DB2: `upgrades/dbpatches/2.0/ibm_db2/patch.sql`

These scripts are only for upgrading Zabbix 1.8.x to 2.0! For upgrading from earlier versions first use upgrade scripts from Zabbix 1.6.x or Zabbix 1.8.x.

7 Install new Zabbix web interface

Follow [installation instructions](#).

8 Start new Zabbix binaries

Start new binaries. Check log files to see if the binaries have started successfully.

7 Upgrade notes for 2.0.0

1 Requirement changes

- Dropping MSIE 6 and 7 support for Zabbix frontend
- MySQL 5.0
- Windows 2000 (should specify minimum SP level)
- Minimum supported PHP version has been changed to 5.1.6
- Note that with the use of jQuery for the Zabbix frontend browser compatibility might have suffered somewhat - most notably, Konqueror does not work anymore.
- Minimum supported PostgreSQL version has been changed to 8.1
- Frontend now requires:
 - PHP gettext
 - PHP XML Reader
 - PHP XML Writer

2 Case-sensitive MySQL database

A case-sensitive MySQL database is required for proper server work. It is **recommended** to create a case-sensitive MySQL database during new installations. If you created a MySQL database with the utf8 character set previously, in order to support case sensitivity of stored data, you need to convert the charset to utf8_bin.

3 Frontend specific upgrade notes

User group access property "API access" has been removed. Starting with Zabbix 2.0, all users have full access to the API.

Variables TRIGGER_FALSE_PERIOD and TRIGGER_BLINK_PERIOD in include/defines.inc.php have been removed. These settings may now be configured in *Administration* → *General* → *Trigger displaying options*.

Graph item type *Aggregated* has been removed. If such items are present, upgrade procedure will convert them to normal items.

Sorting by 'State' column in *Configuration* → *Maintenance* has been removed. Sorting by 'IP range' and 'Delay' columns in *Configuration* → *Discovery* have been removed.

After upgrading, sorting in all locations (tables) will be reset to defaults.

Multiple element selecting in maps while holding down *Shift* was removed. On Mac OSX multiple elements can now be selected by holding down *Command*.

Dropdown choice in *Administration* → *General* is not remembered anymore.

4 Database upgrading notes

Database upgrade to 2.0 for MySQL will run UPDATE statement that fixes web monitoring item units for all items including the ones with correct unit. This should be harmless and only result in some additional UPDATES.

For better performance, especially with large databases on MySQL, increasing [innodb_buffer_pool_size](#) before upgrading is highly suggested. Also, disable binlog, if it is not needed, or increase [binlog_cache_size](#).

During database upgrade with MySQL *ERROR 1436 (HY000): Thread stack overrun* - increase **thread_stack** in **my.cnf**. 256K should be enough.

5 Zabbix server related notes

Upgrading to 2.0.0 may increase Zabbix server configuration cache usage, because it now additionally caches general configuration and trigger information.

6 Changed default configuration directories

- **configuration files location**

/etc/zabbix → /usr/local/etc (depends on compile-time installation variable *sysconfdir*)

- **external scripts**

/etc/zabbix/externalscripts → /usr/local/share/zabbix/externalscripts (depends on compile-time installation variable *datadir*)

- **alert scripts**

/home/zabbix/bin → /usr/local/share/zabbix/alertscripts (depends on compile-time installation variable *datadir*)

- **include directories** (recommended)

/etc/zabbix/zabbix_server → /usr/local/etc/zabbix_server.conf.d

/etc/zabbix/zabbix_proxy → /usr/local/etc/zabbix_proxy.conf.d

/etc/zabbix/zabbix_agent → /usr/local/etc/zabbix_agent.conf.d

/etc/zabbix/zabbix_agentd → /usr/local/etc/zabbix_agentd.conf.d

7 Installation procedure changes

Schema file **sqlite.sql** has been renamed to **sqlite3.sql**.

Image file **images_pgsql.sql** has been renamed to **images_postgresql.sql**.

Database creation files have been moved in **database** directory. It contains one directory per supported database type.

Option --with-pgsql for **configure** script has been changed to --with-postgresql.

Directory **misc/conf** was moved to be just **conf** in the top level source directory.

8 Item changes

net.tcp.dns and **net.tcp.dns.query** have been **renamed** to **net.dns** and **net.dns.record**, respectively. Old syntax still supported as an alias.

All items that return text (character, log, text types of information), including user parameters and external checks, now can return whitespace only as well, setting the return value to an empty string.

Internal item **zabbix[log]** has been removed.

External check parameter handling was changed. Previously, only one parameter was accepted. Starting with Zabbix 2.0, key syntax conforms to other types of items and multiple comma-separated parameters may be passed.

Multi-line values can now be stored for the output of external checks.

Hardcoded first parameter {HOST.CONN} for external checks has been removed and now external check parameters may be fully customised. Upgrade process adds this macro as a first parameter to all existing external checks.

Attention:

After upgrading, it is recommended to manually review all external checks that have more than one parameter in the key for any possible conversion errors to correct.

The format of **simple checks** has been changed. The possibility of using IP addresses in simple checks has been added. Database upgrade process converts all existing simple check keys to the new syntax.

A **status** item has been removed. Instead of it a new **internal item** - **zabbix[host,agent,available]** has been added. The database upgrade process **doesn't convert** all existing 'status' keys and associated triggers to the new syntax, so they have to be converted manually.

9 Upgrade notes for proxies

Warning:

Older Zabbix proxies (1.6 and 1.8) cannot be used with Zabbix 2.0 server; similarly, a Zabbix 2.0 proxy cannot be used with Zabbix 1.8 or older server. All proxies must be upgraded at the same time as the server.

Zabbix 1.8 proxies may continue collecting data while the server database is upgraded. In this case proxy database must be converted before the new server is started up. Proxy will send all collected and converted values to the server once both server and proxy are upgraded and started up.

10 Upgrade notes for agents

Server parameter has been changed. Entries in it will never be used for active checks, only for passive. Only entries from **ServerActive** will be used for active checks.

Parameter **ServerPort** has been removed. If it will be still specified in the configuration file, agent will not start up.

Parameters **DisableActive** and **DisablePassive** have been removed. Instead, not specifying **ServerActive** will not do any active item processing, and specifying **StartAgents=0** will disable processing of incoming connections for passive items.

If these parameters will be still specified in the configuration file, agent will not start up.

See also: [Version compatibility](#).

11 Upgrade notes for web monitoring

If a host was assigned to be monitored by a proxy before upgrading to 2.0 and it had a web scenario assigned to it, this web scenario would be performed from the Zabbix server. After upgrade such a web scenario may stop working. After upgrading to 2.0 the proxy that this host is assigned to should be compiled with web monitoring support and be able to access websites, used in the scenario.

12 General notes

{IPADDRESS} macro deprecated, use {HOST.IP}

"Host profiles" were renamed to "host inventory". New host inventory macros {INVENTORY.*} were introduced. {PROFILE.*} macros are still supported but it's highly recommended to change those to {INVENTORY.*} . See [this page](#) for details.

8 Upgrade notes for 2.0.1

- In 2.0.0, XML export used a mix of short and long forms for empty tags (both `<tag></tag>` and `<tag/>`). Since 2.0.1, all empty tags use the empty form (`<tag/>`).
- Updating Zabbix agent from 2.0.0 to 2.0.1 on Windows systems might result in historical data loss. See [ZBX-5094](#) for more details.

9 Upgrade notes for 2.0.2

This minor version does not have any upgrade notes.

10 Upgrade notes for 2.0.3

Accepted data limit of 128MB was introduced when using Zabbix protocol. Any other data (including older Zabbix protocols) stays limited at 16MB.

11 Upgrade notes for 2.0.4

Requirement changes

Zabbix frontend will now require PHP session auto start disabled (session.auto_start parameter in php.ini file).

API changes

API version has been bumped to 2.0.4 and will match Zabbix version from now on.

12 Upgrade notes for 2.0.5

This minor version does not have any upgrade notes.

13 Upgrade notes for 2.0.6

This minor version does not have any upgrade notes.

14 Upgrade notes for 2.0.7

This minor version does not have any upgrade notes.

15 Upgrade notes for 2.0.8

Configuration changes

Upgrade to Zabbix 2.0.8 changes the ODBC login timeout for *Database monitor* items from a fixed value, 5 seconds, to the value of configurable "Timeout" parameter in server/proxy configuration file (3 seconds by default).

Daemon security fixes

Zabbix server now correctly enables SSL host verification when using Ez Texting service to send alerts.

16 Upgrade notes for 2.0.9

This minor version does not have any upgrade notes.

17 Upgrade notes for 2.0.10

This minor version does not have any upgrade notes.

18 Upgrade notes for 2.0.11

Syslog application names change

If Zabbix is logging to syslog then after an upgrade to Zabbix 2.0.11 you will see changes in the application names appearing in syslog:

```
Zabbix agent  → zabbix_agent
Zabbix Agent  → zabbix_agentd
Zabbix proxy   → zabbix_proxy
Zabbix server  → zabbix_server
Zabbix get     → zabbix_get
Zabbix Sender  → zabbix_sender
```

The old, incorrect names (on the left) contained a space which is not allowed by RFC 5424 for APP-NAME. If you are using regular expressions in monitoring of syslog you may want to adjust them for the new application names.

19 Upgrade notes for 2.0.12

IT services unlinked when deleting triggers

Previously, when removing triggers, the IT services linked to those triggers would also be removed. Now the IT services are simply unlinked from the removed triggers and their SLA calculation disabled.

Log file handling

Zabbix agent startup log level was changed in order to log information about started agent processes at DebugLevel set to 0.

20 Upgrade notes for 2.0.13

Template changes

Item name "mpTenured" has been fixed to be "mp Tenured" in *Template JMX Generic*.

In order to fix it, import this template from https://www.zabbix.org/wiki/Zabbix_Templates/Official_Templates.

Daemon changes

Java gateway now uses Android JSON library instead of JSON.org library. When upgrading, apart from the gateway itself, it is necessary to replace the JSON library file and update `startup.sh` script. See [Java gateway file overview](#) for details.

21 Upgrade notes for 2.0.14

This minor version does not have any upgrade notes.

22 Upgrade notes for 2.0.15

Template changes

Value type was changed from "Numeric (unsigned)" to "Numeric (float)" for items **system.stat[kthr,b]** and **system.stat[kthr,r]** in *Template OS AIX*. Both items were also added to "Performance" application.

In order to fix it, import this template from https://www.zabbix.org/wiki/Zabbix_Templates/Official_Templates.

Daemon changes

Previously, if Zabbix could not send ICMP ping packets to a particular host, all ICMP ping items would attain a value of 0 in some cases. Now, they always become unsupported.

Validation of global regular expressions in LLD rules

A check for valid reference has been added for global regular expressions in LLD rules. If entered reference is not valid, due to misspelling or missing referenced global regular expression, the respective LLD rule will become unsupported and appropriate error message will be displayed.

23 Upgrade notes for 2.0.16

Template changes

Item **vm.memory.size[total]** moved from "Filesystems" to "Memory" application in *Template OS Windows*.

In order to fix it, import this template from https://www.zabbix.org/wiki/Zabbix_Templates/Official_Templates.

24 Upgrade notes for 2.0.17

Item changes

Correct resolution of low-level discovery **macros** has been improved in calculated item **formulas**. Function parameters now will be quoted if, after resolving low-level discovery macros, they contain `,` `)` characters or start with `"`, `<space>` characters.

4. Quickstart

Please use the sidebar to access content in the Quickstart section.

1 Login and configuring user

Overview

In this section you will learn how to log in and set up a system user in Zabbix.

Login



This is the Zabbix "Welcome" screen. Enter the user name **Admin** with password **zabbix** to log in as a **Zabbix superuser**.

When logged in, you will see 'Connected as Admin' in the lower right corner of the page. Access to *Configuration* and *Administration* menus will be granted.

Protection against brute force attacks

In case of five consecutive failed login attempts, Zabbix interface will pause for 30 seconds in order to prevent brute force and dictionary attacks.

The IP address of a failed login attempt will be displayed after a successful login.

Adding user

To view information about users, go to *Administration* → *Users* and select *Users* in the dropdown.

☐	Alias	Name	Surname	User type	Groups	Is online?	Login	Frontend access	Debug mode	Status
☐	Admin	Zabbix	Administrator	Zabbix Super Admin	Zabbix administrators	Yes (Wed, 04 Jan 2012 15:39:51 +0200)	Ok	System default	Disabled	Enabled
☐	guest	Default	User	Zabbix User	Guests	Yes (Wed, 04 Jan 2012 15:33:42 +0200)	Ok	System default	Disabled	Enabled

Initially there are only two users defined in Zabbix.

- 'Admin' user is a Zabbix superuser, which has full permissions.
- 'Guest' user is a special default user. If you are not logged in, you are accessing Zabbix with "guest" permissions. By default, "guest" has no permissions on Zabbix objects.

To add a new user, click on *Create user*.

In the new user form, make sure to add your user to one of the existing user groups, for example 'Network administrators'.

UserMediaPermissions

Alias

user

Name

New

Surname

User

Password

.....

Password (once again)

.....

Groups

Network administrators

Add

Delete selected

Language

English (en_GB)

Theme

System default

Auto-login

☐

Auto-logout (min 90 seconds)

☐90

Refresh (in seconds)

30

Rows per page

50

URL (after login)

Save

Cancel

By default, new users have no media (notification delivery methods) defined for them. To create one, go to the 'Media' tab and click on *Add*.

New media ?

Type: Email

Send to: user@domain.tld

When active: 1-7,00:00-24:00

Use if severity:

- ☒ Not classified
- ☒ Information
- ☒ Warning
- ☒ Average
- ☒ High
- ☒ Disaster

Status: Enabled

Add
Cancel

In this pop-up, enter an e-mail address for the user.

You can specify a time period when the medium will be active (see [Time period specification](#) page for description of the format), by default a medium is always active. You can also customise [trigger severity](#) levels for which the medium will be active, but leave all of them enabled for now.

Click on *Add*, then click *Save* in the user properties form. The new user appears in the userlist.

☐	Alias ↑	Name	Surname	User type	Groups	Is online?	Login	Frontend access	Debug mode	Status
☐	Admin	Zabbix	Administrator	Zabbix Super Admin	Zabbix administrators	Yes (Wed, 04 Jan 2012 15:42:02 +0200)	Ok	System default	Disabled	Enabled
☐	guest	Default	User	Zabbix User	Guests	Yes (Wed, 04 Jan 2012 15:33:42 +0200)	Ok	System default	Disabled	Enabled
☐	user	New	User	Zabbix User	Network administrators	No	Ok	System default	Disabled	Enabled

Adding permissions

By default, a new user has no permissions to access hosts. To grant the user rights, click on the group of the user in the *Groups* column (in this case - 'Network administrators'). In the group properties form, go to the *Permissions* tab.

Permissions

Composing permissions

Read-write	Read only	Deny
Add Delete selected	Add Delete selected	Add Delete selected

This user is to have read-only access to *Linux servers* group, so click on *Add* below the 'Read only' listbox.

Read only

☐ Name

☐ Discovered hosts

☒ Linux servers

☐ Templates

☐ Windows servers

☐ Zabbix servers

Select

In this pop-up, mark the checkbox next to 'Linux servers', then click *Select*. *Linux servers* should be displayed in the respective box. In the user group properties form, click *Save*.

Attention:

In Zabbix, access rights to hosts are assigned to **user groups**, not individual users.

Done! You may try to log in using the credentials of the new user.

2 New host

Overview

In this section you will learn how to set up a new host.

A host in Zabbix is a networked entity (physical, virtual) that you wish to monitor. The definition of what can be a "host" in Zabbix is quite flexible. It can be a physical server, a network switch, a virtual machine or some application.

Adding host

Information about configured hosts in Zabbix is available in *Configuration* → *Hosts*. There is already one pre-defined host, called 'Zabbix server', but we want to learn adding another.

To add a new host, click on *Create*. This will present us with a host configuration form.

Host

Templates

IPMI

Macros

Host inventory

Host name

New host

Visible name

Groups

In groups

Linux servers

Other groups

Discovered hosts

Switches

Templates

Zabbix servers

New host group

Agent interfaces

IP address

127.0.0.1

DNS name

Connect to

IP

DNS

Port

10050

Default

☒

Remove

Add

SNMP interfaces

Add

JMX interfaces

Add

IPMI interfaces

Add

Monitored by proxy

(no proxy)

Status

Monitored

Save

Cancel

The bare minimum to enter here is:

Host name

- Enter a host name. Alphanumerics, spaces, dots, dashes and underscores are allowed.

Groups

- Select one or several groups from the right hand side selectbox and click on « to move them to the 'In groups' selectbox.

Note:

All access permissions are assigned to host groups, not individual hosts. That is why a host must belong to at least one group.

IP address

- Enter the IP address of the host. Note that if this is the Zabbix server IP address, it must be specified in the Zabbix agent configuration file 'Server' directive.

Other options will suit us with their defaults for now.

When done, click **Save**. Your new host should be visible in the hostlist.

Note:

If the Z icon in the *Availability* column is red, there is some error with communication - move your mouse cursor over it to see the error message. If that icon is gray, no status update has happened so far. Check that Zabbix server is running, and try refreshing the page later as well.

3 New item

Overview

In this section you will learn how to set up an item.

Items are the basis of gathering data in Zabbix. Without items, there is no data - because only an item defines a single metric or what data to get off of a host.

Adding item

All items are grouped around hosts. That is why to configure a sample item we go to *Configuration* → *Hosts* and find the 'New host' we have created.

The *Items* link in the row of 'New host' should display a count of '0'. Click on the link, and then click on *Create item*. This will present us with an item definition form.

Item :

Host

Name

Type

Key

Host interface

Type of information

Units

Use custom multiplier ☐

Update interval (in sec)

Flexible intervals

Interval	Period	Action
No flexible intervals defined.		

New flexible interval

Interval (in sec)	Period	Action
50	1-7,00:00-24:00	Add

Keep history (in days)

Keep trends (in days)

Store value

Show value [show value mappings](#)

New application

Applications

Populates host inventory field

Description

Status

For our sample item, the essential information to enter is:

Name

- Enter *CPU Load* as the value. This will be the item name displayed in lists and elsewhere.

Key

- Enter *system.cpu.load* as the value. This is a technical name of an item that identifies the type of information that will be gathered. The particular key is just one of **pre-defined keys** that come with Zabbix agent.

Type of information


- Select *Numeric (float)* here. This attribute defines the format of expected data.

Note:

You may also want to reduce the amount of days **item history** will be kept, to 7 or 14. This is good practice to relieve the database from keeping lots of historical values.

Other options will suit us with their defaults for now.

When done, click **Save**. The new item should appear in the itemlist. Click on *Details* above the list to view what exactly was done.



Seeing data

With an item defined, you might be curious if it is actually gathering data. For that, go to *Monitoring* → *Latest data*, click on the + before - **other** - and expect your item to be there and displaying data.

Name	Last check	Last value	Change	History
- other - (1 Items)				
CPU Load	05 Jan 2012 14:48:38	0.47	+0.37	Graph

With that said, first data may take up to 60 seconds to arrive. That, by default, is how often the server reads configuration changes and picks up new items to execute.

If you see no value in the 'Change' column, maybe only one value has been received so far. Wait 30 seconds for another value to arrive.


If you do not see information about the item as in the screenshot, make sure that:

- you entered item 'Key' and 'Type of information' fields exactly as in the screenshot
- both agent and server are running
- host status is 'Monitored' and its availability icon is green
- host is selected in the host dropdown, item is active

Graphs

With the item working for a while, it might be time to see something visual. **Simple graphs** are available for any monitored numeric item without any additional configuration. These graphs are generated on runtime.

To view the graph, go to *Monitoring* → *Latest data* and click on the 'Graph' link next to the item.



4 New trigger

Overview

In this section you will learn how to set up a trigger.

Items only collect data. To automatically evaluate incoming data we need to define triggers. A trigger contains an expression that defines a threshold of what is an acceptable level for the data.

If that level is surpassed by the incoming data, a trigger will "fire" or go into a 'Problem' state - letting us know that something has happened that may require attention. If the level is acceptable again, trigger returns to an 'Ok' state.

Adding trigger

To configure a trigger for our item, go to *Configuration* → *Hosts*, find 'New host' and click on *Triggers* next to it and then on *Create trigger*. This presents us with a trigger definition form.

Trigger

Dependencies

Name

CPU load too high on 'New host' for 3 minutes

Expression

{New host:system.cpu.load.avg(180)}>2

Add

Expression constructor

Multiple PROBLEM events generation

☐

Description

URL

Severity

Not classified

Information

Warning

Average

High

Disaster

Enabled

☒

For our trigger, the essential information to enter here is:

Name

- Enter *CPU load too high on 'New host' for 3 minutes* as the value. This will be the trigger name displayed in lists and elsewhere.

Expression

- Enter: `{New host:system.cpu.load.avg(180)}>2`

This is the trigger expression. Make sure that the expression is entered right, down to the last symbol. The item key here (system.cpu.load) is used to refer to the item. This particular expression basically says that the problem threshold is exceeded when the CPU load average value for 3 minutes is over 2. You can learn more about the [syntax of trigger expressions](#).

When done, click **Save**. The new trigger should appear in the trigger list.

Displaying trigger status

With a trigger defined, you might be interested to see its status.

For that, go to *Monitoring* → *Triggers*. After 3 minutes or so (we asked to evaluate a 3-minute average after all) your trigger should appear there, presumably with a green 'OK' flashing in the 'Status' column.

	Severity	Status	Info	Last change	Age	Duration	Acknowledged	Host	Name	Comments
	Not classified	OK		06 Jan 2012 14:06:38	9m 43s		Acknowledged	New host	CPU load too high on 'New host' for 3 minutes	Add

The flashing indicates a recent change of trigger status, one that has taken place in the last 30 minutes.

If a red 'PROBLEM' is flashing there, then obviously the CPU load has exceeded the threshold level you defined in the trigger.

5 Receiving problem notification

Overview

In this section you will learn how to set up alerting in the form of notifications in Zabbix.

With items collecting data and triggers designed to "fire" upon problem situations, it would also be useful to have some alerting mechanism in place that would notify us about important events even when we are not directly looking at Zabbix frontend.

This is what notifications do. E-mail being the most popular delivery method for problem notifications, we will learn how to set up an e-mail notification.

E-mail settings

Initially there are several predefined notification **delivery methods** in Zabbix. **E-mail** is one of those.

To configure e-mail settings, go to *Administration* → *Media types* and click on *Email* in the list of pre-defined media types.

Media types

Displaying 1 to 3 of 3 found

☐	Description 	Type	Used in actions	Details
☐	Email	Email	-	SMTP server: "mail.company.com", SMTP helo: "company.com", SMTP email: "zabbix@company.com"
☐	Jabber	Jabber	-	Jabber identifier: "jabber@company.com"
☐	SMS	SMS	-	GSM modem: "/dev/ttyS0"

This will present us with the e-mail settings definition form.

Media

Description

Email

Type

Email

SMTP server

mail.company.com

SMTP helo

company.com

SMTP email

zabbix@company.com

Save

Delete

Cancel

Set the values of SMTP server, SMTP helo and SMTP e-mail to the appropriate for your environment.

Note:

'SMTP email' will be used as the 'From' address for the notifications sent from Zabbix.

Press **Save** when ready.

Now you have configured 'Email' as a working media type. A media type must be linked to users by defining specific delivery addresses (like we did when **configuring a new user**), otherwise it will not be used.

New action

Delivering notifications is one of the things **actions** do in Zabbix. Therefore, to set up a notification, go to *Configuration* → *Actions* and click on *Create action*.

ActionConditionsOperations

Name

Test action

Default escalation period (minimum 60 seconds)

3600 (seconds)

Default subject

{TRIGGER.STATUS}: {TRIGGER.NAME}

Default message

Trigger: {TRIGGER.NAME}

Trigger status: {TRIGGER.STATUS}

Trigger severity: {TRIGGER.SEVERITY}

Trigger URL: {TRIGGER.URL}

Item values:

Recovery message

☐

Enabled

☒

Save

Clone

Delete

Cancel

In this form, enter a name for the action.

{TRIGGER.STATUS} and {TRIGGER.NAME} macros (or variables), visible in the *Default subject* and *Default message* fields, will be replaced with the actual trigger status and trigger name values.

In the most simple case, if we do not add any more specific **conditions**, the action will be taken upon any trigger change from 'Ok' to 'Problem'.

We still should define what the action should do - and that is done in the *Operations* tab. Click on *New* in there, which opens a new operation form.

ActionConditionsOperations

Action operations

☐ Steps Details Period (sec) Delay Action

No operations defined.

Operation details

Step

From

1

To

1 (0 - infinitely)

Escalation period

0 (minimum 60 seconds, 0 - use action default)

Operation type

Send message

Send to User groups

Add

Send to Users

user Remove

Add

Send only to

Email

Default message

☒

Conditions

No conditions defined.

New

Add Cancel

Save

Clone

Delete

Cancel

Here, click on *Add* in the *Send to Users* block and select the user ('user') we have defined. Select 'Email' as the value of *Send only*

to. When done with this, click on *Add*.

That is all for a simple action configuration, so click *Save* in the action form.

Receiving notification

Now, with delivering notifications configured it would be fun to actually receive one. To help with that, we might on purpose increase the load on our host - so that our **trigger** "fires" and we receive a problem notification.

Open the console on your host and run:

```
cat /dev/urandom | md5sum
```

You may run one or several of [these processes](#).

Now go to *Monitoring* → *Latest data* and see how the values of 'CPU Load' have increased. Remember, for our trigger to *fire*, the 'CPU Load' value has to go over '2' for 3 minutes running. Once it does:

- in *Monitoring* → *Triggers* you should see the trigger with a flashing 'Problem' status
- you should receive a problem notification in your e-mail

Attention:

If notifications do not work:

- verify once again that both the e-mail settings and the action have been configured properly
- make sure the user you created has at least read permissions on the host which generated the event, as noted in the *Adding user* step. The user, being part of the 'Network administrators' user group must have at least read access to 'Linux servers' host group that our host belongs to.
- Additionally, you can check out the action log by going to *Administration* → *Audit*, and choosing *Actions* in the dropdown, located in the upper right corner.

6 New template

Overview

In this section you will learn how to set up a template.

Previously we learned how to set up an item, a trigger and how to get a problem notification for the host.

While all of these steps offer a great deal of flexibility in themselves, it may appear like a lot of steps to take if needed for, say, a thousand hosts. Some automation would be handy.

This is where templates come to help. Templates allow to group useful items, triggers and other entities so that those can be reused again and again by applying to hosts in a single step.

When a template is linked to a host, the host inherits all entities of the template. So, basically a pre-prepared bunch of checks can be applied very quickly.

Adding template

To start working with templates, we must first create one. To do that, in *Configuration* → *Templates* click on *Create*. This will present us with a template configuration form.

The required parameters to enter here are:

Template name

- Enter a template name. Alpha-numericals, spaces and underscores are allowed.

Groups

- Select one or several groups from the right hand side selectbox and click on « to move them to the 'In groups' selectbox. The template must belong to a group.

When done, click Save. Your new template should be visible in the list of templates.

As you may see, the template is there, but it holds nothing in it - no items, triggers or other entities.

Adding item to template

To add an item to the template, go to the item list for 'New host'. In *Configuration* → *Hosts* click on *Items* next to 'New host'.

Then:

- mark the checkbox of the 'CPU Load' item in the list
- select *Copy selected to...* in the dropdown below the list and click on *Go*
- select the template to copy item to

- click on *Copy*


If you now go to *Configuration → Templates*, 'New template' should have one new item in it.

We will stop at one item only for now, but similarly you can add any other items, triggers or other entities to the template until it's a fairly complete set of entities for given purpose (monitoring OS, monitoring single application).

Linking template to host

With a template ready, it only remains to add it to a host. For that, go to *Configuration → Hosts*, click on 'New host' to open its property form and go to the **Templates** tab.

There, click on *Add*, mark the template we have created ('New template') and click on *Select*. The template should appear in the form.



Click *Save* in the form to save the changes. The template is now added to the host, with all entities that it holds.

As you may have guessed, this way it can be applied to any other host as well. Any changes to the items, triggers and other entities at the template level will propagate to the hosts the template is linked to.

Linking pre-defined templates to hosts

As you may have noticed, Zabbix comes with a set of predefined templates for various OS, devices and applications. To get started with monitoring very quickly, you may link the appropriate one of them to a host, but beware that these templates need to be fine-tuned for your environment. Some checks may not be needed, and polling intervals may be way too frequent.

More information about [templates](#) is available.

5. Zabbix appliance

As an alternative to setting up manually or reusing existing server for Zabbix, users may [download Zabbix appliance](#).

To get started, boot the appliance and point your browser at the IP it has received over DHCP.

|<| |<| |-|

Zabbix appliance versions are based upon the following OpenSUSE versions:

Zabbix appliance version	OpenSUSE version
2.0.0	12.1
2.0.2	
2.0.3	
2.0.4	12.2*
2.0.5	

* Initial 2.0.4 appliance builds were with OpenSUSE 12.1, final builds are with OpenSUSE 12.2.

It is available in the following formats:

- vmdk (VMware/Virtualbox)
- OVF (Open Virtualisation Format)
- VHD (Hyper-V)**
- CD iso
- HDD/flash image
- [Preload ISO](#)
- Xen guest

* Since appliance 2.0.5.

It has Zabbix server configured and running on MySQL, as well as frontend available.

The appliance has been built using [SUSE Studio](#).

1 Changes to SUSE configuration There are some changes applied to the base OpenSUSE configuration.

1.1 MySQL configuration changes

- Binary log is disabled;
- InnoDB is configured to store data for each table in a separate file.

1.2 Using a static IP address

By default the appliance uses DHCP to obtain IP address. To specify a static IP address:

- Log in as root user;
- Open file `/etc/sysconfig/network/ifcfg-eth0` in your favourite editor;
- Set **BOOTPROTO** variable to **static**;
- Set **IPADDR**, **NETMASK** and any other parameters as required for your network;
- Create file `/etc/sysconfig/network/routes`. For the default route, use **default 192.168.1.1 - -** (replacing with your gateway address).
- Run the command **rcnetwork restart**.

To configure DNS, add nameserver entries in `/etc/resolv.conf`, specifying each nameserver on its own line: **nameserver 192.168.1.2**.

Alternatively, just use **yast** configuration utility to update network settings.

1.3 Changing time zone

By default the appliance uses UTC for the system clock. To change the time zone, copy appropriate file from `/usr/share/zoneinfo` to `/etc/localtime`, for example:

```
cp /usr/share/zoneinfo/Europe/Riga /etc/localtime
```

1.4 Other changes

- Network is configured to use DHCP to obtain IP address;
- Utility **fping** is set to have permissions 4710 and is owned by group **zabbix** - suid and only allowed to be used by zabbix group;
- ntpd configured to synchronise to the public pool servers;
- Various basic utilities have been added that could make working with Zabbix and monitoring in general easier.

2 Zabbix configuration Appliance Zabbix setup has the following passwords and other configuration changes:

2.1 Passwords

System:

- root:zabbix
- zabbix:zabbix

Database:

- root:zabbix
- zabbix:zabbix

Zabbix frontend:

- Admin:zabbix

Attention:

If you change frontend password, do not forget to update password setting web monitoring (*Configuration → Web*).

To change the database user password it has to be changed in the following locations:

- MySQL;
- zabbix_server.conf;
- zabbix.conf.php.

2.2 File locations

- Configuration files are placed in **/etc**.
- Zabbix logfiles are placed in **/var/log/zabbix**.
- Zabbix frontend is placed in **/usr/share/zabbix**.
- Home directory for user **zabbix** is **/var/lib/zabbix**.

2.3 Changes to Zabbix configuration

- Server name for Zabbix frontend set to "Zabbix 2.0 Appliance";
- Frontend timezone is set to Europe/Riga, Zabbix home (this can be modified in **/etc/php5/apache2/php.ini**);
- Disabled triggers and web scenarios are shown by default to reduce confusion.

2.4 Preserving configuration

If you are running live CD version of the appliance or for some other reason can't have persistent storage, you can create a backup of whole database, including all configuration and gathered data.

To create the backup, run:

```
mysqldump zabbix | bzip2 -9 > dbdump.bz2
```

Now you can transfer file **dbdump.bz2** to another machine.

To restore from the backup, transfer it to the appliance and execute:

```
bzcat dbdump.bz2 | mysql zabbix
```

Attention:

Make sure that Zabbix server is stopped while performing the restore.

3 Frontend access Access to frontend by default is allowed from:

- 127.0.0.1
- 192.168.0.0/16
- 10.0.0.0/8
- ::1

Root (/) is redirected to /zabbix on the webserver, thus frontend can be accessed both as **http://<host>** and **http://<host>/zabbix**.

This can be customised in **/etc/apache2/conf.d/zabbix.conf**. You have to restart webserver after modifying this file. To do so, log in using SSH as **root** user and execute:

```
service apache2 restart
```

4 Firewall By default, only two ports are open - 22 (SSH) and 80 (HTTP). To open additional ports - for example, Zabbix server and agent ports - modify iptables rules with **SuSEfirewall2** utility:

```
SuSEfirewall2 open EXT TCP zabbix-trapper zabbix-agent
```

Then reload the firewall rules:

```
SuSEfirewall2 stop
SuSEfirewall2 start
```

5 Monitoring capabilities Zabbix server is compiled with support for the following:

- SNMP;
- IPMI;
- Web monitoring;
- SSH2;
- IPv6.

In the provided configuration Zabbix server itself is monitored with the help of locally installed agent for some base parameters, additionally Zabbix frontend is monitored as well using web monitoring.

|<|<|<|

Note:

Note that web frontend monitoring logs in - this can add lots of entries to the audit log.

6 Naming, init and other scripts Appropriate init scripts are provided. To control Zabbix server, use any of these:

```
service zabbix_server status
rczabbix_server status
/etc/init.d/zabbix_server status
```

Replace **server** with **agentd** for Zabbix agent daemon.

6.1 Scheduled scripts

There is a scheduled script, run from the crontab every 10 minutes that restarts Zabbix server if it is not running, **/var/lib/zabbix/bin**. It logs timestamped problems and starting attempts at **/var/log/zabbix/server_problems.log**.

Attention:

Make sure to disable this crontab entry if stopping of Zabbix server is desired.

6.2 Increasing available diskpace

Warning:

Create a backup of all data before attempting any of the steps.

Available diskpace on the appliance might not be sufficient. In that case it is possible to expand the disk. To do so, first expand the block device in your virtualisation environment, then follow these steps.

Start *fdisk* to change the partition size. As *root*, execute:

```
fdisk /dev/sda
```

This will start *fdisk* on disk *sda*. Next, switch to sectors by issuing:

```
u
```

Attention:

Don't disable DOS compatibility mode by entering **c**. Proceeding with it disabled will damage the partition.

Then delete the existing partition and create new one with desired size. In majority of cases you will accept the available maximum, which will expand the filesystem to whatever size you made available for the virtual disk. To do so, enter the following sequence in *fdisk* prompt:

```
d
n
p
1
(accept default 63)
(accept default max)
```

If you wish to leave some space for additional partitions (swap etc), you can enter another value for *last sector*. When done, save the changes by issuing:

```
w
```

Reboot the virtual machine (as the partition we modified is in use currently). After reboot, filesystem resizing can take place.

```
resize2fs /dev/sda1
```

That's it, filesystem should be grown to the partition size now.

7 Format-specific notes 7.1 Xen

To use images in Xen server, run:

```
xm create -c file-with-suffix.xenconfig
```

See the following pages for more information on using Xen images:

- http://en.opensuse.org/openSUSE:How_to_use_downloaded_SUSE_Studio_appliances#Using_Xen_guests
- http://old-en.opensuse.org/SUSE_Studio_Xen_Howtos

Converting image for XenServer

To use Xen images with Citrix XenServer you have to convert the disk image. To do so:

- Create a virtual disk which is at least as large as the image
- Find out the UUID for this disk

```
xe vdi-list params=all
```

- If there are lots of disks, they can be filtered by name parameter *name-label*, as assigned when creating the virtual disk
- Import the image

```
xe vdi-import filename="image.raw" uuid="<UUID>"
```

Instructions from Brian Radford blog.

7.2 VMware

The images in *vmdk* format are usable directly in VMware Player, Server and Workstation products. For use in ESX, ESXi and vSphere they must be converted using [VMware converter](#).

7.3 HDD/flash image (raw)

See http://en.opensuse.org/openSUSE:SUSE_Studio_Disc_Image_Howtos for more information on disk images.

8 Known issues 8.1 Extracting on Windows

Windows archive management software is known to mishandle the appliance archives. If extraction fails, try different software. Open source tool [7-zip](#) might work.

8.2 For the 2.0.0 appliance

- Web monitoring is not preconfigured.
- Locales other than English do not work. To enable other locales, run the following commands:

```
/usr/share/zabbix/locale/make_mo.sh
rcapache2 restart
```

- Static IP address might not work because of a [systemd problem](#). Updating the system with **zypper up** should solve it.

8.3 For the 2.0.2 appliance

- Web monitoring is not preconfigured.
- Locales other than English do not work. See 2.0.0 notes for solution.

8.4 For the 2.0.4 appliance

- SSH daemon could not be started in the initial appliance builds. Please download latest builds (post-2013.01.27).
- Web monitoring is not preconfigured.

6. Configuration

Please use the sidebar to access content in the Configuration section.

1 Configuring a template

Overview

Configuring a template requires that you first create a template by defining its general parameters and then you add entities (items, triggers, graphs etc.) to it.

Creating a template

To create a template, do the following:

- Go to *Configuration → Templates*
- Click on *Create template*
- Edit template attributes

The **Template** tab contains general template attributes.

Template attributes:

Parameter	Description
<i>Template name</i>	Unique template name.
<i>Visible name</i>	If you set this name, it will be the one visible in lists, maps, etc.
<i>Groups</i>	Host/template groups the template belongs to.
<i>New group</i>	A new group can be created to hold the template.
<i>Hosts/Templates</i>	Ignored, if empty. List of hosts/templates the template is applied to.

The **Linked templates** tab allows you to link one or more "nested" templates to this template. All entities (items, triggers, graphs etc.) will be inherited from the linked templates.

To link a new template, click on *Add*. To unlink a template, use one of the two options:

- *Unlink* - unlink the template, but preserve its items, triggers and graphs
- *Unlink and clear* - unlink the template and remove all its items, triggers and graphs

The **Macros** tab allows you to define template-level **user macros**.

Buttons:

Save

Save the template. The saved template should appear in the list.

Clone

Create another template based on the properties of the current template, including the entities (items, triggers, etc) inherited from linked templates.

Full clone	Create another template based on the properties of the current template, including the entities (items, triggers, etc) both inherited from linked templates and directly attached to the current template.
Delete	Delete the template; entities of the template (items, triggers, etc) remain with the linked hosts.
Delete and clear	Delete the template and all its entities from linked hosts.
Cancel	Cancel the editing of template properties.

With a template created, it is time to add some entities to it.

Attention:

Items have to be added to a template first. Triggers and graphs cannot be added without the corresponding item.

Adding items, triggers, graphs

To add items to the template, do the following:

- Go to *Configuration → Hosts (or Templates)*
- Click on *Items* in the row of the required host/template
- Mark the checkboxes of items you want add to the template
- Select *Copy selected to...* below the item list and click on *Go*
- Select the template (or group of templates) the items should be copied to and click on *Copy*

All the selected items should be copied to the template.

Adding triggers and graphs is done in similar fashion (from the list of triggers and graphs respectively), again, keeping in mind that they can only be added if the required items are added first.

Adding screens

To add screens to a template in *Configuration → Templates*, do the following:

- Click on *Screens* in the row of the template
- Configure a screen following the usual method of [configuring screens](#)

Attention:

The elements that can be included in a template screen are: simple graph, custom graph, clock, plain text, URL.

Configuring low-level discovery rules

See the [low-level discovery](#) section of the manual.

2 Linking/unlinking

Overview

Linking is a process whereby templates are applied to hosts, whereas unlinking removes the association with the template from a host.

Attention:

Templates are linked directly to individual hosts and not to host groups. Simply adding a template to a host group will not link it. Host groups are used only for logical grouping of hosts and templates.

Linking a template

To link a template to the host, do the following:

- Go to *Configuration → Hosts*
- Click on the required host and switch to the *Templates* tab

- Click on *Add*
- Select one or several templates in the popup window
- Click on *Save* in the host attributes form

The host will now have all the items, triggers, graphs, applications, screens and low-level discovery rules of the template.

Attention:

Linking multiple templates to the same host will fail if in those templates there are items with the same item key or applications with the same name. And, as triggers and graphs use items, they cannot be linked to a single host from multiple templates either, if using identical item keys.

When entities (items, triggers, graphs etc.) are added from the template:

- previously existing identical entities on the host are updated as entities of the template
- entities from the template are added
- any directly linked entities that, prior to template linkage, existed only on the host remain untouched

In the lists, all entities from the template now are prefixed by the template name, indicating that these belong to the particular template. The template name itself (in grey text) is a link allowing to access the list of those entities on the template level.

If some entity (item, trigger, graph etc.) is not prefixed by the template name, it means that it existed on the host before and was not added by the template.

Entity uniqueness criteria

When adding entities (items, triggers, graphs etc.) from a template it is important to know what of those entities already exist on the host and need to be updated and what entities differ. The uniqueness criteria for deciding upon the sameness/difference are:

- for items - the item key
- for triggers - trigger name and expression
- for custom graphs - graph name and its items
- for applications - application name

Linking templates to several hosts

There are some ways of mass-applying templates (to many hosts at once):

- To link a template to many hosts, in *Configuration* → *Templates*, click on the template, then select hosts from the respective group in the *Other* box, click on « and save the template.

Vice versa, if you select the linked hosts in the *In* box, click on » and save the template, you unlink the template from these hosts (while the hosts will still inherit the items, triggers, graphs etc. from the template).

- To update template linkage of many hosts, in *Configuration* → *Hosts* select some hosts by marking their checkboxes, then choose **Mass update** below the list, click on *Go* and then from all the options select to link additional templates:

Link templates ☒

Name	Action
C_Template_Linux	Remove
Add	
☒ Replace	
☐ Clear when unlinking	

Click on *Add* to select the templates to link. The *Replace* option allows to unlink the previously linked templates, before linking the new ones. When choosing to replace, there is also the *Clear when unlinking* option. Marking that will remove all the entities (items, triggers, graphs, etc) previously linked to the hosts by the linked templates.

Note:

Zabbix offers a sizable set of predefined templates. You can use these for reference, but beware of using them unchanged in production as they may contain too many items and poll for data too often. If you feel like using them, finetune them to fit you real needs.

Editing linked entities

If you try to edit an item or trigger that was linked from the template, you may realize that many key options are disabled for editing. This makes sense as the idea of templates is that things are edited in one-touch manner on the template level. However,

you still can, for example, enable/disable an item on the individual host and set the update interval, history length and some other parameters.

If you want to edit the entity fully, you have to edit it on the template level (template level shortcut is displayed in the form name), keeping in mind that these changes will affect all hosts that have this template linked to them.

Unlinking a template

To unlink a template from a host, do the following:

- Go to *Configuration* → *Hosts*
- Click on the required host and switch to the *Templates* tab
- Click on *Unlink* or *Unlink and clear* next to the template to unlink
- Click on *Save* in the host attributes form

Choosing the *Unlink* option will simply remove association with the template, while leaving all its entities (items, triggers, graphs etc.) with the host.

Choosing the *Unlink and clear* option will remove both the association with the template and all its entities (items, triggers, graphs etc.).

3 Nesting

Overview

Nesting is a way of one template encompassing one or more other templates.

As it makes sense to separate out on individual templates entities for various services, applications etc. you may end up with quite a few templates all of which may need to be linked to quite a few hosts. To simplify the picture, it is possible to link some templates together, in one "nested" template.

The benefit of nesting is that then you have to link only the one template to the host and the host will inherit all entities of the linked templates automatically.

Configuring a nested template

If you want to link some templates, to begin with you can take an existing template or a new one, then:

- Open the template properties form
- Look for the *Linked templates* tab
- Click on *Add*, select the templates to link in the popup window
- Click on *Save* in the template properties form

Now the template should have all the entities (items, triggers, custom graphs etc.) of the linked templates.

To unlink any of the linked templates, in the same form use the *Unlink* or *Unlink and clear* buttons and click on *Save*.

Choosing the *Unlink* option will simply remove the association with the other template, while not removing all its entities (items, triggers, graphs etc.).

Choosing the *Unlink and clear* option will remove both the association with the other template and all its entities (items, triggers, graphs etc.).

Permission issues

- You may have a setup where an Admin level user has *Read-write* access to some Template A while not having *Read-write* access to Template B that holds Template A in a nested setup. In this case, an item created on Template A, while inherited by the hosts of Template A, **will not** be inherited by the hosts of Template B. Thus, creating a trigger for such an item will fail altogether, because of missing corresponding items on hosts of Template B.

1 Hosts and host groups

What is a "host"?

Typical Zabbix hosts are the devices you wish to monitor (servers, workstations, switches, etc.).

Creating hosts is one of the first monitoring tasks in Zabbix. For example, if you want to monitor some parameters on a server "x", you must first create a host called, say, "Server X" and then you can look to add monitoring items to it.

Hosts are organized into host groups.

Proceed to **creating and configuring a host**.

1 Configuring a host

Overview

To configure a host in Zabbix frontend, do the following:

- Go to: *Configuration* → *Hosts*
- Click on *Create* to the right (or on the host name to edit an existing host)
- Enter parameters of the host in the form

You can also use the *Clone* and *Full clone* buttons in the form of an existing host to create a new host. Clicking on *Clone* will retain all host parameters and template linkage (keeping all entities from those templates). *Full clone* will additionally retain directly attached entities (applications, items, triggers, graphs and low-level discovery rules).

Note: When a new host is cloned, it will retain all template entities as they are originally on the template. Any changes to those entities made on the existing host level (such as changed item interval, modified regular expression or added prototypes to the low-level discovery rule) will not be cloned to the new host; instead they will be as on the template.

Configuration

The **Host** tab contains general host attributes:

Host

Templates

IPMI

Macros

Host inventory

Host name

Zabbix server

Visible name

Groups

In groups

Zabbix servers

Other groups

Discovered hosts

Linux servers

Templates

New host group

Agent interfaces

IP address

DNS name

Connect to

Port

Default

192.168.3.2

IP

DNS

10050

☒

Remove

Add

SNMP interfaces

Add

JMX interfaces

Add

IPMI interfaces

Add

Monitored by proxy

(no proxy)

Status

Monitored

Save

Clone

Full clone

Delete

Cancel

Parameter	Description
Host name	Enter a unique host name. Alphanumerics, spaces, dots, dashes and underscores are allowed. <i>Note:</i> With Zabbix agent running on the host you are configuring, the agent configuration file parameter <i>Hostname</i> must have the same value as the host name entered here. The name in the parameter is needed in the processing of active checks .
Visible name	If you set this name, it will be the one visible in lists, maps, etc. This attribute has UTF-8 support.
Groups	Select host groups the host belongs to. A host must belong to at least one host group.
New host group	A new group can be created and linked to the host. Ignored, if empty.

Parameter	Description
<i>Interfaces</i>	Several host interface types are supported for a host: <i>Agent</i> , <i>SNMP</i> , <i>JMX</i> and <i>IPMI</i> . To add a new interface, click on <i>Add</i> in the <i>Interfaces</i> block and enter <i>IP/DNS</i> , <i>Connect to</i> and <i>Port</i> info. <i>Note</i> : Interfaces that are used in any items cannot be removed and link <i>Remove</i> is greyed out for them.
<i>IP address</i>	Host IP address (optional).
<i>DNS name</i>	Host DNS name (optional).
<i>Connect to</i>	Clicking the respective button will tell Zabbix server what to use to retrieve data from agents: IP - Connect to the host IP address (recommended) DNS - Connect to the host DNS name
<i>Port</i>	TCP/UDP port number. Default values are: 10050 for Zabbix agent, 161 for SNMP agent, 12345 for JMX and 623 for IPMI.
<i>Default</i>	Check the radio button to set the default interface.
<i>Monitored by proxy</i>	The host can be monitored either by Zabbix server or one of Zabbix proxies: (no proxy) - host is monitored by Zabbix server Proxy name - host is monitored by Zabbix proxy "Proxy name"
<i>Status</i>	Host status: Monitored - Host is active, ready to be monitored Not monitored - Host is not active, thus not monitored

The **Templates** tab allows you to link **templates** to the host. All entities (items, triggers, graphs and applications) will be inherited from the template.

To link a new template, click on *Add*. To unlink a template, use one of the two options:

- *Unlink* - unlink the template, but preserve its items, triggers and graphs
- *Unlink and clear* - unlink the template and remove all its items, triggers and graphs

The **IPMI** tab contains IPMI management attributes.

Parameter	Description
<i>Authentication algorithm</i>	Select the authentication algorithm.
<i>Privilege level</i>	Select the privilege level.
<i>Username</i>	User name for authentication.
<i>Password</i>	Password for authentication.

The **Macros** tab allows you to define host-level **user macros**.

The **Host inventory** tab allows you to manually enter **inventory** information for the host. You can also select to enable *Automatic* inventory population, or disable inventory population for this host.

Configuring a host group

To configure a host group in Zabbix frontend, do the following:

- Go to: *Configuration* → *Host groups*
- Click on *Create Group* in the upper right corner of the screen
- Enter parameters of the group in the form

Parameter	Description
<i>Group name</i>	Enter a unique host group name. The name must be unique within a Zabbix node.
<i>Hosts</i>	Select hosts, members of the group. A host group may have zero, one or more hosts.

2 Inventory

Overview

You can keep the inventory of networked devices in Zabbix.

There is a special *Inventory* menu in the Zabbix frontend. However, you will not see any data there initially and it is not where you enter data. Building inventory data is done manually when configuring a host or automatically by using some automatic population options.

Building inventory

Manual mode

When **configuring a host**, in the *Host inventory* tab you can enter such details as the type of device, serial number, location, responsible person, etc - data that will populate inventory information.

If a URL is included in host inventory information and it starts with 'http' or 'https', it will result in a clickable link in the *Inventory* section.

Automatic mode

Host inventory can also be populated automatically. For that to work, when configuring a host the inventory mode in the *Host inventory* tab must be set to *Automatic*.

Then you can **configure host items** to populate any host inventory field with their value, indicating the destination field with the respective attribute (called *Item will populate host inventory field*) in item configuration.

Items that are especially useful for automated inventory data collection:

- system.hw.chassis[full|type|vendor|model|serial] - default is [full], root permissions needed
- system.hw.cpu[all|cpunum,full|maxfreq|vendor|model|curfreq] - default is [all,full]
- system.hw.devices[pci|usb] - default is [pci]
- system.hw.macaddr[interface,short|full] - default is [all,full], interface is regexp
- system.sw.arch
- system.sw.os[name|short|full] - default is [name]
- system.sw.packages[package,manager,short|full] - default is [all,all,full], package is regexp

Inventory overview

The overview of all inventory data entered is available in the *Inventory* menu.

In *Inventory* → *Overview* you can group the display of available data by various fields of the inventory.

In *Inventory* → *Hosts* you can see all hosts that have inventory information. Clicking on the host name will reveal all the details in a form.

Inventory macros

There are host inventory macros {INVENTORY.*} available for use in notifications, for example:

"Server in {INVENTORY.LOCATION1} has a problem, responsible person is {INVENTORY.CONTACT1}, phone number {INVENTORY.POC.PRIMARY.PHONE.A1}."

Attention:

{PROFILE.*} macros from previous Zabbix versions are still supported but it's highly recommended to change those to {INVENTORY.*}

For more details, see the [Macros supported by location](#) page.

3 Mass update

Overview

Sometimes you may want to change some attribute for a number of hosts at once. Instead of opening each individual host for editing, you may use the mass update function for that.

Using mass update

To mass-update some hosts, do the following:

- Mark the checkboxes of the hosts to update in the list
- Select *Mass update* from the dropdown below and click on *Go*
- Mark the checkboxes of the attributes to update
- Enter new values for the attributes and click on *Save*

Mass update

Replace host groups☐

Original

New host group☐

Original

Monitored by proxy☒

Netherlands

Status☐

Original

Link templates☒

Name	Action
C_Template	Remove
Add	
☐ Replace	
☐ Clear when unlinking	

IPMI authentication algorithm☐

Original

IPMI privilege level☐

Original

IPMI username☐

Original

IPMI password☐

Original

Inventory mode☐

Original

Update

Cancel

2 Items

Overview

Items are the ones that gather data from a host.

Once you have configured a host, you need to add some monitoring items to start getting actual data.

An item is an individual metric. One way of quickly adding many items is to attach one of the predefined templates to a host. For optimized system performance though, you may need to fine-tune the templates to have only as many items and as frequent monitoring as is really necessary.

In an individual item you specify what sort of data will be gathered from the host.

For that purpose you use the **item key**. Thus an item with the key name **system.cpu.load** will gather data of the processor load, while an item with the key name **net.if.in** will gather incoming traffic information.

To specify further parameters with the key, you include those in square brackets after the key name. Thus, **system.cpu.load[avg5]** will return processor load average for the last 5 minutes, while **net.if.in[eth0]** will show incoming traffic in the interface eth0.

Note:

For all supported item types and item keys, see individual sections of **item types**.

Proceed to **creating and configuring an item**.

1 Creating an item

Overview

To create an item in Zabbix frontend, do the following:

- Go to: *Configuration* → *Hosts*
- Click on *Items* in the row of the host
- Click on *Create item* in the upper right corner of the screen
- Enter parameters of the item in the form

Configuration

Item : Available memory

Host

Zabbix server

Select

Name

Available memory

Type

Zabbix agent

Key

vm.memory.size[available]

Select

Host interface

192.168.3.41 : 10050

Type of information

Numeric (unsigned)

Data type

Decimal

Units

B

Use custom multiplier

☐

1

Update interval (in sec)

60

Flexible intervals

Interval	Period	Action
No flexible intervals defined.		

New flexible interval

Interval (in sec)

50

Period

1-7,00:00-24:00

Add

Keep history (in days)

7

Keep trends (in days)

365

Store value

As is

Show value

As is

[show value mappings](#)

New application

Applications

-None-

CPU

Filesystems

General

Memory

Network interfaces

Populates host inventory field

-None-

Description

Available memory is defined as free+cached+buffers memory.

Status

Enabled

Save

Cancel

Item attributes:

Parameter	Description
<i>Host</i>	Select the host or template.
<i>Name</i>	This is how the item will be named. The following macros can be used: \$1, \$2...\$9 - referring to the first, second... ninth parameter of the item key For example: Free disk space on \$1 If the item key is "vfs.fs.size[/,free]", the description will automatically change to "Free disk space on /"
<i>Type</i>	Item type. See individual item type sections.
<i>Key</i>	Item key. The supported item keys can be found in individual item type sections. The key must be unique within a single host. If key type is 'Zabbix agent', 'Zabbix agent (active)', 'Simple check' or 'Zabbix aggregate', the key value must be supported by Zabbix agent or Zabbix server. See also: the correct key format.

Parameter	Description
<i>Host interface</i>	Select the host interface. This field is available when editing an item on the host level.
<i>Type of information</i>	Type of data as stored in the database after performing conversions, if any. Numeric (unsigned) - 64bit unsigned integer Numeric (float) - floating point number Negative values can be stored. Allowed range (for MySQL): -999999999999.9999 to 999999999999.9999 (double(16,4)). Character - character (string) data limited to 255 bytes Log - log file. Must be set for log*, eventlog item keys. Text - text of unlimited size
<i>Data type</i>	Data type is used for integer items in order to specify the expected data type: Boolean - textual representation translated into either 0 or 1. Thus, 'TRUE' is stored as 1 and 'FALSE' is stored as 0. All values are matched in a case-insensitive way. Currently recognized values are, for: <i>TRUE</i> - true, t, yes, y, on, up, running, enabled, available <i>FALSE</i> - false, f, no, n, off, down, unused, disabled, unavailable Additionally, any non-zero numeric value is considered to be TRUE and zero is considered to be FALSE. Octal - data in octal format Decimal - data in decimal format Hexadecimal - data in hexadecimal format Zabbix will automatically perform the conversion to numeric. The conversion is done by Zabbix server (even when a host is monitored by Zabbix proxy).
<i>Units</i>	If a unit symbol is set, Zabbix will add post processing to the received value and display it with the set unit postfix. By default, if the raw value exceeds 1000, it is divided by 1000 and displayed accordingly. For example, if you set <i>bps</i> and receive a value of 881764, it will be displayed as 881.76 Kbps. Special processing is used for B (byte), Bps (bytes per second) units, which are divided by 1024. Thus, if units are set to B or Bps Zabbix will display: 1 as 1B/1Bps 1024 as 1KB/1KBps 1536 as 1.5KB/1.5KBps Special processing is used if the following time-related units are used: unixtime - translated to "yyyy.mm.dd hh:mm:ss". To translate correctly, the received value must be a <i>Numeric (unsigned)</i> type of information. uptime - translated to "hh:mm:ss" or "N days, hh:mm:ss" For example, if you receive the value as 881764 (seconds), it will be displayed as "10 days, 04:56:04" s - translated to "yyy mmm ddd hhh mmm sss ms"; parameter is treated as number of seconds. For example, if you receive the value as 881764 (seconds), it will be displayed as "10d 4h 56m" Only 3 upper major units are shown, like "1m 15d 5h" or "2h 4m 46s". If there are no days to display, only two levels are displayed - "1m 5h" (no minutes, seconds or milliseconds are shown). Will be translated to "< 1 ms" if the value is less than 0.001. See also the unit blacklist.
<i>Use custom multiplier</i>	If you enable this option, all received values will be multiplied by the integer or floating-point value set in the value field. Use this option to convert values received in KB, MBps, etc into B, Bps. Otherwise Zabbix cannot correctly set prefixes (K, M, G etc).

Parameter	Description
<i>Update interval (in sec)</i>	Retrieve a new value for this item every N seconds. <i>Note:</i> If set to '0', the item will not be polled. However, if a flexible interval also exists with a non-zero value, the item will be polled during the flexible interval duration.
<i>Flexible intervals</i>	You can create exceptions to <i>Update interval</i>. For example: Interval: 10, Period: 1-5,09:00-18:00 - will check the item every 10 seconds during working hours. Interval: 0, Period: 1-7,00:00-7:00 - will disable checking the item at night. Interval: 0, Period: 7-7,00:00-24:00 - will disable checking the item on Sundays. To check an item once per day at a specific time (say, 12:00), set the default <i>Update interval</i> to '0', but specify 60 in the flexible interval and a period like 1-7,12:00-12:01 Up to seven flexible intervals can be defined. If multiple flexible intervals overlap, the smallest <i>Interval</i> value is used for the overlapping period. Note that if the smallest value of overlapping flexible intervals is '0', no polling will take place. Outside the flexible intervals the default update interval is used. See the page about setting time periods for description of the <i>Period</i> format. Note that if the flexible interval equals the length of the period, the item will be checked exactly once. If the flexible interval is greater than the period, the item might be checked once or it might not be checked at all (thus such configuration is not advisable). If the flexible interval is less than the period, the item will be checked at least once. If the flexible interval is set to '0', the item is not polled during the flexible interval period and resumes polling according to the default <i>Update interval</i> once the period is over. <i>Note:</i> Not available for Zabbix agent active items.
<i>Keep history
(in days)</i>	Number of days to keep detailed history in the database. Older data will be removed by the housekeeper. It is recommended to keep the recorded values for the smallest possible number of days to reduce the size of value history in the database. Instead of keeping long history of values, you can keep longer data of trends. See also History and trends.
<i>Keep trends
(in days)</i>	Keep aggregated (hourly min, max, avg, count) detailed history for N days in the database. Older data will be removed by the housekeeper. <i>Note:</i> Keeping trends is not available for non-numeric data - character, log and text. See also History and trends.

Parameter	Description
<i>Store value</i>	As is - no pre-processing Delta (speed per second) - evaluate value as $(\text{value} - \text{prev\_value}) / (\text{time} - \text{prev\_time})$, where <i>value</i> - current value <i>value_prev</i> - previously received value <i>time</i> - current timestamp <i>prev_time</i> - timestamp of previous value This setting is extremely useful to get speed per second for a constantly growing value. If current value is smaller than the previous value, Zabbix discards that difference (stores nothing) and waits for another value. This helps to work correctly with, for instance, a wrapping (overflow) of 32-bit SNMP counters. <i>Note:</i> As this calculation may produce floating point numbers, it is recommended to set the 'Type of information' to <i>Numeric (float)</i>, even if the incoming raw values are integers. This is especially relevant for small numbers where the decimal part matters. If the floating point values are large and may exceed the 'float' field length in which case the entire value may be lost, it is actually suggested to use <i>Numeric (unsigned)</i> and thus trim only the decimal part. Delta (simple change) - evaluate as $(\text{value} - \text{prev\_value})$, where <i>value</i> - current value <i>value_prev</i> - previously received value This setting can be useful to measure a constantly growing value. If the current value is smaller than the previous value, Zabbix discards that difference (stores nothing) and waits for another value.
<i>Show value</i>	Apply value mapping to this item. Value mapping does not change received values, it is for displaying data only. It works with integer items only. For example, "Windows service states".
<i>Log time format</i>	Available for items of type Log only. Supported placeholders: * y: Year (1970-2038) * M: Month (01-12) * d: Day (01-31) * h: Hour (00-23) * m: Minute (00-59) * s: Second (00-59) If left blank the timestamp will not be parsed. For example, consider the following line from the Zabbix agent log file: " 23480:20100328:154718.045 Zabbix agent started. Zabbix 1.8.2 (revision 11211)." It begins with six character positions for PID, followed by date, time, and the rest of the line. Log time format for this line would be "pppppp:yyyyMMdd:hhmmss". Note that "p" and ":" chars are just placeholders and can be anything but "yMdhms".
<i>New application</i>	Enter the name of a new application for the item.
<i>Applications</i>	Link item to one or more existing applications.
<i>Populates host inventory field</i>	You can select a host inventory field that the value of item will populate. This will work if automatic inventory population is enabled for the host.
<i>Description</i>	Enter an item description.

Parameter	Description
<i>Status</i>	Enabled - the item will be processed. Disabled - the item will not be processed. Not supported - the item is not supported. This item will not be processed, however Zabbix may try to periodically set the status of such items to <i>Enabled</i> according to the interval set for refreshing unsupported items.

You can also create an item by opening an existing one, pressing the *Clone* button and then saving under a different name.

Note:

When editing an existing **template** level item on a host level, a number of fields are read-only. You can use the link in the form header and go to the template level and edit them there, keeping in mind that the changes on a template level will change the item for all hosts that the template is linked to.

Unit blacklist

By default, specifying a unit for an item will result in a multiplier prefix being added - for example, value 2048 with unit B would be displayed as 2KB. For a pre-defined, hardcoded list of units this is prevented:

- ms
- RPM
- rpm
- %

Note that both lowercase and uppercase **rpm** (*rpm* and *RPM*) strings are blacklisted.

Unsupported items

An item can become unsupported if its value cannot be retrieved for some reason. Such items are still rechecked at a fixed interval, configurable in **Administration section**.

1 Item key


1.1 Flexible and non-flexible parameters

A flexible parameter is a parameter which accepts an argument. For example, in `vfs.fs.size[*]` the asterisk symbol '*' indicates a flexible parameter. '*' is any string that will be passed as an argument to the parameter. Correct definition examples:

- `vfs.fs.size[/]`
- `vfs.fs.size[/opt]`

1.2 Key format

Item key format, including key parameters, must follow syntax rules. The following illustrations depict the supported syntax. Allowed elements and characters at each point can be determined by following the arrows - if some block can be reached through the line, it is allowed, if not - it is not allowed.



To construct a valid item key, one starts with specifying the key name, then there's a choice to either have parameters or not - as depicted by the two lines that could be followed.

Key name


The key name itself has a limited range of allowed characters, which just follow each other. Allowed characters are:

0-9a-zA-Z_ - .

Which means:


- all numbers;

- all lowercase letters;
- all uppercase letters;
- underscore;
- dash;
- dot.




Key parameters

An item key can have multiple parameters that are comma separated.




Each key parameter can be either a quoted string, an unquoted string or an array.



The parameter can also be left empty, thus using the default value. In that case, the appropriate number of commas must be added if any further parameters are specified. For example, item key **icmpping[,200,500]** would specify that the interval between individual pings is 200 milliseconds, timeout - 500 milliseconds, and all other parameters are left at their defaults.

Parameter - quoted string

If the key parameter is a quoted string, any Unicode character is allowed, and included double quotes must be backslash escaped.




Warning:

To quote item key parameters, use double quotes only. Single quotes are not supported.


Parameter - unquoted string

If the key parameter is an unquoted string, any Unicode character is allowed except comma and right square bracket (]).



Parameter - array

If the key parameter is an array, it is again enclosed in square brackets, where individual parameters come in line with the rules and syntax of specifying multiple parameters.



2 Item types

Overview

Item types cover various methods of acquiring data from your system. Each item type comes with its own set of supported item keys and required parameters.

The following items types are currently offered by Zabbix:

- Zabbix agent checks
- SNMP agent checks
- SNMP traps
- IPMI checks
- Simple checks
- Log file monitoring
- Calculated items
- Zabbix internal checks
- SSH checks
- Telnet checks
- External checks
- Aggregate checks
- Trapper items

- [JMX monitoring](#)
- [ODBC checks](#)

Details for all item types are included in the subpages of this section. Even though item types offer a lot of options for data gathering, there are further options through [user parameters](#).

Some checks are performed by Zabbix server alone (as agent-less monitoring) while others require Zabbix agent or even Zabbix Java gateway (with JMX monitoring).

Attention:

If a particular item type requires a particular interface (like an IPMI check needs an IPMI interface on the host) that interface must exist in the host definition.

Starting with Zabbix 2.0 multiple interfaces can be set in the host definition: Zabbix agent, SNMP agent, JMX and IPMI. If an item can use more than one interface, it will search the available host interfaces (in the order: Agent→SNMP→JMX→IPMI) for the first appropriate one to be linked with.

All items that return text (character, log, text types of information) can return whitespace only as well (where applicable) setting the return value to an empty string (supported since 2.0).

1 Zabbix agent

Overview

These checks use the communication with Zabbix agent for data gathering.

There are [passive](#) and [active](#) agent checks. When configuring an item, you can select the required type:

- *Zabbix agent* - for passive checks
- *Zabbix agent (active)* - for active checks

Supported item keys

The table provides details on the item keys that you can use with Zabbix agent items.

See also:

- [Items supported by platform](#)
- [item keys specific for WIN32 agent](#)

**** Mandatory and optional parameters ****

Parameters without angle brackets are mandatory. Parameters marked with angle brackets < > are optional.

Key	Description	Return value	Parameters	Comments
▲ agent.hostname	Agent host name.	String		Returns the actual value of the agent hostname from a configuration file.
agent.ping	Agent availability check.	Nothing - unavailable 1 - available		Use function nodata() to check for host unavailability.
agent.version	Version of Zabbix agent.	String		Example of returned value: 1.8.2
kernel.maxfiles	Maximum number of opened files supported by OS.	Integer		

Key				
kernel.maxproc	Maximum number of processes supported by OS.	Integer		
log[file,<regex>,<encoding>,<maxlines>,<mode>]	Monitoring of log file.	Log	file - full path and name of log file regex - regular expression describing the required pattern encoding - code page identifier maxlines - maximum number of new lines per second the agent will send to Zabbix server or proxy. This parameter overrides the value of 'MaxLinesPerSecond' in <code>zabbix_agentd.conf</code> mode - possible values: <i>all</i> (default), <i>skip</i> - skip processing of older data (affects only newly created items that have not returned any data yet). The mode parameter is supported from version 2.0.	The item must be configured as an active check. If file is missing or permissions do not allow access, item turns unsupported. <i>Example key:</i> log[/home/zabbix/logs/logfile,, See a more detailed description.
logrt[file_regex,<regex>,<encoding>,<maxlines>,<mode>]				

	Monitoring of log file with log rotation support.	Log	file_regexp - absolute path to file and regexp describing the file name pattern regexp - regular expression describing the required content pattern encoding - code page identifier maxlines - maximum number of new lines per second the agent will send to Zabbix server or proxy. This parameter overrides the value of 'MaxLinesPerSecond' in <code>zabbix_agentd.conf</code> mode - possible values: <i>all</i> (default), <i>skip</i> - skip processing of older data (affects only newly created items that have not returned any data yet). The mode parameter is supported from version 2.0.	The item must be configured as an active check. <i>Examples:</i> <code>logrt[/home/zabbix/logs/^log9]{1,3}\$",,,100]</code> - will match a file like "logfile1" (will not match ".logfile1") <code>logrt[/home/user/^logfile_.*_9]{1,3}\$", "pattern_to_match"8",100]</code> - will collect data from files such "logfile_abc_1" or "logfile__001". Log rotation is based on the last modification time of files. See a more detailed description.

net.dns[<ip>,name,<type>,<timeout>,<count>]

net.dns.record[<ip>,name,<type>,<timeout>,<count>]	Checks if DNS service is up.	0 - DNS is down (server did not respond or DNS resolution failed) 1 - DNS is up	ip - IP address of DNS server (leave empty for the default DNS server, ignored on Windows) name - DNS name to query type - record type to be queried (default is SOA) timeout (ignored on Windows) - timeout for the request in seconds (default is 1 second) count (ignored on Windows) - number of tries for the request (default is 2)	<i>Example key:</i> net.dns[8.8.8.8,zabbix.com,M] The possible values for type are: ANY, A, NS, CNAME, MB, MG, MR, PTR, MD, MF, MX, SOA, NULL, WKS (except for Windows), HINFO, MINFO, TXT, SRV SRV record type is supported since Zabbix agent versions 1.8.6 (Unix) and 2.0.0 (Windows). Internationalized domain names are not supported, please use IDNA encoded names instead. Naming before Zabbix 2.0 (still supported): <i>net.tcp.dns</i>

Key

	Performs a DNS query.	Character string with the required type of information	ip - IP address of DNS server (leave empty for the default DNS server, ignored on Windows) name - DNS name to query type - record type to be queried (default is SOA) timeout (ignored on Windows) - timeout for the request in seconds (default is 1 second) count (ignored on Windows) - number of tries for the request (default is 2)	<i>Example key:</i> net.dns.record[8.8.8.8,zabbix. The possible values for type are: ANY, A, NS, CNAME, MB, MG, MR, PTR, MD, MF, MX, SOA, NULL, WKS (except for Windows), HINFO, MINFO, TXT, SRV SRV record type is supported since Zabbix agent versions 1.8.6 (Unix) and 2.0.0 (Windows). Internationalized domain names are not supported, please use IDNA encoded names instead. Naming before Zabbix 2.0 (still supported): net.tcp.dns.query
net.if.collisions[if]	Number of out-of-window collisions.	Integer	if - interface	
net.if.discovery	List of network interfaces. Used for low-level discovery.	JSON object		Supported since Zabbix agent version 2.0. Some Windows versions (for example, Server 2008) might require the latest updates installed to support non-ASCII characters in interface names.
net.if.in[if,<mode>]				

	Incoming traffic statistics on network interface.	Integer	if - network interface name (Unix); network interface full description or IPv4 address (Windows) mode - possible values: <i>bytes</i> - number of bytes (default) <i>packets</i> - number of packets <i>errors</i> - number of errors <i>dropped</i> - number of dropped packets	Multi-byte interface names on Windows are supported since Zabbix agent version 1.8.6. Examples: => net.if.in[eth0,errors] => net.if.in[eth0] You may obtain network interface descriptions on Windows with net.if.discovery or net.if.list items. You may use this key with a <i>Delta (speed per second)</i> store value in order to get bytes per second statistics.
net.if.out[if,<mode>]				

	Outgoing traffic statistics on network interface.	Integer	if - network interface name (Unix); network interface full description or IPv4 address (Windows) mode - possible values: <i>bytes</i> - number of bytes (default) <i>packets</i> - number of packets <i>errors</i> - number of errors <i>dropped</i> - number of dropped packets	Multi-byte interface names on Windows are supported since Zabbix agent 1.8.6 version. Examples: => net.if.out[eth0,errors] => net.if.out[eth0] You may obtain network interface descriptions on Windows with net.if.discovery or net.if.list items. You may use this key with a <i>Delta (speed per second)</i> store value in order to get bytes per second statistics.
net.if.total[if,<mode>]				

	Sum of incoming and outgoing traffic statistics on network interface.	Integer	if - network interface name (Unix); network interface full description or IPv4 address (Windows) mode - possible values: <i>bytes</i> - number of bytes (default) <i>packets</i> - number of packets <i>errors</i> - number of errors <i>dropped</i> - number of dropped packets	Examples: => net.if.total[eth0,errors] => net.if.total[eth0] You may obtain network interface descriptions on Windows with net.if.discovery or net.if.list items. You may use this key with a <i>Delta (speed per second)</i> store value in order to get bytes per second statistics. Note that dropped packets are supported only if both net.if.in and net.if.out work for dropped packets on your platform.
net.tcp.listen[port]	Checks if this TCP port is in LISTEN state.	0 - it is not in LISTEN state 1 - it is in LISTEN state	port - TCP port number	Example: net.tcp.listen[80] On Linux supported since Zabbix agent version 1.8.4
net.tcp.port[<ip>,port]				

Key

Checks if it is possible to make TCP connection to port number port.

0 - cannot connect

1 - can connect

ip - IP address (default is 127.0.0.1)
port - port number

Example:
net.tcp.port[,80]
can be used to test availability of web server running on port 80.
Old naming:
check_port[*]
For simple TCP performance testing use
net.tcp.service.perf[tcp,<ip>,<port>]
Note that these checks may result in additional messages in system daemon logfiles (SMTP and SSH sessions being logged usually).

net.tcp.service[service,<ip>,<port>]

Checks if service is running and accepting TCP connections.

0 - service is down
1 - service is running

service - either of: *ssh, ntp, ldap, smtp, ftp, http, pop, nntp, imap, tcp, https, telnet* (see [details](#))
ip - IP address (default is 127.0.0.1)
port - port number (by default standard service port number is used)

Example key:
`net.tcp.service[ftp,,45]`
- can be used to test the availability of FTP server on TCP port 45. Note that these checks may result in additional messages in system daemon logfiles (SMTP and SSH sessions being logged usually).
Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use `net.tcp.port` for checks like these.
Checking of LDAP and HTTPS by Windows agent is currently not supported.
Note that the telnet check looks for a login prompt (': ' at the end).
Old naming: `check_service[*]`
https and **telnet** services are supported since Zabbix 2.0.
ntp service only works since Zabbix 2.0.15, despite being available in earlier versions.

`net.tcp.service.perf[service,<ip>,<port>]`

	Checks performance of service.	0 - service is down seconds - the number of seconds spent while connecting to the service	service - either of: <i>ssh, ntp, ldap, smtp, ftp, http, pop, nntp, imap, tcp, https, telnet</i> (see details) ip - IP address (default is 127.0.0.1) port - port number (by default standard service port number is used)	<i>Example key:</i> net.tcp.service.perf[ssh] - can be used to test the speed of initial response from SSH server. Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use net.tcp.service.perf[tcp,<ip>], for checks like these. Checking of LDAP and HTTPS by Windows agent is currently not supported. Note that the telnet check looks for a login prompt (':' at the end). Old naming: check_service_perf[*] https and telnet services are supported since Zabbix 2.0. ntp service only works since Zabbix 2.0.15, despite being available in earlier versions.
net.udp.listen[port]	Checks if this UDP port is in LISTEN state.	0 - it is not in LISTEN state 1 - it is in LISTEN state	port - UDP port number	Example: net.udp.listen[68] On Linux supported since Zabbix agent version 1.8.4
proc.mem[<name>,<user>,<mode>,<cmdline>]				

Key

Memory used
by process in
bytes.

Integer - with
mode as *max*,
min, *sum*

Float - with
mode as *avg*

name -
process name
(default is *all*
processes)

user - user
name (default
is *all users*)

mode -
possible
values:
avg, *max*, *min*,
sum (default)

cmdline - filter
by command
line

Example keys:
proc.mem[,root]
- memory used
by all
processes
running under
the "root" user
proc.mem[zabbix_server,zabb
- memory used
by all
zabbix_server
processes
running under
the zabbix user
proc.mem[,oracle,max,oracle2
- memory used
by the most
memory-
hungry process
running under
oracle having
oracleZABBIX
in its command
line

Note: When
several
processes use
shared
memory, the
sum of
memory used
by processes
may result in
large,
unrealistic
values.

`proc.num[<name>,<user>,<state>,<cmdline>]`

	The number of processes.	Integer	name - process name (default is <i>all processes</i>) user - user name (default is <i>all users</i>) state - possible values: <i>all</i> (default), <i>run</i> , <i>sleep</i> , <i>zomb</i> cmdline - filter by command line	<i>Example keys:</i> proc.num[,mysql] - number of processes running under the mysql user proc.num[apache2,www-data] - number of apache2 processes running under the www-data user proc.num[,oracle,sleep,oracle] - number of processes in sleep state running under oracle having oracleZABBIX in its command line On Windows, only the <i>name</i> and <i>user</i> parameters are supported.
sensor[device,sensor,<mode>]	Hardware sensor reading.	Float	device - device name (if <mode> is used, it is a regular expression) sensor - sensor name (if <mode> is used, it is a regular expression) mode - possible values: <i>avg</i> , <i>max</i> , <i>min</i> (if this parameter is omitted, device and sensor are treated verbatim).	On Linux 2.4, reads <i>/proc/sys/dev/sensors</i> . Example key: sensor[w83781d-i2c-0-2d,temp1] Prior to Zabbix 1.8.4, the sensor[temp1] format was used. On OpenBSD, reads the <i>hw.sensors</i> MIB. Example keys: sensor[cpu0,temp0] - temperature of one CPU sensor["cpu[0-2]\$",temp,avg] - average temperature of the first three CPU's Supported on OpenBSD since Zabbix 1.8.4.
system.boottime	System boot time.	Integer (Unix timestamp)		

Key				
system.cpu.intr	Device interrupts.	Integer		
system.cpu.load[<cpu>,<mode>]	CPU load.	Float	cpu - possible values: <i>all</i> (default), <i>percpu</i> (total load divided by online CPU count) mode - possible values: <i>avg1</i> (one-minute average, default), <i>avg5</i> (5-minute average), <i>avg15</i> (an average within 15 minutes)	<i>Example key:</i> system.cpu.load[,avg5] Old naming: sys-tem.cpu.loadX Parameter percpu is supported since Zabbix 2.0.0.
system.cpu.num[<type>]	Number of CPUs.	Integer	type - possible values: <i>online</i> (default), <i>max</i>	<i>Example key:</i> system.cpu.num
system.cpu.switches	Count of context switches.	Integer		Old naming: sys-tem[switches]
system.cpu.util[<cpu>,<type>,<mode>]	CPU utilisation in percent.	Float	cpu - CPU number (default is all CPUs) type - possible values: <i>idle</i> , <i>nice</i> , <i>user</i> (default), <i>system</i> (default for Windows), <i>iowait</i> , <i>interrupt</i> , <i>softirq</i> , <i>steal</i> mode - possible values: <i>avg1</i> (one-minute average, default), <i>avg5</i> (5-minute average), <i>avg15</i> (an average within 15 minutes)	<i>Example key:</i> system.cpu.util[0,user,avg5] Old naming: sys-tem.cpu.idleX, sys-tem.cpu.niceX, sys-tem.cpu.systemX, sys-tem.cpu.userX
system.hostname[<type>]				

Key	System host name.	String	type (only on Windows, ignored on other systems) - possible values: <i>netbios</i> (default) or <i>host</i>	The value is acquired by either GetComputerName() (for netbios) or gethostname() (for host) functions on Windows and by "hostname" command on other systems. The type parameter for this item is supported since 1.8.6 version. Examples of returned values: on Linux: <i>system.hostname</i> → linux-w7x1 <i>system.hostname</i> → www.zabbix.com on Windows: <i>system.hostname</i> → WIN-SERV2008-I6 <i>system.hostname[host]</i> → Win-Serv2008-I6LonG See also a more detailed description.
system.hw.chassis[<info>]				

	Chassis information.	String	info - one of full (default), model, serial, type or vendor	Example: system.hw.chassis[full] Hewlett-Packard HP Pro 3010 Small Form Factor PC CZXXXXXXXX Desktop] This key depends on the availability of the SMBIOS table in memory. Root permissions are required because the value is acquired by reading from memory. Supported since Zabbix agent version 2.0.
system.hw.cpu[<cpu>,<info>]	CPU information.	String or integer	cpu - CPU number or all (default) info - one of full (default), curfreq, maxfreq, model or vendor	Example: system.hw.cpu[0,vendor] AuthenticAMD Gathers info from /proc/cpuinfo and /sys/devices/system/cpu/[cpu*] If a CPU number and curfreq or maxfreq is specified, a numeric value is returned (Hz). Supported since Zabbix agent version 2.0.
system.hw.devices[<type>]				

	Listing of PCI or USB devices.	Text	type - pci (default) or usb	Example: system.hw.devices[pci] 00:00.0 Host bridge: Advanced Micro Devices [AMD] RS780 Host Bridge [...] Returns the output of either lspci or lsusb utility (executed without any parameters) Supported since Zabbix agent version 2.0.
system.hw.macaddr[<interface>,<format>]	Listing of MAC addresses.	String	interface - all (default) or a regular expression format - full (default) or short	Example: system.hw.macaddr["eth0\$",full] [eth0] 00:11:22:33:44:55 Lists MAC addresses of the interfaces whose names match the given interface regexp (<i>all</i> lists for all interfaces). If format is specified as short, interface names and identical MAC addresses are not listed. Supported since Zabbix agent version 2.0.
system.localtime[<type>]				

	System time.	Integer - with type as <i>utc</i>	utc - (default) the time since the Epoch (00:00:00 UTC, January 1, 1970), measured in seconds.	Example: <code>system.localtime[local]</code> - create an item using this key and then use it to display host time in the <i>Clock</i> screen element .
		String - with type as <i>local</i>	local - the time in the 'yyyy-mm-dd,hh:mm:ss.nnn, +hh:mm' format Parameters for this item supported from version 2.0.	
<code>system.run[command,<mode>]</code>	Run specified command on the host.			
		1 - with mode as <i>nowait</i> (regardless of command result)	command - command for execution mode - one of wait (default, wait end of execution), nowait (do not wait)	Up to 512KB of data can be returned (64KB before Zabbix 2.0.5), including trailing whitespace that is truncated. To be processed correctly, the output of the command must be text. Example: <code>system.run[ls -l /]</code> - detailed file list of root directory. <i>Note:</i> To enable this functionality, agent configuration file must have <code>EnableRemoteCommands=1</code> option. See also: Command execution .
<code>system.stat[resource,<type>]</code>				

System
statistics

Integer or float

ent - number
of processor
units this
partition is
entitled to
receive (float)**kthr,<type>** -
information
about kernel
thread states:**r** - average
number of
runnable
kernel threads
(float)**b** - average
number of
kernel threads
placed in the
Virtual Memory
Manager wait
queue (float)**memory,<type>**
- information
about the
usage of virtual
and real
memory:**avm** - active
virtual pages
(integer)**fre** - size of the
free list
(integer)**page,<type>**
- information
about page
faults and
paging activity:**fi** - file
page-ins per
second (float)**fo** - file
page-outs per
second (float)**pi** - pages
paged in from
paging space
(float)**po** - pages
paged out to
paging space
(float)**fr** - pages
freed (page
replacement)
(float)**sr** - pages
scanned by
page-
replacement
algorithm
(float)**faults,<type>**
- trap and

Key				
system.sw.arch	Software architecture information.	String		Example: system.sw.arch i686 Info is acquired from uname() function. Supported since Zabbix agent version 2.0.
system.sw.os[<info>]	Operating system information.	String	info - one of full (default), short or name	Example: system.sw.os[short] Ubuntu 2.6.35-28.50-generic 2.6.35.11 Info is acquired from (note that not all files are present in all distributions): [full] - /proc/version [short] - /proc/version_signature [name] - /etc/issue.net Supported since Zabbix agent version 2.0.
system.sw.packages[<package>,<manager>,<format>]				

	Listing of installed packages.	Text	package - all (default) or a regular expression manager - all (default) or a package manager format - full (default) or short	Example: system.sw.packages[mini,dpkg,sl python-minimal, python2.6-minimal, ubuntu-minimal Lists (alphabetically) installed packages whose names match the given package regexp (<i>all</i> lists them all). Supported packages managers: manager (executed command) dpkg (dpkg --get-selections) pkgtool (ls /var/log/packages) rpm (rpm -qa) pacman (pacman -Q) If format is specified as full, packages are grouped by package managers (each manager on a separate line beginning with its name in square brackets). If format is specified as short, packages are not grouped and are listed on a single line. Supported since Zabbix agent version 2.0.
system.swap.in[<device>,<type>]				

	Swap in (from device into memory) statistics.	Integer	device - device used for swapping (default is all) type - possible values: <i>count</i> (number of swapins), <i>sectors</i> (sectors swapped in), <i>pages</i> (pages swapped in). See supported by platform for details on defaults.	<i>Example key:</i> system.swap.in[,pages] <i>The source of this information is:</i> Linux 2.4: /proc/swaps, /proc/partitions, /proc/stat Linux 2.6: /proc/swaps, /proc/diskstats, /proc/vmstat
system.swap.out[<device>,<type>]	Swap out (from memory onto device) statistics.	Integer	device - device used for swapping (default is all) type - possible values: <i>count</i> (number of swapouts), <i>sectors</i> (sectors swapped out), <i>pages</i> (pages swapped out). See supported by platform for details on defaults.	<i>Example key:</i> system.swap.out[,pages] <i>The source of this information is:</i> Linux 2.4: /proc/swaps, /proc/partitions, /proc/stat Linux 2.6: /proc/swaps, /proc/diskstats, /proc/vmstat
system.swap.size[<device>,<type>]				

Key

	Swap space size in bytes or in percentage from total.	Integer - for bytes Float - for percentage ¹	device - device used for swapping (default is all) type - possible values: <i>free</i> (free swap space, default), <i>pfree</i> (free swap space, in percent), <i>pusd</i> (used swap space, in percent), <i>total</i> (total swap space), <i>used</i> (used swap space)	<i>Example key:</i> system.swap.size[,pfree] - free swap space percentage If <i>device</i> is not specified Zabbix agent will only take into account swap devices (files), physical memory will be ignored. For example, on Solaris systems <i>swap -s</i> command includes a portion of physical memory and swap devices (unlike <i>swap -l</i>). Old naming: sys-tem.swap.free, sys-tem.swap.total
system.uname	Detailed host information.	String		Example of returned value: FreeBSD localhost 4.4-RELEASE FreeBSD 4.4-RELEASE #0: Tue Sep 18 11:57:08 PDT 2001 mur-ray@builder.FreeBSD.org: /usr/src/sys/compile/GENERIC i386
system.uptime	System uptime in seconds.	Integer		In item configuration , use s or uptime units to get readable values.
system.users.num	Number of users logged in.	Integer		who command is used on the agent side to obtain the value.
vfs.dev.read[<device>,<type>,<mode>]				

	Disk read statistics.	Integer - with type in <i>sectors, operations, bytes</i> Float - with type in <i>sps, ops, bps</i>	device - disk device (default is <i>all</i> ²) type - possible values: <i>sectors, operations, bytes, sps, ops, bps</i> (must be specified, since defaults differ under various OSes). <i>sps, ops, bps</i> stand for: sectors, operations, bytes per second, respectively mode - possible values: <i>avg1</i> (one-minute average, default), <i>avg5</i> (five-minute average), <i>avg15</i> (15-minute average). <i>Note:</i> The third parameter is supported only if the type is in: <i>sps, ops, bps</i> .	Default values of 'type' parameter for different OSes: FreeBSD - <i>bps</i> Linux - <i>sps</i> OpenBSD - <i>operations</i> Solaris - <i>bytes</i> <i>Example key:</i> <i>vfs.dev.read[,operations]</i> <i>Old naming:</i> <i>io[*]</i> Usage of the type parameters <i>ops, bps</i> and <i>sps</i> on supported platforms is limited to 8 devices (7 individual devices and one <i>all</i>). Starting with Zabbix 2.0.1 this limit is increased to 1024 (1023 individual devices and one for <i>all</i>). Supports LVM since Zabbix 1.8.6. Until Zabbix 1.8.6, only relative device names may be used (for example, sda), since 1.8.6 an optional /dev/ prefix may be used (for example, /dev/sda)
vfs.dev.write[<device>,<type>,<mode>]				

	Disk write statistics.	Integer - with type in <i>sectors, operations, bytes</i> Float - with type in <i>sps, ops, bps</i>	device - disk device (default is <i>all</i> ²) type - one of sectors, operations, bytes, sps, ops, bps (must specify exactly which parameter to use, since defaults are different under various OSes). sps, ops, bps means: sectors, operations, bytes per second respectively mode - one of avg1 (default), avg5 (average within 5 minutes), avg15. <i>Note:</i> The third parameter is supported only if the type is in: sps, ops, bps.	Default values of 'type' parameter for different OSes: FreeBSD - bps Linux - sps OpenBSD - operations Solaris - bytes Example: vfs.dev.write[,operations] Old naming: io[*] The type parameters ops, bps and sps on supported platforms are limited to 8 devices (7 individual devices and one <i>all</i>). Starting with Zabbix 2.0.1 this limit is increased to 1024 (1023 individual devices and one for <i>all</i>). Supports LVM since Zabbix 1.8.6. Until Zabbix 1.8.6, only relative device names may be used (for example, sda), since 1.8.6 optional /dev/ prefix may be used (for example, /dev/sda)
vfs.file.cksum[file]				

Key				
	File checksum, calculated by the UNIX cksum algorithm.	Integer	file - full path to file	Example of returned value: 1938292000 Example: vfs.file.cksum[/etc/passwd] Old naming: cksum The file size limit depends on large file support.
vfs.file.contents[file,<encoding>]	Retrieving contents of a file.	Text	file - full path to file encoding - code page identifier	Returns EOF if the file is empty or contains LF/CR characters only. Example: vfs.file.contents[/etc/passwd] This item is limited to files no larger than 64 Kbytes. Supported since Zabbix agent version 2.0.
vfs.file.exists[file]	Checks if file exists.	0 - not found 1 - regular file or a link (symbolic or hard) to regular file exists	file - full path to file	Example: vfs.file.exists[/tmp/application] The return value depends on what S_ISREG POSIX macro returns. The file size limit depends on large file support.
vfs.file.md5sum[file]				

	MD5 checksum of file.	Character string (MD5 hash of the file)	file - full path to file	Example of returned value: b5052decb577e0fffd622d6dd
				Example: vfs.file.md5sum[/usr/local/etc/
				The file size limit (64 MB) for this item was removed in version 1.8.6.
				The file size limit depends on large file support .
vfs.file.regexp[file,regexp,<encoding>]	Find string in a file.	The whole line from file containing the matched string or EOF if expression not found	file - full path to file regexp - GNU regular expression encoding - code page identifier	Only the first matching line is returned. Example: vfs.file.regexp[/etc/passwd,zal
vfs.file.regmatch[file,regexp,<encoding>]	Find string in a file.	0 - expression not found 1 - found	file - full path to file regexp - GNU regular expression encoding - code page identifier	Example: vfs.file.regmatch[/var/log/app.
vfs.file.size[file]	File size (in bytes).	Integer	file - full path to file	File must have read permissions for user zabbix Example: vfs.file.size[/var/log/syslog] The file size limit depends on large file support .
vfs.file.time[file,<mode>]	File time information.	Integer (Unix timestamp)	file - full path to the file mode - one of modify (default, modification time), access - last access time, change - last change time	Example: vfs.file.time[/etc/passwd,modi The file size limit depends on large file support .

Key

vfs.fs.discovery	List of mounted filesystems. Used for low-level discovery.	JSON object		Supported since Zabbix agent version 2.0.
vfs.fs.inode[fs,<mode>]	Number or percentage of inodes.	Integer - for number Float - for percentage	fs - filesystem mode - one of total (default), free, used, pfree (free, percentage), pused (used, percentage)	Example: vfs.fs.inode[,pfree] Old naming: vfs.fs.inode.free[*], vfs.fs.inode.pfree[*], vfs.fs.inode.total[*]
vfs.fs.size[fs,<mode>]	Disk space in bytes or in percentage from total.	Integer - for bytes Float - for percentage	fs - filesystem mode - one of total (default), free, used, pfree (free, percentage), pused (used, percentage)	In case of a mounted volume, disk space for local file system is returned. Example: vfs.fs.size[/tmp,free] Reserved space of a file system is taken into account and not included when using the <i>free</i> mode. Old naming: vfs.fs.free[*], vfs.fs.total[*], vfs.fs.used[*], vfs.fs.pfree[*], vfs.fs.pused[*]
vm.memory.size[<mode>]				

	Memory size in bytes or in percentage from total.	Integer - for bytes Float - for percentage	mode - one of total (default), active, anon, buffers, cached, exec, file, free, inactive, pinned, shared, wired, used, pused, available, pavailable	Old naming: vm.memory.buffers, vm.memory.cached, vm.memory.free, vm.memory.shared, vm.memory.total Item vm.memory.size[] accepts three categories of parameters. First category consists of total - total amount of memory. Second category contains platform-specific memory types: active, anon, buffers, cached, exec, file, free, inactive, pinned, shared, wired. Third category are user-level estimates on how much memory is used and available: used, pused, available, pavailable. See a more detailed description of vm.memory.size parameters.
web.page.get[host,<path>,<port>]	Get content of web page.	Web page source as text	host - hostname path - path to HTML document (default is /) port - port number (default is 80)	Returns EOF on fail. Example: web.page.get[www.zabbix.com]
web.page.perf[host,<path>,<port>]				

Key				
	Loading time of full web page (in seconds).	Float	host - hostname path - path to HTML document (default is /) port - port number (default is 80)	Returns 0 on fail. Example: web.page.perf[www.zabbix.co
web.page.regex[host,<path>,<port>,<regexp>,<length>]	Get first occurrence of regexp in web page.	Matched string	host - hostname path - path to HTML document (default is /) port - port number (default is 80) regexp - GNU regular expression length - maximum number of characters to return	Returns EOF on fail (no match). Example: web.page.regex[www.zabbix

Note:

[1] The system.swap.size key might report incorrect data on virtualized (VMware ESXi, VirtualBox) Windows platforms. In this case use perf_counter[700(_Total)\702] key to obtain correct swap usage percentage.

Note:

[2] If default *all* is used for the first parameter of **vfs.dev.*** keys then the keys will return summary statistics, including: all block devices like sda, sdb and their partitions sda1, sda2, sdb3 ... and multiple devices (MD raid) based on those block devices/partitions and logical volumes (LVM) based on those block devices/partitions.
In such cases returned values should be considered only as relative value (dynamic in time) but not as absolute values.

Note:

A Linux-specific note. Zabbix agent must have read-only access to filesystem */proc*. Kernel patches from www.grsecurity.org limit access rights of non-privileged users.

Available encodings

The `encoding` parameter is used to specify encoding for processing corresponding item checks, so that data acquired will not be corrupted. For a list of supported encodings (code page identifiers), please consult respective documentation, such as documentation for [libiconv](http://libiconv.org) (GNU Project) or Microsoft Windows SDK documentation for "Code Page Identifiers".

If empty `encoding` is passed, then UTF-8 (default locale for newer Unix/Linux distributions, see your system's settings) or ANSI with system-specific extension (Windows) is used by default.

Windows-specific item keys

Item keys

The table provides details on the item keys that you can use with Zabbix Windows agent only.

Key

▲	Description	Return value	Parameters	Comments
	eventlog[name,<regexp>,<severity>,<source>,<eventid>,<maxlines>,<mode>]			

Event log monitoring.	Log	name - name of event log regex - regular expression describing the required pattern severity - regular expression describing severity The parameter accepts the following values: "Information", "Warning", "Error", "Failure Audit", "Success Audit" source - source identifier eventid - regular expression describing the event identifier(s) maxlines - maximum number of new lines per second the agent will send to Zabbix server or proxy. This parameter overrides the value of 'MaxLinesPerSecond' in <code>zabbix_agentd.win.conf</code> mode - possible values: <i>all</i> (default), <i>skip</i> - skip processing of older data (affects only newly created items that have not returned any data yet). The mode parameter is supported from version 2.0.	The item must be configured as an active check. Examples: <pre>eventlog[Application] eventlog[Security,"Failure Audit","^(529 680)\$] eventlog[System,"Warning Error","^(529 680)\$"] eventlog[System,,,,^1\$] eventlog[System,,,,@TWOSHORT]</pre> - here a custom regular expression named TWOSHORT is referenced (defined as a Result is TRUE type, the expression itself being ^1\$ ^70\$). See also additional information on log monitoring.

Key

net.if.list

Network
interface list
(includes
interface type,
status, IPv4
address,
description).

Text

Supported
since Zabbix
agent version
1.8.1.
Multi-byte
interface
names
supported
since Zabbix
agent version
1.8.6. Disabled
interfaces are
not listed.

Note that en-
abling/disabling
some
components
may change
their ordering
in the Windows
interface
name.

Some Windows
versions (for
example,
Server 2008)
might require
the latest
updates
installed to
support
non-ASCII
characters in
interface
names.

perf_counter[counter,<interval>]

	Value of any Windows performance counter.	Integer, float, string or text (depending on the request)	counter - path to the counter interval - last N seconds for storing the average value. The <code>interval</code> must be between 1 and 900 seconds (included) and the default value is 1.	Performance Monitor can be used to obtain list of available counters. Until version 1.6 this parameter will return correct value only for counters that require just one sample (like \System\Threads). It will not work as expected for counters that require more than one sample - like CPU utilisation. Since 1.6 interval is used, so the check returns an average value for last "interval" seconds every time. See also: Windows performance counters .
<code>proc_info[process,<attribute>,<type>]</code>				

Various
information
about specific
process(es).

Float

process -
process name
attribute -
requested
process
attribute.
type -
representation
type
(meaningful
when more
than one
process with
the same
name exists)

The following
attributes are
currently
supported:
vmsize
(default) - Size
of process
virtual memory
in Kbytes
wkset - Size of
process
working set
(amount of
physical
memory used
by process) in
Kbytes
pf - Number of
page faults
ktime -
Process kernel
time in
milliseconds
utime -
Process user
time in
milliseconds
io_read_b -
Number of
bytes read by
process during
I/O operations
io_read_op -
Number of
read operation
performed by
process
io_write_b -
Number of
bytes written
by process
during I/O
operations
io_write_op -
Number of
write operation
performed by
process
io_other_b -
Number of
bytes
transferred by
process during
operations
other than read
and write
operations
io_other_op -
Number of I/O
operations
performed by
process, other
than read and
write
operations

Key

service_state[service]

State of a service.

0 - running
1 - paused
2 - start pending
3 - pause pending
4 - continue pending
5 - stop pending
6 - stopped
7 - unknown
255 - no such service

service - a real service name or its display name as seen in MMC Services snap-in

services[<type>,<state>,<exclude>]

Listing of services.

0 - if empty

Text - list of services separated by a newline

type - one of all (default), automatic, manual, disabled
state - one of all (default), stopped, started, start_pending, stop_pending, running, continue_pending, pause_pending, paused
exclude - list of services to exclude it from the result. Excluded services should be written in double quotes, separated by comma, without spaces. This parameter is supported starting with Zabbix **1.8.1**.

Examples:
services[,started] - list of started services
services[automatic,stopped] - list of stopped services, that should be run
services[automatic,stopped,"service1,service2,service3"] - list of stopped services, that should be run, excluding services with names service1, service2 and service3

Monitoring Windows services

This tutorial provides step-by-step instructions for setting up the monitoring of Windows services. It is assumed that Zabbix server and agent are configured and operational.

To monitor the up/down status of a service you need to perform the following steps:

Step 1

Get the service name.

You can get that name by going to the services mmc and bringing up the properties of the service. In the General tab you should see a field called 'Service name'. The value that follows is the name you will use when setting up an item for monitoring.

For example, if you wanted to monitor the "workstation" service then your service might be: **lanmanworkstation**.

Step 2

Configure an item for monitoring the service, with:

- **Key:** service_state[lanmanworkstation]
- **Type of information:** Numeric (unsigned)
- **Show value:** select the *Windows service state* value mapping

2 SNMP agent

Overview

You may want to use SNMP monitoring on devices such as printers, network switches, routers or UPS that usually are SNMP-enabled and on which it would be impractical to attempt setting up complete operating systems and Zabbix agents.

To be able to retrieve data provided by SNMP agents on these devices, Zabbix server must be **initially configured** with SNMP support.

SNMP checks are performed over the UDP protocol only.

Warning:

If monitoring SNMPv3 devices, make sure that msgAuthoritativeEngineID (also known as snmpEngineID or "Engine ID") is never shared by two devices. According to [RFC 2571](#) (section 3.1.1.1) it must be unique for each device.

Note:

For SNMPv3 privacy and authentication currently MD5 and DES protocols are supported.

Configuring SNMP monitoring

To start monitoring a device through SNMP, the following steps have to be performed:

Step 1

Create a **host** for the device with an SNMP interface.

Enter the IP address. Set the host status to NOT MONITORED. You can use one of the provided SNMP templates (*Template SNMP Device* and others) that will automatically add a set of items. However, the template may not be compatible with the host.

Note:

SNMP checks do not use *Agent port*, it is ignored.

Step 2

Find out the SNMP string (or OID) of the item you want to monitor.

To get a list of SNMP strings, use the **snmpwalk** command (part of [ucd-snmp/net-snmp](#) software which you should have installed as part of the Zabbix installation) or equivalent tool:

```
shell> snmpwalk -v 2c -c public <host IP> .
```

As '2c' here stands for SNMP version, you may also substitute it with '1', to indicate SNMP Version 1 on the device.

This should give you a list of SNMP strings and their last value. If it doesn't then it is possible that the SNMP 'community' is different from the standard 'public' in which case you will need to find out what it is.

You can then go through the list until you find the string you want to monitor, e.g. if you wanted to monitor the bytes coming in to your switch on port 3 you would use the IF-MIB::ifInOctets.3 string from this line:

```
IF-MIB::ifInOctets.3 = Counter32: 3409739121
```

You may now use the **snmpget** command to find out the numeric OID for 'IF-MIB::ifInOctets.3':

```
shell> snmpget -v 2c -c public -On 10.62.1.22 IF-MIB::ifInOctets.3
```

Note that the last number in the string is the port number you are looking to monitor. See also: [Dynamic indexes](#).

This should give you something like the following:

```
.1.3.6.1.2.1.2.2.1.10.3 = Counter32: 3472126941
```

Again, the last number in the OID is the port number.

Note:

3COM seem to use port numbers in the hundreds, e.g. port 1 = port 101, port 3 = port 103, but Cisco use regular numbers, e.g. port 3 = 3.

Note:

Some of the most used SNMP OIDs are **translated automatically to a numeric representation** by Zabbix.

Step 3

Create an item for monitoring.

So, now go back to Zabbix and click on Items, selecting the SNMP host you created earlier. Depending on whether you used a template or not when creating your host, you will have either a list of SNMP items associated with your host or just a new item box. We will work on the assumption that you are going to create the item yourself using the information you have just gathered using snmpwalk and snmpget, so enter a plain English description in the 'Description' field of the new item box. Make sure the 'Host' field has your switch/router in it and change the 'Type' field to "SNMPv* agent". Enter the community (usually public) and enter the textual or numeric OID that you retrieved earlier into the 'SNMP OID' field, for example: .1.3.6.1.2.1.2.2.1.10.3

Enter the 'SNMP port' as 161 and the 'Key' as something meaningful, e.g. SNMP-InOctets-Bps. Choose a Multiplier if you want one and enter an 'update interval' and 'keep history' if you want it to be different from the default. Set the 'Status' to Monitored, the 'Type of information' to *Numeric (float)* and the 'Store value' to *Delta (speed per second)* (important otherwise you will get cumulative values from the SNMP device instead of the latest change).

Now save the item and go back to the hosts area of Zabbix. From here change the SNMP device status to 'Monitored' and check in *Latest data* for your SNMP data!

Example 1

General example:

Parameter	Description
Community	public
OID	1.2.3.45.6.7.8.0 (or .1.2.3.45.6.7.8.0)
Key	<Unique string to be used as reference to triggers> For example, "my_param".

Note that OID can be given in either numeric or string form. However, in some cases, string OID must be converted to numeric representation. Utility snmpget may be used for this purpose:

```
shell> snmpget -On localhost public enterprises.ucdavis.memory.memTotalSwap.0
```

Monitoring of SNMP parameters is possible if either --with-net-snmp or --with-ucd-snmp flag was specified while configuring Zabbix sources.

Example 2

Monitoring of uptime:

Parameter	Description
Community	public
Oid	MIB::sysUpTime.0
Key	router.uptime
Value type	Float
Units	uptime
Multiplier	0.01

1 Dynamic indexes**Overview**

While you may find the required index number (for example, of a network interface) among the SNMP OIDs, sometimes you may not completely rely on the index number always staying the same.

Index numbers may be dynamic - they may change over time and your item may stop working as a consequence.

To avoid this scenario, it is possible to define an OID which takes into account the possibility of an index number changing.

For example, if you need to retrieve the index value to append to **ifInOctets** that corresponds to the **GigabitEthernet0/1** interface on a Cisco device, use the following OID:

```
ifInOctets["index","ifDescr","GigabitEthernet0/1"]
```

The syntax

A special syntax for OID is used:

<OID of data>["index", "<base OID of index>", "<string to search for>"]

Parameter	Description
OID of data	Main OID to use for data retrieval on the item.
index	Method of processing. Currently one method is supported: index - search for index and append it to the data OID
base OID of index	This OID will be looked up to get the index value corresponding to the string.
string to search for	The string to use for an exact match with a value when doing lookup. Case sensitive.

Example

Getting memory usage of apache process.

If using this OID syntax:

```
HOST-RESOURCES-MIB::hrSWRunPerfMem["index", "HOST-RESOURCES-MIB::hrSWRunPath", "/usr/sbin/apache2"]
```

the index number will be looked up here:

```
...
HOST-RESOURCES-MIB::hrSWRunPath.5376 = STRING: "/sbin/getty"
HOST-RESOURCES-MIB::hrSWRunPath.5377 = STRING: "/sbin/getty"
HOST-RESOURCES-MIB::hrSWRunPath.5388 = STRING: "/usr/sbin/apache2"
HOST-RESOURCES-MIB::hrSWRunPath.5389 = STRING: "/sbin/sshd"
...
```

Now we have the index, 5388. The index will be appended to the data OID in order to receive the value we are interested in:

```
HOST-RESOURCES-MIB::hrSWRunPerfMem.5388 = INTEGER: 31468 KBytes
```

Note:

Dynamic indexes are cached since Zabbix version 1.6.3.

Note:

Using dynamic indexes leads to more SNMP queries in Zabbix versions up to 1.7. Dynamic index lookup and data retrieval is performed in single connection since Zabbix version 1.7.

2 Special OIDs

Some of the most used SNMP OIDs are translated automatically to a numeric representation by Zabbix. For example, **ifIndex** is translated to **1.3.6.1.2.1.2.2.1.1**, **ifIndex.0** is translated to **1.3.6.1.2.1.2.2.1.1.0**.

The table contains list of the special OIDs.

Special OID	Identifier	Description
ifIndex	1.3.6.1.2.1.2.2.1.1	A unique value for each interface.

Special OID	Identifier	Description
ifDescr	1.3.6.1.2.1.2.2.1.2	A textual string containing information about the interface. This string should include the name of the manufacturer, the product name and the version of the hardware interface.
ifType	1.3.6.1.2.1.2.2.1.3	The type of interface, distinguished according to the physical/link protocol(s) immediately 'below' the network layer in the protocol stack.
ifMtu	1.3.6.1.2.1.2.2.1.4	The size of the largest datagram which can be sent / received on the interface, specified in octets.
ifSpeed	1.3.6.1.2.1.2.2.1.5	An estimate of the interface's current bandwidth in bits per second.
ifPhysAddress	1.3.6.1.2.1.2.2.1.6	The interface's address at the protocol layer immediately 'below' the network layer in the protocol stack.
ifAdminStatus	1.3.6.1.2.1.2.2.1.7	The current administrative state of the interface.
ifOperStatus	1.3.6.1.2.1.2.2.1.8	The current operational state of the interface.
ifInOctets	1.3.6.1.2.1.2.2.1.10	The total number of octets received on the interface, including framing characters.
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
ifInNUcastPkts	1.3.6.1.2.1.2.2.1.12	The number of non-unicast (i.e., subnetwork- broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol.
ifInDiscards	1.3.6.1.2.1.2.2.1.13	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
ifInErrors	1.3.6.1.2.1.2.2.1.14	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
ifInUnknownProtos	1.3.6.1.2.1.2.2.1.15	The number of packets received via the interface which were discarded because of an unknown or unsupported protocol.
ifOutOctets	1.3.6.1.2.1.2.2.1.16	The total number of octets transmitted out of the interface, including framing characters.
ifOutUcastPkts	1.3.6.1.2.1.2.2.1.17	The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.
ifOutNUcastPkts	1.3.6.1.2.1.2.2.1.18	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.

Special OID	Identifier	Description
ifOutDiscards	1.3.6.1.2.1.2.2.1.19	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
ifOutErrors	1.3.6.1.2.1.2.2.1.20	The number of outbound packets that could not be transmitted because of errors.
ifOutQLen	1.3.6.1.2.1.2.2.1.21	The length of the output packet queue (in packets).

3 SNMP traps

Overview

Receiving SNMP traps is the opposite to querying SNMP-enabled devices.

In this case the information is sent from a SNMP-enabled device and is collected or "trapped" by Zabbix.

Usually traps are sent upon some condition change and the agent connects to the server on port 162 (as opposed to port 161 on the agent side that is used for queries). Using traps may detect some short problems that occur amidst the query interval and may be missed by the query data.

Receiving SNMP traps in Zabbix is designed to work with **snmptrapd** and one of the built-in mechanisms for passing the traps to Zabbix - either a perl script or SNMPTT.

The workflow of receiving a trap:

1. **snmptrapd** receives a trap
2. snmptrapd passes the trap to SNMPTT or calls Perl trap receiver
3. SNMPTT or Perl trap receiver parses, formats and writes the trap to a file
4. Zabbix SNMP trapper reads and parses the trap file
5. For each trap Zabbix finds all "SNMP trapper" items with host interfaces matching the received trap address. Note that only the selected "IP" or "DNS" in host interface is used during the matching.
6. For each found item, the trap is compared to regexp in "snmptrap[regexp]". The trap is set as the value of **all** matched items. If no matching item is found and there is an "snmptrap.fallback" item, the trap is set as the value of that.
7. If the trap was not set as the value of any item, Zabbix by default logs the unmatched trap. (This is configured by "Log unmatched SNMP traps" in Administration → General → Other.)

3.1 Configuring SNMP traps

Configuring the following fields in the frontend is specific for this item type:

- Your host must have an SNMP interface

In *Configuration → Hosts*, in the **Host interface** field set an SNMP interface with the correct IP or DNS address. The address from each received trap is compared to the IP and DNS addresses of all SNMP interfaces to find the corresponding hosts.

- Configure the item

In the **Key** field use one of the SNMP trap keys:

Key		
Description	Return value	Comments
snmptrap[regexp]		

Key		
Catches all SNMP traps from a corresponding address that match the regular expression specified in regexp	SNMP trap	This item can be set only for SNMP interfaces. This item is supported starting from version 2.0.0. <i>Note:</i> Starting with Zabbix 2.0.5, user macros and global regular expressions are supported in the parameter of this item key.
snmptrap.fallback Catches all SNMP traps from a corresponding address that were not caught by any of the snmptrap[] items for that interface	SNMP trap	This item can be set only for SNMP interfaces. This item is supported starting from version 2.0.0.

Note:

Multi-line regexp matching is not supported at this time.

Set the **Type of information** to be 'Log' for the timestamps to be parsed. Note that other formats such as 'Numeric' are also acceptable but might require a custom trap handler.

Note:

For SNMP trap monitoring to work, it must first be correctly set up.

3.2 Setting up SNMP trap monitoring

Configuring Zabbix server/proxy

To read the traps, Zabbix server or proxy must be configured to start the SNMP trapper process and point to the trap file that is being written by SNMPTT or a perl trap receiver. To do that, edit the configuration file (**zabbix_server.conf** or **zabbix_proxy.conf**):

1. StartSNMPTrapper=1
2. SNMPTTrapperFile=[TRAP FILE]

Warning:

If systemd parameter **PrivateTmp** is used, this file is unlikely to work in */tmp*.

Configuring SNMPTT

At first, snmptrapd should be configured to use SNMPTT.

Note:

For the best performance, SNMPTT should be configured as a daemon using **snmptthandler-embedded** to pass the traps to it. See instructions for configuring SNMPTT in its homepage:

<http://snmptt.sourceforge.net/docs/snmptt.shtml>

When SNMPTT is configured to receive the traps, configure SNMPTT to log the traps:

1. log traps to the trap file which will be read by Zabbix:
log_enable = 1
log_file = [TRAP FILE]
2. set the date-time format:
date_time_format = %H:%M:%S %Y/%m/%d = [DATE TIME FORMAT]

Now format the traps for Zabbix to recognise them (edit snmptt.conf):

1. Each FORMAT statement should start with "ZBXTRAP [address]", where [address] will be compared to IP and DNS addresses of SNMP interfaces on Zabbix. E.g.:
EVENT coldStart .1.3.6.1.6.3.1.1.5.1 "Status Events" Normal
FORMAT ZBXTRAP \$aA Device reinitialized (coldStart)
2. See more about SNMP trap format below.

Attention:

Do not use unknown traps - Zabbix will not be able to recognise them. Unknown traps can be handled by defining a general event in snmptt.conf:
EVENT general .* "General event" Normal

Configuring Perl trap receiver

Requirements: Perl, Net-SNMP compiled with --enable-embedded-perl (done by default since Net-SNMP 5.4)

Perl trap receiver (look for misc/snmptrap/zabbix_trap_receiver.pl) can be used to pass traps to Zabbix server directly from snmptrapd. To configure it:

- add the perl script to snmptrapd configuration file (snmptrapd.conf), e.g.:
perl do "[FULL PATH TO PERL RECEIVER SCRIPT]";
- configure the receiver, e.g:
\$SNMPTrapperFile = '[TRAP FILE]';
\$DateTimeFormat = '[DATE TIME FORMAT]';

Note:

If script name is not quoted, snmptrapd will refuse to start up with messages, similar to these:
Regexp modifiers "/l" and "/a" are mutually exclusive at (eval 2) line 1, at end of line
Regexp modifier "/l" may not appear twice at (eval 2) line 1, at end of line

SNMP trap format

All customised perl trap receivers and SNMPTT trap configuration must format the trap in the following way:

```
[timestamp] [the trap, part 1] ZBXTRAP [address] [the trap, part 2]
```

where

- [timestamp] - timestamp used for log items
- ZBXTRAP - header that indicates that a new trap starts in this line
- [address] - IP address used to find the host for this trap

Note that "ZBXTRAP" and "[address]" will be cut out from the message during processing. If the trap is formatted otherwise, Zabbix might parse the traps unexpectedly.

Example trap line in the file:

```
11:30:15 2011/07/27 .1.3.6.1.6.3.1.1.5.3 Normal "Status Events" localhost - ZBXTRAP 192.168.1.1 Link down
```

This will result in the following trap for SNMP interface with IP=192.168.1.1:

```
11:30:15 2011/07/27 .1.3.6.1.6.3.1.1.5.3 Normal "Status Events" localhost - Link down on interface 2. Adm
```

3.3 System requirements

Log rotation

Zabbix does not provide any log rotation system - that should be handled by the user. The log rotation should first rename the old file and only later delete it so that no traps are lost:

1. Zabbix opens the trap file at the last known location and goes to step 3
2. Zabbix checks if the currently opened file has been rotated by comparing the inode number to the define trap file's inode number. If there is no opened file, Zabbix resets the last location and goes to step 1.
3. Zabbix reads the data from the currently opened file and sets the new location.
4. The new data are parsed. If this was the rotated file, the file is closed and goes back to step 2.
5. If there was no new data, Zabbix sleeps for 1 second and goes back to step 2.

Attention:

The maximum log file size supported by Zabbix is 2 gigabytes. The log file must be rotated before reaching this limit.

File system

Because of the trap file implementation, Zabbix needs the file system to support inodes to differentiate files (the information is acquired by a `stat()` call).

3.4 Setup example

This example uses `snmptrapd` + `SNMPTT` to pass traps to Zabbix server. Setup:

1. **zabbix_server.conf** - configure Zabbix to start SNMP trapper and set the trap file:
`StartSNMPTrapper=1`
`SNMPTrapperFile=/tmp/my_zabbix_traps.tmp`
2. **snmptrapd.conf** - add `SNMPTT` as the trap handler:
`traphandle default snmptt`
3. **snmptt.ini** - configure output file and time format:
`log_file = /tmp/my_zabbix_traps.tmp`
`date_time_format = %H:%M:%S %Y/%m/%d`
4. **snmptt.conf** - define a default trap format:
`EVENT general .* "General event" Normal`
`FORMAT ZBXTRAP $aA $ar`
5. Create an SNMP item TEST:
Host's SNMP interface IP: 127.0.0.1
Key: `snmptrap["General"]`
Log time format: `hh:mm:ss yyyy/MM/dd`

This results in:

1. Command used to send a trap:
`snmptrap -v 1 -c public 127.0.0.1 '.1.3.6.1.6.3.1.1.5.3' '0.0.0.0' 6 33 '55' .1.3.6.1.6.3.1.1.5.3 s "teststring000"`
2. The received trap:
`15:48:18 2011/07/26 .1.3.6.1.6.3.1.1.5.3.0.33 Normal "General event" localhost - ZBXTRAP 127.0.0.1 127.0.0.1`
3. Value for item TEST:
`15:48:18 2011/07/26 .1.3.6.1.6.3.1.1.5.3.0.33 Normal "General event" localhost - 127.0.0.1`

Note:

This simple example uses `SNMPTT` as **traphandle**. For better performance on production systems, use embedded Perl to pass traps from `snmptrapd` to `SNMPTT` or directly to Zabbix.

4 IPMI checks

Overview

You can monitor the health and availability of Intelligent Platform Management Interface (IPMI) devices in Zabbix. To perform IPMI checks Zabbix server must be initially **configured** with IPMI support.

IPMI is a standardized interface for remote "lights-out" or "out-of-band" management of computer systems. It allows to monitor hardware status directly from the so-called "out-of-band" management cards, independently from the operating system or whether the machine is powered on at all.

Zabbix IPMI monitoring works only for devices having IPMI support (HP iLO, DELL DRAC, IBM RSA, Sun SSP, etc).

See also **known issues** for IPMI checks.

Configuration

Host configuration

A host must be configured to process IPMI checks. An IPMI interface must be added, with the respective IP and port numbers, and IPMI authentication parameters must be defined.

See the **configuration of hosts** for more details.

Server configuration

By default, the Zabbix server is not configured to start any IPMI pollers, thus any added IPMI items won't work. To change this, open the Zabbix server configuration file (**zabbix_server.conf**) as root and look for the following line:

```
# StartIPMIPollers=0
```

Uncomment it and set poller count to, say, 3, so that it reads:

StartIPMIPollers=3

Save the file and restart zabbix_server afterwards.

Item configuration

When **configuring an item** on a host level:

- For *Host interface* select the IPMI IP and port
- Select 'IPMI agent' as the *Type*
- Specify the *IPMI sensor* (for example 'FAN MOD 1A RPM' on Dell Poweredge)
- Enter an item **key** that is unique within the host (say, ipmi.fan.rpm)
- Select the respective type of information ('Numeric (float)' in this case), units (most likely 'rpm') and any other required item attributes

Timeout and session termination

IPMI message timeouts and retry counts are defined in OpenIPMI library. Due to the current design of OpenIPMI, it is not possible to make these values configurable in Zabbix, neither on interface nor item level.

IPMI session inactivity timeout for LAN is 60 +/-3 seconds. Currently it is not possible to implement periodic sending of Activate Session command with OpenIPMI. If there are no IPMI item checks from Zabbix to a particular BMC for more than the session timeout configured in BMC then the next IPMI check after the timeout expires will time out due to individual message timeouts, retries or receive error. After that a new session is opened and a full rescan of the BMC is initiated. If you want to avoid unnecessary rescans of the BMC it is advised to set the IPMI item polling interval below the IPMI session inactivity timeout configured in BMC.

5 Simple checks

5.1 Overview

Simple checks are normally used for remote agent-less checks of services.

Note that Zabbix agent is not needed for simple checks. Zabbix server/proxy is responsible for the processing of simple checks (making external connections, etc).

Examples of using simple checks:

```
net.tcp.service[ftp,,155]
net.tcp.service[http]
net.tcp.service.perf[http,,8080]
```

5.2 Supported simple checks

List of supported simple checks:

Key				
▲	Description	Return value	Parameters	Comments
icmpping[<target>,<packets>,<interval>,<size>,<timeout>]	Host accessibility by ICMP ping.	0 - ICMP ping fails 1 - ICMP ping successful	target - host IP or DNS name packets - number of packets interval - time between successive packets in milliseconds size - packet size in bytes timeout - timeout in milliseconds	Example: icmpping[,4] - if at least one packet of the four is returned, the item will return 1. See also: table of default values .
icmppingloss[<target>,<packets>,<interval>,<size>,<timeout>]				

	Percentage of lost packets.	Float.	target - host IP or DNS name packets - number of packets interval - time between successive packets in milliseconds size - packet size in bytes timeout - timeout in milliseconds	See also: table of default values .
icmppingsec[<target>,<packets>,<interval>,<size>,<timeout>,<mode>]	ICMP ping response time (in seconds).	Float.	target - host IP or DNS name packets - number of packets interval - time between successive packets in milliseconds size - packet size in bytes timeout - timeout in milliseconds mode - one of min, max, avg (default)	If host is not available (timeout reached), the item will return 0. See also: table of default values .
net.tcp.service[service,<ip>,<port>]				

	Check if service is running and accepting TCP connections.	0 - service is down 1 - service is running	service - one of ssh, ntp, ldap, smtp, ftp, http, pop, nntp, imap, tcp, https, telnet (see details) ip - IP address or DNS name (by default, host IP/DNS is used) port - port number (by default standard service port number is used).	Example: net.tcp.service[ftp,,45] can be used to test the availability of FTP server on TCP port 45. Note that with tcp service indicating the port is mandatory. Note that these checks may result in additional messages in system daemon logfiles (SMTP and SSH sessions being logged usually). Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use <code>net.tcp.service[tcp,<ip>,<port>]</code> for checks like these. Services https and telnet supported since Zabbix 2.0. Service ntp only works since Zabbix 2.0.15, despite being available in earlier versions.

net.tcp.service.perf[service,<ip>,<port>]

Service performance check.	0 - service is down sec - number of seconds spent while connecting to the service	service - one of ssh, ntp, ldap, smtp, ftp, http, pop, nntp, imap, tcp, https, telnet (see details) ip - IP address or DNS name (by default, host IP/DNS is used) port - port number (by default standard service port number is used).	Example: net.tcp.service.perf[ssh] can be used to test the speed of initial response from SSH server. Note that with tcp service indicating the port is mandatory. Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use <code>net.tcp.service.perf[tcp,<ip>]</code> , for checks like these. Services https and telnet supported since Zabbix 2.0. Service ntp only works since Zabbix 2.0.15, despite being available in earlier versions. Called <code>tcp_perf</code> before Zabbix 2.0.

Timeout processing

Zabbix will not process a simple check longer than the Timeout seconds defined in the Zabbix server/proxy configuration file.

5.3 ICMP pings

Zabbix uses external utility **fping** for processing of ICMP pings.

The utility is not part of Zabbix distribution and has to be additionally installed. If the utility is missing, has wrong permissions or its location does not match the location set in the Zabbix server/proxy configuration file ('FpingLocation' parameter), ICMP pings (**icmpping**, **icmppingloss**, **icmppingsec**) will not be processed.

See also: [known issues](#)

fping must be executable by the user Zabbix daemons run as and setuid root. Run these commands as user **root** in order to set up correct permissions:

```
shell> chown root:zabbix /usr/sbin/fping
shell> chmod 4710 /usr/sbin/fping
```

After performing the two commands above check ownership of the **fping** executable. In some cases the ownership can be reset by executing the chmod command.

Defaults, limits and description of values for ICMP check parameters:

Parameter	Unit	Description	Fping's flag	Defaults set by	Allowed limits by Zabbix

Warning:

Warning: fping defaults can differ depending on platform and version - if in doubt, check fping documentation.

Zabbix writes IP addresses to be checked by any of three *icmping** keys to a temporary file, which is then passed to **fping**. If items have different key parameters, only ones with identical key parameters are written to a single file.

All IP addresses written to the single file will be checked by fping in parallel, so Zabbix icmp pinger process will spend fixed amount of time disregarding the number of IP addresses in the file.

6 Log file monitoring

Overview

Zabbix can be used for centralized monitoring and analysis of log files with/without log rotation support.

Notifications can be used to warn users when a log file contains certain strings or string patterns.

To monitor a log file you must have:

- Zabbix agent running on the host
- log monitoring item set up

Attention:

The size limit of a monitored log file depends on [large file support](#).

Configuration

Verify agent parameters

Make sure that in the [agent configuration file](#):

- 'Hostname' parameter matches the host name in the frontend
- Servers in the 'ServerActive' parameter are specified for the processing of active checks

Item configuration

Configure a log monitoring [item](#):

Item "New host : Log item"

Host: New host [Select]

Name: Log item

Type: Zabbix agent (active) [v]

Key: log[/var/log/syslog,error] [Select]

Type of information: Log [v]

Update interval (in sec): 10

Keep history (in days): 7 [Clear history]

Status: Active [v]

Log time format: ppppddphh:mm:ss

Specifically for log monitoring items you enter:

Type

Select **Zabbix agent (active)** here.

Key

Set either:

log[/path/to/file/file_name,<regex>,<encoding>,<maxlines>,<mode>]
or

logrt[/path/to/file/regex_describing_filename_pattern,<regex>,<encoding>,<maxlines>,<mode>]
For example:

log[/var/log/syslog]

log[/var/log/syslog,error]

logrt[/home/zabbix/logs/^logfile[0-9]{1,3}\$",,100] - will match a file like "logfile1" (will not match ".logfile1")

logrt[/home/user/logfile_.*_[0-9]{1,3}\$",,100,"pattern_to_match","UTF-8",100] - will collect data from files such as "logfile_abc_1" or "logfile__001".

For more details see **log** and **logrt** entries in the **supported agent item keys** section.

Make sure that the file has read permissions for user 'zabbix' otherwise the item status will be set to 'unsupported'.

Zabbix agent will filter entries of log file by the regex if present.

Select **Log** here.

The parameter defines how often Zabbix agent will check for any changes in the log file. Setting it to 1 second will make sure that you get new records as soon as possible.

In this field you may optionally specify the pattern for parsing the log line timestamp.

If left blank the timestamp will not be parsed.

Supported placeholders:

* **y**: Year (0001-9999)

* **M**: Month (01-12)

* **d**: Day (01-31)

* **h**: Hour (00-23)

* **m**: Minute (00-59)

* **s**: Second (00-59)

For example, consider the following line from the Zabbix agent log file:

" 23480:20100328:154718.045 Zabbix agent started. Zabbix 1.8.2 (revision 11211)."

It begins with six character positions for PID, followed by date, time, and the rest of the line.

Log time format for this line would be

"pppppp:yyyyMMdd:hhmmss".

Note that "p" and ":" chars are just placeholders and can be anything but "yMdhms".

Type of information

Update interval (in sec)

Log time format

Important notes

- The server and agent keep the trace of a monitored log's size and last modification time (for logrt) in two counters. Additionally, since Zabbix **2.0.13**:
- * The agent also internally uses inode numbers (on UNIX/GNU/Linux), file indexes (on Microsoft Windows)
- * On UNIX/GNU/Linux systems it is assumed that the file systems where log files are stored report inode numbers
- * On Microsoft Windows Zabbix agent determines the file system type the log files reside on and uses:
 - * On NTFS file systems 64-bit file indexes.
 - * On ReFS file systems (only from Microsoft Windows Server 2012) 128-bit file IDs.
 - * On file systems where file indexes change (e.g. FAT32, exFAT) a fall-back algorithm is used to track file changes.
- * The inode numbers, file indexes and MD5 sums are internally collected by Zabbix agent. They are not modified.
- * Do not modify the last modification time of log files with 'touch' utility, do not copy a log file with 'cp' utility.
- * If there are several matching log files for 'logrt[]' item and Zabbix agent is following the most recent one.
- * Zabbix ****2.0.15**** fixes an issue [<https://support.zabbix.com/browse/ZBX-9290>] (unexpected results when log files are rotated).
- * The agent starts reading the log file from the point it stopped the previous time.
- * The number of bytes already analyzed (the size counter) and last modification time (the time counter) are stored in the Zabbix database.
- * Whenever the log file becomes smaller than the log size counter known by the agent, the counter is reset to the actual size of the log file.
- * For 'logrt' items, if there are several matching files with the same last modification time in the directory:
 - * before Zabbix ****2.0.13**** the agent will read lexicographically the smallest one.

- * since Zabbix **2.0.13**:
- * The agent tries to correctly analyze all log files with the same modification time and avoid skipping.
- * The agent does not assume any particular log file rotation scheme nor determines one. When presented with a log file, the agent will process all records.
- * Zabbix agent processes new records of a log file once per `//Update interval//` seconds.
- * Zabbix agent does not send more than `**maxlines**` of a log file per second. The limit prevents overloading the server.
- * To find the required string Zabbix will process 4 times more new lines than set in `MaxLinesPerSecond`. The limit is applied to the number of lines processed per second.
- * Additionally, log values are always limited to 50% of the agent send buffer size, even if there are no new records.
- * In the absence of log items all agent buffer size is used for non-log values. When log values come in the buffer, the non-log values are discarded.
- * For log file records longer than 256kB, only the first 256kB are matched against the regular expression.
- * Special note for `"\"` path separators: if `file_format` is `"file\log"`, then there should not be a `"file"` directory.
- * Regular expressions for `'logrt'` are supported in filename only, directory regular expression matching is not supported.
- * On UNIX platforms a `'logrt[]'` item becomes `NOTSUPPORTED` if a directory where the log files are expected does not exist.
- * On Microsoft Windows if a directory does not exist the item does not become `NOTSUPPORTED` (for example, `logrt[C:\log\]`).
- * An absence of log files for `'logrt[]'` item does not make it `NOTSUPPORTED` (before Zabbix 2.0.12 it caused `NOTSUPPORTED`).
- * Errors of reading log files for `'logrt[]'` item are logged as warnings into Zabbix agent log file but do not affect the item value.
- * Zabbix agent log file can be helpful to find out why a `'log[]'` or `'logrt[]'` item became `NOTSUPPORTED`.

7 Calculated items

7.1 Overview

With calculated items you can create calculations on the basis of other items.

Thus, calculated items are a way of creating virtual data sources. The values will be periodically calculated based on an arithmetical expression.

The resulting data will be stored in the Zabbix database as for any other item - this means storing both history and trend values for fast graph generation. Calculated items may be used in trigger expressions, referenced by macros or other entities same as any other item type.

To use calculated items, choose the item type **Calculated**.

7.2 Configurable fields

The **key** is a unique item identifier (per host). You can create any key name using supported symbols.

Calculation definition should be entered in the **Formula** field (named 'Expression' in 1.8.1 and 1.8.2). There is virtually no connection between the formula and the key. The key parameters are not used in formula in any way.

The correct syntax of a simple formula is:

```
func(<key>|<hostname:key>,<parameter1>,<parameter2>,...)
```

Where:

ARGUMENT	DEFINITION
func	One of the functions supported in trigger expressions: last, min, max, avg, count, etc
key	The key of another item whose data you want to use. It may be defined as key or hostname:key . <i>Note:</i> Putting the whole key in double quotes ("...") is strongly recommended to avoid incorrect parsing because of spaces or commas within the key. If there are also quoted parameters within the key, those double quotes must be escaped by using the backslash (\). See Example 5 below.
parameter(s)	Function parameter(s), if required.

Note:

All items that are referenced from the calculated item formula must exist and be collecting data. Also, if you change the item key of a referenced item, you have to manually update any formulas using that key.

Attention:

User macros in the formula will be expanded if used to reference a function parameter or a constant. User macros will NOT be expanded if referencing a function, host name, item key, item key parameter or operator.

A more complex formula may use a combination of functions, operators and brackets. You can use all functions and **operators** supported in trigger expressions. Note that the syntax is slightly different, however logic and operator precedence are exactly the same.

Unlike trigger expressions, Zabbix processes calculated items according to the item update interval, not upon receiving a new value.

Note:

If the calculation result is a float value it will be trimmed to an integer if the calculated item type of information is *Numeric (unsigned)*.

A calculated item may become unsupported in several cases:

1. referenced item(s) not found
2. no data to calculate a function
3. division by zero
4. incorrect syntax used

Note:

Support for calculated items was introduced in Zabbix 1.8.1

7.3 Usage examples

Example 1

Calculating percentage of free disk space on '/'.

Use of function **last**:

```
100*last("vfs.fs.size[/,free]"/last("vfs.fs.size[/,total]"))
```

Zabbix will take the latest values for free and total disk spaces and calculate percentage according to the given formula.

Example 2

Calculating a 10-minute average of the number of values processed by Zabbix.

Use of function **avg**:

```
avg("Zabbix Server:zabbix[wcache,values]",600)
```

Note that extensive use of calculated items with long time periods may affect performance of Zabbix server.

Example 3

Calculating total bandwidth on eth0.

Sum of two functions:

```
last("net.if.in[eth0,bytes]")+last("net.if.out[eth0,bytes]")
```

Example 4

Calculating percentage of incoming traffic.

More complex expression:

```
100*last("net.if.in[eth0,bytes]"/(last("net.if.in[eth0,bytes]")+last("net.if.out[eth0,bytes]"))
```

Example 5

Using aggregated items correctly within a calculated item.

Take note of how double quotes are escaped within the quoted key:

```
last("grpsum[\"video\", \"net.if.out[eth0,bytes]\", \"last\", \"0\"]") / last("grpsum[\"video\", \"nginx_stat.
```

8 Internal checks

8.1 Overview

Internal checks allow to monitor the internal processes of Zabbix. Internal checks are calculated by Zabbix server.

To use this item, choose the **Zabbix internal** item type.

Note:

Internal checks are processed by Zabbix pollers.

8.2 Supported checks

- Parameters without angle brackets are constants - for example, 'host' and 'available' in `zabbix[host,<type>,available]`. Use them in the item key *as is*.

Key	Description	Return value	Comments
▲ zabbix[boottime]	Startup time of Zabbix server process in seconds.	Integer.	In seconds since the epoch.
zabbix[history]	Number of values stored in table HISTORY	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used!
zabbix[history_log]	Number of values stored in table HISTORY_LOG	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used! This item is supported starting from version 1.8.3.
zabbix[history_str]	Number of values stored in table HISTORY_STR	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used!
zabbix[history_text]	Number of values stored in table HISTORY_TEXT	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used! This item is supported starting from version 1.8.3.
zabbix[history_uint]			

Key			
	Number of values stored in table HISTORY_UINT	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used! This item is supported starting from version 1.8.3.
zabbix[host,<type>,available]	Returns availability of a particular type of checks on the host. The value of this item corresponds to availability icons in the host list.	0 - not available, 1 - available, 2 - unknown.	Valid types are: agent , snmp , ipmi , jmx . The item value is calculated according to configuration parameters regarding host unreachability/unavailability . This item is supported starting from version 2.0.0.
zabbix[items]	Number of items in Zabbix database	Integer.	
zabbix[items_unsupported]	Number of unsupported items in Zabbix database	Integer.	
zabbix[java,,<param>]			

Key

Returns
information
associated
with Zabbix
Java
gateway.

If <param>
is **ping**, "1"
is returned.
Can be used
to check
Java
gateway
availability
using
nodata()
trigger
function.

Valid values
for
<param>
are: *ping*,
version

Second
parameter
must be
empty and is
reserved for
future use.

If <param>
is **version**,
version of
Java
gateway is
returned.
Example:
"2.0.0".

**This item is
supported
starting
from
version
2.0.0.**

zabbix[process,<type>,<mode>,<state>]

Time a particular Zabbix process or a group of processes (identified by <type> and <mode>) spent in <state> in percentage. It is calculated for last minute only.

If <mode> is Zabbix process number that is not running (for example, with 5 pollers running <mode> is specified to be 6), such an item will turn into unsupported state. Minimum and maximum refers to the usage percentage for a single process. So if in a group of 3 pollers usage percentages per process were 2, 18 and 66, min would return 2 and max would return 66. Processes report what they are doing in shared memory and the self-monitoring process summarizes that data

Percentage of time. Float.

The following process types are currently supported:

- alerter** - process for sending notifications
- configuration syncer** - process for managing in-memory cache of configuration data
- db watchdog** - sender of a warning message in case DB is not available
- discoverer** - process for discovery of devices
- escalator** - process for escalation of actions
- history syncer** - history DB writer
- housekeeper** - process for removal of old historical data
- http poller** - web monitoring poller
- icmp pinger** - poller for icmping checks
- ipmi poller** - poller for IPMI checks
- java poller** - poller for Java checks
- node watcher** - process for sending historical data and configuration changes

Key			
zabbix[proxy,<name>,<param>]	Access to Zabbix proxy related information.	Integer.	<name> - Proxy name List of supported parameters (<param>): lastaccess - timestamp of last heart beat message received from Proxy For example, zabbix[proxy,"Germany",lastaccess]
zabbix[queue,<from>,<to>]	Number of server monitored items in the Queue which are delayed by <from> to <to> seconds, inclusive.	Integer.	<from> - default: 6 seconds <to> - default: infinity Time-unit symbols (s,m,h,d,w) are supported for these parameters. Parameters from and to are supported starting from version 1.8.3.
zabbix[rcache,<cache>,<mode>]	Availability statistics of Zabbix configuration cache.	Integer (for size); float (for percentage).	Cache: buffer Mode: total - total size of buffer free - size of free buffer pfree - percentage of free buffer used - size of used buffer
zabbix[requiredperformance]			

Key			
	Required performance of the Zabbix server, in new values per second expected.	Float.	Approximately correlates with "Required server performance, new values per second" in <i>Reports</i> → <i>Status of Zabbix</i> . Supported since Zabbix 1.6.2.
zabbix[trends]	Number of values stored in table TRENDS	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used!
zabbix[trends_uint]	Number of values stored in table TRENDS_UINT	Integer.	Do not use if MySQL InnoDB, Oracle or PostgreSQL is used! This item is supported starting from version 1.8.3.
zabbix[triggers]	Number of enabled triggers in Zabbix database, with at least one enabled item on an enabled host.	Integer.	
zabbix[uptime]	Uptime of Zabbix server process in seconds.	Integer.	
zabbix[wcache,<cache>,<mode>]	Statistics and availability of Zabbix write cache. Cache		Mode

values	all	Total number of values processed by Zabbix server, except unsupported items.	Integer.	Counter.
	float	Number of processed float values.	Integer.	Counter.
	uint	Number of processed unsigned integer values.	Integer.	Counter.
	str	Number of processed character/string values.	Integer.	Counter.
	log	Number of processed log items.	Integer.	Counter.
	text	Number of processed text items.	Integer.	Counter.
	not supported	Number of processed unsupported items.	Integer.	Counter. <i>Not supported</i> mode is supported starting with Zabbix 1.8.6 .
	history			A low number indicates performance problems on the database side.
history	pfree	Percentage of free history buffer.	Float.	
	free	Size of free history buffer.	Integer.	
	total	Total size of history buffer.	Integer.	
	used	Size of used history buffer.	Integer.	
trend	pfree	Percentage of free trend buffer.	Float.	
	free	Size of free trend buffer.	Integer.	
	total	Total size of trend buffer.	Integer.	
	used	Size of used trend buffer.	Integer.	

Key				
	text	pfree	Percentage of free text history buffer.	Float.
		free	Size of free text history buffer.	Integer.
		total	Total size of text history buffer.	Integer.
		used	Size of used text history buffer.	Integer.

9 SSH checks

9.1 Overview

SSH checks are performed as agent-less monitoring. Zabbix agent is not needed for SSH checks.

To perform SSH checks Zabbix server must be **initially configured** with SSH2 support.

Attention:

The minimum supported libssh2 library version is 1.0.0.

9.2 Configuration

9.2.1 Server configuration

By default, Zabbix server is not configured to perform SSH checks, thus any added SSH items won't work. To change this, open the Zabbix server configuration file (**zabbix_server.conf**) as root and look for the following line:

```
# SSHKeyLocation=
```

Uncomment it and set full path to a folder where public and private keys will be located:

```
SSHKeyLocation=/home/zabbix/.ssh
```

Save the file and restart zabbix_server afterwards.

/home/zabbix here is the home directory for the *zabbix* user account and *.ssh* is a directory where by default public and private keys will be generated by a **ssh-keygen** command inside the home directory.

Usually installation packages of zabbix-server from different OS distributions create the *zabbix* user account with a home directory in not very well-known places (as for system accounts). For example, for CentOS it's */var/lib/zabbix*, for Debian it's */var/run/zabbix*.

Before starting to generate the keys, an approach to reallocate the home directory to a better known place (intuitively expected) could be considered. This will correspond with the *SSHKeyLocation* Zabbix server configuration parameter mentioned above.

These steps can be skipped if *zabbix* account has been added manually according to the **installation section** because in this case most likely the home directory is already located at */home/zabbix*.

To change the setting for the *zabbix* user account all working processes which are using it have to be stopped:

```
# service zabbix-agent stop
# service zabbix-server stop
```

To change the home directory location with an attempt to move it (if it exists) a command should be executed:

```
# usermod -m -d /home/zabbix zabbix
```

It's absolutely possible that a home directory did not exist in the old place (in the CentOS for example), so it should be created at the new place. A safe attempt to do that is:

```
# test -d /home/zabbix || mkdir /home/zabbix
```

To be sure that all is secure, additional commands could be executed to set permissions to the home directory:

```
# chown zabbix:zabbix /home/zabbix
# chmod 700 /home/zabbix
```

Previously stopped processes now can be started again:

```
# service zabbix-agent start
# service zabbix-server start
```

Now steps to generate public and private keys can be performed by a command:

```
# sudo -u zabbix ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/zabbix/.ssh/id_rsa):
Created directory '/home/zabbix/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/zabbix/.ssh/id_rsa.
Your public key has been saved in /home/zabbix/.ssh/id_rsa.pub.
The key fingerprint is:
90:af:e4:c7:e3:f0:2e:5a:8d:ab:48:a2:0c:92:30:b9 zabbix@it0
The key's randomart image is:
+--[ RSA 2048 ]-----+
|          |
|      .   |
|     o    |
|  .   o   |
|+    . S  |
|. +   o =  |
|E .   * =  |
|=o . . .*  |
|... oo.o+  |
+-----+
```

Note: public and private keys (*id_rsa.pub* and *id_rsa* respectively) have been generated by default in the */home/zabbix/.ssh* directory which corresponds to the Zabbix server *SSHKeyLocation* configuration parameter.

Attention:

Key types other than "rsa" may be supported by the ssh-keygen tool and SSH servers but they may not be supported by libssh2, used by Zabbix.

9.2.2 Shell configuration form

This step should be performed only once for every host that will be monitored by SSH checks.

By using the following command the **public** key file can be installed on a remote host *10.10.10.10* so that then SSH checks can be performed with a *root* account:

```
# sudo -u zabbix ssh-copy-id root@10.10.10.10
The authenticity of host '10.10.10.10 (10.10.10.10)' can't be established.
RSA key fingerprint is 38:ba:f2:a4:b5:d9:8f:52:00:09:f7:1f:75:cc:0b:46.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.10.10.10' (RSA) to the list of known hosts.
root@10.10.10.10's password:
Now try logging into the machine, with "ssh 'root@10.10.10.10'", and check in:
  .ssh/authorized_keys
to make sure we haven't added extra keys that you weren't expecting.
```

Now it's possible to check the SSH login using the default private key (*/home/zabbix/.ssh/id_rsa*) for *zabbix* user account:

```
# sudo -u zabbix ssh root@10.10.10.10
```

If the login is successful, then the configuration part in the shell is finished and remote SSH session can be closed.

9.2.3 Item configuration

Actual command(s) to be executed must be placed in the **Executed script** field in the item configuration.

Multiple commands can be executed one after another by placing them on a new line. In this case returned values also will be formatted as multi lined.

Item parameter	Description	Comments
Key	Unique (per host) item key in format ssh.run[<unique short description>,<ip>,<port>,<encoding>]	<unique short description> is required and should be unique for all SSH items per host Default port is 22, not the port specified in the interface to which this item is assigned
Authentication method	One of the "Password" or "Public key"	
User name	User name to authenticate on remote host. Required	
Public key file	File name of public key if <i>Authentication method</i> is "Public key". Required	Example: <i>id_rsa.pub</i> - default public key file name generated by a command ssh-keygen
Private key file	File name of private key if <i>Authentication method</i> is "Public key". Required	Example: <i>id_rsa</i> - default private key file name
Password or Key passphrase	Password to authenticate or Passphrase if it was used for the private key	Leave the <i>Key passphrase</i> field empty if passphrase was not used See also known issues regarding passphrase usage
Executed script	Executed shell command(s) using SSH remote session	Examples: <i>date +%s</i> <i>service mysql-server status</i> <i>ps auxww grep httpd wc -l</i>

The resulting item configuration should look like this:

Item "Test host: SSH test check (without passphrase)"

Host	Test host		
Name	SSH test check (without passphrase)		
Type	SSH agent ▼		
Key	ssh.run[clear]	Select	
Host interface	10.10.10.10 : 10050 ▼		
Authentication method	Public key ▼		
User name	root		
Public key file	id_rsa.pub		
Private key file	id_rsa		
Key passphrase			
Executed script	service mysql-server status		
Type of information	Text ▼		
Units			
Use custom multiplier	☐		1
Update interval (in sec)	60		

Attention:

libssh2 library may truncate executable scripts to ~32kB.

10 Telnet checks

10.1 Overview

Telnet checks are performed as agent-less monitoring. Zabbix agent is not needed for Telnet checks.

10.2 Configurable fields

Actual command(s) to be executed must be placed in the **Executed script** field in the item configuration.

Multiple commands can be executed one after another by placing them on a new line. In this case returned value also will be formatted as multi lined.

Supported characters that the shell prompt can end with:

- \$
- #
- .
- %

Note:

A telnet prompt line which ended with one of these characters will be removed from the returned value, but only for the first command in the commands list, i.e. only at a start of the telnet session.

Key	Description	Comments
telnet.run [<uniqueid>,<short description>,<ip>,<port>,<encoding>]	Run a command on a remote device using telnet connection	

Attention:

If a telnet check returns a value with non-ASCII characters and in non-UTF8 encoding then the *<encoding>* parameter of the key should be properly specified. See [encoding of returned values](#) page for more details.

11 External checks

11.1 Overview

External check is a check executed by Zabbix server by running a shell script or a binary. However, when hosts are monitored by a Zabbix proxy, the external checks are executed by the proxy.

External checks do not require any agent running on a host being monitored.

The syntax of the item key is:

```
script[<parameter1>,<parameter2>,...]
```

Where:

ARGUMENT	DEFINITION
script	Name of a shell script or a binary.
parameter(s)	Optional command line parameters.

If you don't want to pass any parameters to the script you may use:

```
script[] or
script
```

Zabbix server will look in the directory defined as the location for external scripts (parameter 'ExternalScripts' in [Zabbix server configuration file](#)) and execute the command. The command will be executed as the user Zabbix server runs as, so any access permissions or environment variables should be handled in a wrapper script, if necessary, and permissions on the command should allow that user to execute it. Only commands in the specified directory are available for execution.

Zabbix uses the standard output of the script as the value (the full output with trimmed trailing whitespace is returned since Zabbix 2.0). Standard error and exit codes are discarded.

In case the requested script is not found or Zabbix server has no permissions to run it, the item will be marked as unsupported and an according error message will be returned. In case of a timeout, the item will be marked as unsupported as well, an according error message will be displayed and the forked process for the script will be killed.

Warning:

Do not overuse external checks! As each script requires starting a fork process by Zabbix server, running many scripts can decrease Zabbix performance a lot.

11.2 Usage example

Executing the script **check_oracle.sh** with the first parameters "-h". The second parameter will be replaced by IP address or DNS name, depending on the selection in the host properties.

```
check_oracle.sh["-h","{HOST.CONN}"]
```

Assuming host is configured to use IP address, Zabbix will execute:

```
check_oracle.sh "-h" "192.168.1.4"
```

12 Aggregate checks

Overview

In aggregate checks Zabbix server collects aggregate information from items by doing direct database queries.

Aggregate checks do not require any agent running on the host being monitored.

Syntax

The syntax of the aggregate item key is:

```
groupfunc["host group","item key",itemfunc,timeperiod]
```

Supported group functions (groupfunc) are:

Group function	Description
<i>grpavg</i>	Average value
<i>grpmax</i>	Maximum value
<i>grpmin</i>	Minimum value
<i>grpsum</i>	Sum of values

Multiple host groups may be included by inserting a comma-delimited array.

All items that are referenced from the aggregate item key must exist and be collecting data. Only enabled items on enabled hosts are included in the calculations.

Attention:

The key of the aggregate item must be updated manually, if the item key of a referenced item is changed.

Supported item functions (itemfunc) are:

Item function	Description
<i>avg</i>	Average value
<i>count</i>	Number of values
<i>last</i>	Last value
<i>max</i>	Maximum value
<i>min</i>	Minimum value
<i>sum</i>	Sum of values

The **timeperiod** parameter specifies a time period of latest collected values. **Supported unit symbols** can be used in this parameter for convenience, for example '5m' (minutes) instead of '300' (seconds) or '1d' (day) instead of '86400' (seconds).

Warning:

An amount of values (prefixed with #) is not supported in the timeperiod.

Timeperiod is ignored by the server if the third parameter (item function) is *last*.

Note:

If the aggregate results in a float value it will be trimmed to an integer if the aggregated item type of information is *Numeric (unsigned)*.

An aggregate item may become unsupported in several cases:

- none of the referenced items is found (which may happen if the item key is incorrect, none of the items exists or all included groups are incorrect)
- no data to calculate a function

Usage examples

Examples of keys for aggregate checks:

Example 1

Total disk space of host group 'MySQL Servers'.

```
grpsum["MySQL Servers","vfs.fs.size[/,total]",last,0]
```

Example 2

Average processor load of host group 'MySQL Servers'.

```
grpavg["MySQL Servers","system.cpu.load[,avg1]",last,0]
```

Example 3

5-minute average of the number of queries per second for host group 'MySQL Servers'.

```
grpavg["MySQL Servers",mysql.qps,avg,5m]
```

Example 4

Average CPU load on all hosts in multiple host groups.

```
grpavg[["Servers A","Servers B","Servers C"],system.cpu.load,last,0]
```

13 Trapper items

Overview

Trapper items accept incoming data instead of querying for it.

It is useful for any data you might want to "push" into Zabbix.

To use a trapper item you must:

- have a trapper item set up in Zabbix
- send in the data into Zabbix

Configuration

Item configuration

To configure a trapper item:

- Go to: *Configuration* → *Hosts*
- Click on *Items* in the row of the host
- Click on *Create item*
- Enter parameters of the item in the form

The screenshot shows the 'Item "New host : Trapper item"' configuration form in Zabbix. The form has the following fields and values:

- Host:** New host (with a 'Select' button)
- Name:** Trapper item
- Type:** Zabbix trapper (dropdown menu)
- Key:** trap (with a 'Select' button)
- Type of information:** Text (dropdown menu)
- Keep history (in days):** 7 (with a 'Clear history' button)
- Status:** Active (dropdown menu)
- Allowed hosts:** (empty text field)

The fields that require specific information for trapper items are:

Type

Key

Type of information

Select **Zabbix trapper** here.

Enter a key that will be used to recognize the item when sending in data.

Select the type of information that will correspond the format of data that will be sent in.

Allowed hosts

If specified, the trapper will accept incoming data only from this comma-delimited list of hosts. No spaces are allowed. Hosts are identified by IP address/DNS name. For example:
Single IP: 192.168.1.33
List of IP addresses:
192.168.56.5,192.168.56.6,192.168.56.7
Single DNS name: testzabbix.zabbix.com
List of DNS names:
testzabbix,testzabbix.zabbix.com,testzabbix1.zabbix.com

Note:

You may have to wait up to 60 seconds after saving the item until the server picks up the changes from a configuration cache update, before you can send in values.

Sending in data

In the simplest of cases, we may use `zabbix_sender` utility to send in some 'test value':

```
zabbix_sender -z <server IP address> -p 10051 -s "New host" -k trap -o "test value"
```

To send in the value we use these keys:

- z - to specify Zabbix server IP address
- p - to specify Zabbix server port number (10051 by default)
- s - to specify the host (make sure to use the 'technical' **host name** here, instead of the 'visible' name)
- k - to specify the key of the item we just defined
- o - to specify the actual value to send

Attention:

Zabbix trapper process does not expand macros used in the item key in attempt to check corresponding item key existence for targeted host.

Display

This is the result in *Monitoring* → *Latest data*:

Trapper item	13 Jan 2012 11:13:52	test value	-	History

14 JMX monitoring

14.1 Overview

JMX monitoring can be used to monitor JMX counters of a Java application.

In Zabbix 1.8, if you wanted to monitor JMX counters of a Java application, your best choice would have been the Zapcat [JMX Zabbix Bridge](#). You would either modify the source code of your application to reference the Zapcat JAR file and programmatically start a Zabbix agent, or you would install a ready-made Zapcat plugin for applications that support it (such as Jetty or Tomcat).

Zabbix 2.0 adds native support for JMX monitoring by introducing a new Zabbix daemon called "Zabbix Java gateway".

When Zabbix server wants to know the value of a particular JMX counter on a host, it asks the Zabbix **Java gateway**, which in turn uses the [JMX management API](#) to query the application of interest remotely.

Note:

For more details on Zabbix Java gateway, including where to get it and how to set it up see [this section](#) of the manual.

14.2 Enabling remote JMX monitoring for Java application

A Java application does not need any additional software installed, but it needs to be started with the command-line options specified below to have support for remote JMX monitoring.

As a bare minimum, if you just wish to get started by monitoring a simple Java application on a local host with no security enforced, start it with these options:

```
java \
-Dcom.sun.management.jmxremote \
-Dcom.sun.management.jmxremote.port=12345 \
-Dcom.sun.management.jmxremote.authenticate=false \
-Dcom.sun.management.jmxremote.ssl=false \
-jar /usr/share/doc/openjdk-6-jre-headless/demo/jfc/Notepad/Notepad.jar
```

This makes Java listen for incoming JMX connections on port 12345, from local host only, and tells it not to require authentication or SSL.

If you want to allow connections on another interface, set the `-Djava.rmi.server.hostname` parameter to the IP of that interface.

If you wish to be more stringent about security, there are many other Java options available to you. For instance, the next example starts the application with a more versatile set of options and opens it to a wider network, not just local host.

```
java \
-Djava.rmi.server.hostname=192.168.3.14 \
-Dcom.sun.management.jmxremote \
-Dcom.sun.management.jmxremote.port=12345 \
-Dcom.sun.management.jmxremote.authenticate=true \
-Dcom.sun.management.jmxremote.password.file=/etc/java-6-openjdk/management/jmxremote.password \
-Dcom.sun.management.jmxremote.access.file=/etc/java-6-openjdk/management/jmxremote.access \
-Dcom.sun.management.jmxremote.ssl=true \
-Djavax.net.ssl.keyStore=$YOUR_KEY_STORE \
-Djavax.net.ssl.keyStorePassword=$YOUR_KEY_STORE_PASSWORD \
-Djavax.net.ssl.trustStore=$YOUR_TRUST_STORE \
-Djavax.net.ssl.trustStorePassword=$YOUR_TRUST_STORE_PASSWORD \
-Dcom.sun.management.jmxremote.ssl.need.client.auth=true \
-jar /usr/share/doc/openjdk-6-jre-headless/demo/jfc/Notepad/Notepad.jar
```

Most (if not all) of these settings can be specified in `/etc/java-6-openjdk/management/management.properties` (or wherever that file is on your system).

Note that if you wish to use SSL, you have to modify `startup.sh` script by adding `-Djavax.net.ssl.*` options to Java gateway, so that it knows where to find key and trust stores.

See [Monitoring and Management Using JMX](#) for a detailed description.

14.3 Configuring JMX interfaces and items in Zabbix GUI

With Java gateway running, server knowing where to find it and a Java application started with support for remote JMX monitoring, it is time to configure the interfaces and items in Zabbix GUI.

Configuring JMX interface

You begin by creating a JMX-type interface on the host of interest:

The screenshot shows the Zabbix GUI 'Host inventory' tab. The 'Host name' is 'Host with Notepad'. The 'Groups' section shows 'Java' selected. The 'JMX interfaces' section shows a new interface added with IP address '192.168.3.14' and port '12345'.

Agent interfaces	IP address	DNS name	Connect to	Port	Default
Add					

SNMP interfaces	IP address	DNS name	Connect to	Port	Default
Add					

JMX interfaces	IP address	DNS name	Connect to	Port	Default
Add	192.168.3.14		IP	12345	☒ Remove

IPMI interfaces	IP address	DNS name	Connect to	Port	Default
Add					

Monitored by proxy: (no proxy) [Add](#)

Status: Monitored

Adding JMX agent item

For each JMX counter you are interested in you add an item of type **JMX agent** attached to that interface. If you have configured authentication on your Java application, then you also specify username and password.

The key in the screenshot below says `jmx["java.lang:type=Memory", "HeapMemoryUsage.used"]`. The JMX item key syntax is similar to Zapcat items, except that a comma is used for separating arguments instead of "]"["]. The key consists of 2 parameters:

- object name - which represents the object name of an MBean
- attribute name - an MBean attribute name with optional composite data field names separated by dots

See below for more detail on JMX item keys.

If you wish to monitor a Boolean counter that is either "true" or "false", then you specify type of information as "Numeric (unsigned)" and data type as "Boolean". Server will store Boolean values as 1 or 0, respectively.

JMX item keys in more detail

Simple attributes

An MBean object name is nothing but a string which you define in your Java application. An attribute name, on the other hand, can be more complex. In case an attribute returns primitive data type (an integer, a string etc.) there is nothing to worry about, the key will look like this:

```
jmx[com.example:Type=Hello,weight]
```

In this example an object name is "com.example:Type=Hello", attribute name is "weight" and probably the returned value type should be "Numeric (float)".

Attributes returning composite data

It becomes more complicated when your attribute returns composite data. For example: your attribute name is "apple" and it returns a hash representing its parameters, like "weight", "color" etc. Your key may look like this:

```
jmx[com.example:Type=Hello,apple.weight]
```

This is how an attribute name and a hash key are separated, by using a dot symbol. Same way, if an attribute returns nested composite data the parts are separated by a dot:

```
jmx[com.example:Type=Hello,fruits.apple.weight]
```

Problem with dots

So far so good. But what if an attribute name or a hash key contains dot symbol? Here is an example:

```
jmx[com.example:Type=Hello,all.fruits.apple.weight]
```

That's a problem. How to tell Zabbix that attribute name is "all.fruits", not just "all"? How to distinguish a dot that is part of the name from the dot that separates an attribute name and hash keys?

Before **2.0.4** Zabbix Java gateway was unable to handle such situations and users were left with UNSUPPORTED items. Since 2.0.4 this is possible, all you need to do is to escape the dots that are part of the name with a backslash:

```
jmx[com.example:Type=Hello,all\.fruits.apple.weight]
```

Same way, if your hash key contains a dot you escape it:

```
jmx[com.example:Type=Hello,all\.fruits.apple.total\.weight]
```

Other issues

A backslash character should be escaped as well:

```
jmx[com.example:type=Hello,c:\\documents]
```

If the object name or attribute name contains spaces or commas double-quote it:

```
jmx["com.example:Type=Hello","fruits.apple.total weight"]
```

This is actually all there is to it. Happy JMX monitoring!

15 ODBC monitoring

15.1 Overview

ODBC monitoring corresponds to the *Database monitor* item type in the Zabbix frontend.

ODBC is a C programming language middle-ware API for accessing database management systems (DBMS). The ODBC concept was developed by Microsoft and later ported to other platforms.

Zabbix may query any database, which is supported by ODBC. To do that, Zabbix does not directly connect to the databases, but uses the ODBC interface and drivers set up in ODBC. This function allows for more efficient monitoring of different databases for multiple purposes - for example, checking specific database queues, usage statistics and so on. Zabbix supports unixODBC and iODBC, which are the two most commonly used open source ODBC API implementations.

15.2 Installing unixODBC

The suggested way of installing unixODBC is to use the Linux operating system default package repositories. In the most popular Linux distributions unixODBC is included in the package repository by default. If it's not available, it can be obtained at the unixODBC homepage: <http://www.unixodbc.org/download.html>.

unixODBC installation using the *yum* package manager:

```
shell> yum -y install unixODBC unixODBC-devel
```

Note:

The unixODBC-devel package is needed to compile Zabbix with unixODBC support.

15.3 Installing unixODBC drivers

A unixODBC database driver should be installed for the database, which will be monitored. unixODBC has a list of supported databases and drivers: <http://www.unixodbc.org/drivers.html>. In some Linux distributions database drivers are included in package repositories.

For example, a MySQL database driver can be installed using the *yum* package manager:

```
shell> yum install mysql-connector-odbc
```

15.4 Configuring unixODBC

ODBC configuration is done by editing the **odbcinst.ini** and **odbc.ini** files. To verify the configuration file location, type:

```
shell> odbcinst -j
```

odbcinst.ini is used to list the installed ODBC database drivers:

```
[mysql]
Description = ODBC for MySQL
Driver      = /usr/lib/libmyodbc5.so
```

Parameter details:

Attribute	Description
<i>mysql</i>	Database driver name.
<i>Description</i>	Database driver description.
<i>Driver</i>	Database driver library location.

odbc.ini is used to define data sources:

```
[test]
Description = MySQL test database
Driver      = mysql
Server      = 127.0.0.1
User        = root
Password    =
Port        = 3306
Database    = zabbix
```

Parameter details:

Attribute	Description
<i>test</i>	Data source name (DSN).
<i>Description</i>	Data source description.
<i>Driver</i>	Database driver name - as specified in odbcinst.ini
<i>Server</i>	Database server IP/DNS.
<i>User</i>	Database user for connection.
<i>Password</i>	Database user password.
<i>Port</i>	Database connection port.
<i>Database</i>	Database name.

To verify if ODBC connection is working successfully, a connection to database should be tested. That can be done with the **isql** utility (included in the unixODBC package):

```
shell> isql test
+-----+
| Connected! |
|          |
| sql-statement |
| help [tablename] |
| quit      |
|          |
+-----+
SQL>
```

15.5 Compiling Zabbix with ODBC support

To enable ODBC support, Zabbix should be compiled with one of following flags:

```
--with-iodbc[=ARG]    use odbc driver against iODBC package [default=no],
--with-unixodbc[=ARG] use odbc driver against unixODBC package
```

Note:

See more about Zabbix installation from the [source code](#).

15.6 Item configuration in Zabbix frontend

Configure a database monitoring **item**:

Item

Host

Zabbix server

Select

Name

MySQL host count

Type

Database monitor

Key

db.odbc.select[mysql-simple-check]

Select

Additional parameters

DSN=test
user=zabbix
password=
sql=select count(*) from hosts;

Type of information

Numeric (unsigned)

Specifically for database monitoring items you must enter:

Type	Select <i>Database monitor</i> here.
Key	Enter db.odbc.select [unique_description] The unique description will serve to identify the item in triggers etc.
Additional parameters	DSN - data source name (as specified in <code>odbc.ini</code>) user - database user name (optional if user is specified in <code>odbc.ini</code>) password - database user password (optional if password is specified in <code>odbc.ini</code>) sql - SQL query
Type of information	It is important to know what type of information will be returned by the query, so that it is selected correctly here. With an incorrect <i>type of information</i> the item will turn unsupported.

15.7 Important notes

- Zabbix does not limit the query execution time. It is up to the user to choose queries that can be executed in a reasonable amount of time.
- Starting from Zabbix 2.0.8 the **Timeout** parameter value from Zabbix server is used as the ODBC login timeout (note that depending on ODBC drivers the login timeout setting might be ignored).
- The query must return one value only.
- If a query returns more than one column, only the first column is read.
- If a query returns more than one line, only the first line is read.
- The SQL command must begin with `select`.
- The SQL command mustn't contain any line breaks.
- See also [known issues](#) for ODBC checks

A user parameter has the following syntax:

`UserParameter=<key>,<command>`

As you can see, a user parameter also contains a key. The key will be necessary when configuring an item. Enter a key of your choice that will be easy to reference (it must be unique within a host). Restart the agent.

Then, when **configuring an item**, enter the key to reference the command from the user parameter you want executed.

User parameters are commands executed by Zabbix agent. Up to 512KB of data can be returned. The return value of the command is standard output; standard error is discarded. `/bin/sh` is used as a command line interpreter under UNIX operating systems. User parameters obey the agent check timeout; if timeout is reached the forked user parameter process is terminated.

See also:

- [Step-by-step tutorial](#) on making use of user parameters
- [Command execution](#)

Examples of simple user parameters

A simple command:

`UserParameter=ping,echo 1`

The agent will always return '1' for an item with 'ping' key.

A more complex example:

`UserParameter=mysql.ping,mysqladmin -uroot ping|grep -c alive`

The agent will return '1', if MySQL server is alive, '0' - otherwise.

Flexible user parameters

Flexible user parameters accept parameters with the key. This way a flexible user parameter can be the basis for creating several items.

Flexible user parameters have the following syntax:

`UserParameter=key[*],command`

Parameter	Description
Key	Unique item key. The [*] defines that this key accepts parameters within the brackets.
Command	Parameters are given when configuring the item. Command to be executed to evaluate value of the key. Use positional references \$1...\$9 to refer to the respective parameter in the item key. Zabbix parses the parameters enclosed in [] of the item key and substitutes \$1,...,\$9 in the command accordingly. \$0 will be substituted by the original command (prior to expansion of \$0,...,\$9) to be run.

Note:

To use positional references unaltered, specify double dollar sign - for example, `awk '{print $$2}'`. In this case `$$2` will actually turn into `$2` when executing the command.

Note that positional references with the \$ sign are interpreted by Zabbix agent regardless of whether they are enclosed between double (") or single (') quotes.

Attention:

Unless `UnsafeUserParameters` agent daemon configuration option is enabled, it is not allowed to pass flexible parameters containing these symbols: \ ' " * ? [] { } ~ \$! & ; () < > | # @. Additionally, newline is not allowed either.

Note:

User parameters that return text (character, log, text types of information) now can return whitespace only as well, setting the return value to an empty string (supported since 2.0). If non-valid value is returned, `ZBX_NOTSUPPORTED` will be sent back by the agent.

Example 1

Something very simple:

```
UserParameter=ping[*],echo $1
```

We may define unlimited number of items for monitoring all having format ping[something].

- ping[0] - will always return '0'
- ping[aaa] - will always return 'aaa'

Example 2

Let's add more sense!

```
UserParameter=mysql.ping[*],mysqladmin -u$1 -p$2 ping | grep -c alive
```

This parameter can be used for monitoring availability of MySQL database. We can pass user name and password:

```
mysql.ping[zabbix,our_password]
```

Example 3

How many lines matching a regular expression in a file?

```
UserParameter=wc[*],grep -c "$2" $1
```

This parameter can be used to calculate number of lines in a file.

```
wc[/etc/passwd,root]  
wc[/etc/services,zabbix]
```

1 Extending Zabbix agents

This tutorial provides step-by-step instructions on how to extend the functionality of Zabbix agent with the use of a **user parameter**.

Step 1

Write a script or command line to retrieve required parameter.

For example, we may write the following command in order to get total number of queries executed by a MySQL server:

```
mysqladmin -uroot status|cut -f4 -d":"|cut -f1 -d"S"
```

When executed, the command returns total number of SQL queries.

Step 2

Add this command to agent's configuration file.

Add the command to zabbix_agentd.conf:

```
UserParameter=mysql.questions,mysqladmin -uroot status|cut -f4 -d":"|cut -f1 -d"S"
```

mysql.questions is a unique identifier. It can be any string, for example, queries.

Test this parameter by using **zabbix_get** utility.

Step 3

Restart Zabbix agent.

Agent will reload configuration file.

Step 4

Add new item for monitoring.

Add new item with Key=mysql.questions to the monitored host. Type of the item must be either Zabbix Agent or Zabbix Agent (active).

Be aware that type of returned values must be set correctly on Zabbix server. Otherwise Zabbix won't accept them.

5 Windows performance counters

Overview

You can effectively monitor Windows performance counters using the *perf_counter[]* key.

For example:

```
perf_counter["\Processor(0)\Interrupts/sec"]
```

or

```
perf_counter["\Processor(0)\Interrupts/sec", 10]
```

For more information on using this key, see [WIN32-specific item keys](#).

In order to get a full list of performance counters available for monitoring, you may run:

```
typeperf -qx
```

Numeric representation

As the naming of performance counters may differ on different Windows servers, depending on local settings, it introduces a certain problem when creating a template for monitoring several Windows machines having different locales.

At the same time every performance counter can also be referred to by its numeric form, which is unique and exactly the same regardless of language settings, so you might use the numeric representation instead of strings.

To find out the numeric equivalents, run **regedit**, then find *HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\00000000*.

The registry entry contains information like this:

```
1
1847
2
System
4
Memory
6
% Processor Time
10
File Read Operations/sec
12
File Write Operations/sec
14
File Control Operations/sec
16
File Read Bytes/sec
18
File Write Bytes/sec
....
```

Here you can find the corresponding numbers for each string part of the performance counter, like in '\System\% Processor Time':

System → 2

% Processor Time → 6

Then you can use these numbers to represent the path in numbers:

```
\2\6
```

Performance counter parameters

You can deploy some PerfCounter parameters for the monitoring of Windows performance counters.

For example, you can add these to the Zabbix agent configuration file:

```
PerfCounter=UserPerfCounter1,"\Memory\Page Reads/sec",30
or
PerfCounter=UserPerfCounter2,"\4\24",30
```

With such parameters in place, you can then simply use *UserPerfCounter1* or *UserPerfCounter2* as the keys for creating the respective items.

Remember to restart Zabbix agent after making changes to the configuration file.

Troubleshooting

Sometimes Zabbix agent cannot retrieve performance counter values in Windows 2000-based systems, because the pdh.dll file is outdated. It shows up as failure messages in Zabbix agent and server log files. In this case pdh.dll should be updated to a newer 5.0.2195.2668 version.

6 Mass update

Overview

Sometimes you may want to change some attribute for a number of items at once. Instead of opening each individual item for editing, you may use the mass update function for that.

Using mass update

To mass-update some items, do the following:

- Mark the checkboxes of the items to update in the list
- Select *Mass update* from the dropdown below and click on *Go*
- Mark the checkboxes of the attributes to update
- Enter new values for the attributes and click on *Update*

Mass update

Type	☐	Original
Host interface	☐	Original
SNMP community	☐	Original
SNMPv3 security name	☐	Original
SNMPv3 security level	☐	Original
SNMPv3 auth passphrase	☐	Original
SNMPv3 priv passphrase	☐	Original
Port	☐	Original
Type of information	☐	Original
Data type	☐	Original
Units	☐	Original
Authentication method	☐	Original
User name	☐	Original
Public key file	☐	Original
Private key file	☐	Original
Password	☐	Original
Custom multiplier (0 - Disabled)	☐	Original
Update interval (in sec)	☒	30
Flexible intervals	☐	Original
Keep history (in days)	☒	7
Keep trends (in days)	☒	365
Status	☐	Original
Log time format	☐	Original
Store value	☐	Original
Show value	☐	Original
Allowed hosts	☐	Original
Applications	☐	Original
Description	☐	Original

7 Value mapping

Overview

For a more "human" representation of received values, you can use value maps that contain the mapping between numeric values and string representations.

Value mappings can be used in both the Zabbix frontend and notifications sent by email/SMS/jabber etc.

For example, an item which has value '0' or '1' can use value mapping to represent the values in a human-readable form:

- '0' => 'Not Available'
- '1' => 'Available'

Thus, when **configuring items** you can use a value map to "humanize" the way an item value will be displayed. To do that, you refer to the name of a previously defined value map in the *Show value* field.

Note:

Value mapping can only be used with items having type *Numeric (unsigned)*.

Configuration

To define a value map:

- Go to: *Administration* → *General*
- Select *Value mapping* from the dropdown
- Click on *Create value map* (or on the name of an existing map)

Value mapping

Name

Windows service state

Mappings

Value		Mapped to	
0	⇒	Running	Remove
1	⇒	Paused	Remove
2	⇒	Start pending	Remove
3	⇒	Pause pending	Remove
4	⇒	Continue pending	Remove
5	⇒	Stop pending	Remove
6	⇒	Stopped	Remove
7	⇒	Unknown	Remove
255	⇒	No such service	Remove

[Add](#)

Save

Delete

Cancel

Parameters of a value map:

Parameter	Description
<i>Name</i>	Unique name of a set of value mappings.

Parameter	Description
<i>Mapping</i>	Individual mappings - pairs of numeric values and their string representations.
<i>New mapping</i>	A single mapping for addition.

How this works

For example, one of the predefined agent items 'Ping to the server (TCP)' uses an existing value map called 'Service state' to display its values.

In the item **configuration form** you can see a reference to this value map in the *Show value* field:

Show value [show value mappings](#)

So in *Monitoring* → *Latest data* the mapping is put to use to display 'Up' (with the raw value in parentheses).

Ping to the server (TCP)	12 Jan 2012 09:52:35	Up (1)	-	Graph

Note:

A value being displayed in a human-readable form is also easier to understand when receiving notifications.

Without a predefined value map you would only get this:

Ping to the server (TCP)	12 Jan 2012 09:55:35	1	-	Graph

So in this case you would either have to guess what the '1' stands for or do a search of documentation to find out.

8 Applications

Overview

Applications are used to group items in logical groups.

For example, the *MySQL Server* application can hold all items related to the MySQL server: availability of MySQL, disk space, processor load, transactions per second, number of slow queries, etc.

Applications are also used for grouping web scenarios.

If you are using applications, then in *Monitoring* → *Latest data* you will see items and web scenarios grouped under their respective applications.

Configuration

To work with applications you must first create them and then link items or web scenarios to them.

To create an application, do the following:

- Go to *Configuration* → *Hosts* or *Templates*
- Click on *Applications* next to the required host or template
- Click on *Create application*
- Enter the application name and save it

Application

Host

Template OS Linux

Name

CPU

You can also create a new application directly in the item properties form.

Items are linked to applications in the item properties form. Select one or more applications the item will belong to.

Web scenarios are linked to applications in the web scenario definition form. Select the application the scenario will belong to.

9 Queue

Overview

The queue displays items that are waiting for a refresh. The queue is just a **logical** representation of data from the database. There is no IPC queue or any other queue mechanism in Zabbix.

Statistics shown by the queue is a good indicator of the performance of Zabbix server.

Reading the queue

To read the queue, go to *Administration* → *Queue*. *Overview* should be selected in the dropdown to the right.

Items	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes
Zabbix agent	0	1	0	0	1	0
Zabbix agent (active)	0	0	0	0	0	0
Simple check	0	0	0	0	0	0
SNMPv1 agent	0	0	0	0	0	0
SNMPv2 agent	0	0	0	0	0	0
SNMPv3 agent	0	0	0	0	0	0
Zabbix internal	0	0	0	0	0	0
Zabbix aggregate	0	0	0	0	0	0
External check	0	0	0	0	0	0
Database monitor	0	0	0	0	0	0
IPMI agent	0	0	0	0	0	0
SSH agent	0	0	0	0	0	0
TELNET agent	0	0	0	0	0	0
JMX agent	0	0	0	0	0	0
Calculated	0	0	0	0	0	0

The picture here is generally "green" so we may assume that the server is doing fine.

The queue shows one item waiting for 10 seconds and one for 5 minutes. Nice, it would be great to know what items these are.

To do just that, select *Details* in the dropdown in the upper right corner. Now you can see a list of those delayed items.

Next check	Delayed by	Host	Name
20 Jan 2012 10:44:07	7m 44s	Zabbix server	Version of zabbix_agent(d) running
20 Jan 2012 10:51:37	14s	Zabbix server	Ping to the server (TCP)

With these details provided it may be possible to find out why these items might be delayed.

With one or two delayed items there perhaps is no cause for alarm. They might get updated in a second. However, if you see a bunch of items getting delayed for too long, there might be a more serious problem.

Items	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes
Zabbix agent	0	1	0	1	0	45
Zabbix agent (active)	0	0	0	0	0	0

Is the agent down?

Delay for remote node items

Queue information from a child node is not up-to-date. The master node receives historical data with a certain delay (normally, up-to 10 seconds for inter-node data transfer), so the information is delayed.

The information from a child node also depends on:

- performance of the child node
- communications between master and child nodes
- possible local time difference between master and child nodes

Queue item

A special internal item **zabbix[queue,<from>,<to>]** can be used to monitor the health of the queue in Zabbix. It will return the number of items delayed by the set amount of time. For more information see [Internal items](#).

3 Triggers

Overview

Triggers are logical expressions that "evaluate" data gathered by items and represent the current system state.

While items are used to gather system data, it is highly impractical to follow these data all the time waiting for a condition that is alarming or deserves attention. The job of "evaluating" data can be left to trigger expressions.

Trigger expressions allow to define a threshold of what state of data is "acceptable". Therefore, should the incoming data surpass the acceptable state, a trigger is "fired" - or changes status to PROBLEM.

A trigger may have the following status:

VALUE	DESCRIPTION
OK	This is a normal trigger state. Called FALSE in older Zabbix versions.
PROBLEM	Normally means that something happened. For example, the processor load is too high. Called TRUE in older Zabbix versions.

Trigger status (the expression) is recalculated every time Zabbix server receives a new value that is part of the expression.

Triggers are evaluated based on **history** data only; trend data are never considered.

If time-based functions (**nodata()**, **date()**, **dayofmonth()**, **dayofweek()**, **time()**, **now()**) are used in the expression, the trigger is recalculated every 30 seconds by a zabbix *timer* process. If both time-based and non-time-based functions are used in an expression, it is recalculated when a new value is received **and** every 30 seconds.

You can **build trigger expressions** with different degrees of complexity.

1 Configuring a trigger

Overview

To configure a trigger, do the following:

- Go to: *Configuration* → *Hosts*
- Click on *Triggers* in the row of the host
- Click on *Create trigger* to the right (or on the trigger name to edit an existing trigger)
- Enter parameters of the trigger in the form

Configuration

The **Trigger** tab contains all the essential trigger attributes.

TriggerDependencies

Name

Low free disk space on {HOST.NAME} volume /

Expression

{Zabbix server.vfs.fs.size[/,pfree].last(0)}<10

Add

Expression constructor

Multiple PROBLEM events generation

☐

Comments

URL

Severity

Not classified

Information

Warning

Average

High

Disaster

Enabled

☒

Save

Clone

Delete

Cancel

Parameter	Description
<i>Name</i>	Trigger name. The name may contain macros. \$1, \$2...\$9 macros can be used to refer to the first, second...ninth constant of the expression. <i>Note:</i> \$1-\$9 macros will resolve correctly if referring to constants in relatively simple, straightforward expressions. For example, the name "Processor load above \$1 on {HOST.NAME}" will automatically change to "Processor load above 5 on New host" if the expression is {New host:system.cpu.load[percpu,avg1].last(0)}>5
<i>Expression</i>	Logical expression used for calculating the trigger state.
<i>Multiple PROBLEM events generation</i>	By checking this option you can set that an event is generated upon every 'Problem' evaluation of the trigger.
<i>Comments</i>	Text field used to provide more information about this trigger. May contain instructions for fixing specific problem, contact detail of responsible staff, etc.
<i>URL</i>	If not empty, the URL entered here is available as a link when clicking on the trigger name in <i>Monitoring</i> → <i>Triggers</i> . One macro may be used in the trigger URL field - {TRIGGER.ID}.
<i>Severity</i>	Set the required trigger severity by clicking the buttons.
<i>Enabled</i>	Unchecking this box will disable the trigger if required.

The **Dependencies** tab contains all the **dependencies** of the trigger.

Click on *Add* to add a new dependency.

Note:

You can also configure a trigger by opening an existing one, pressing the *Clone* button and then saving under a different name.

2 Trigger expression

Overview

The expressions used in triggers are very flexible. You can use them to create complex logical tests regarding monitored statistics.

A simple useful expression might look like:

{<server>:<key>.<function>(<parameter>)}<operator><constant>

1 Functions

Trigger functions allow to reference the collected values, current time and other factors.

A complete list of **supported functions** is available.

2 Function parameters

Most of numeric functions accept the number of seconds as a parameter.

You may use the prefix **#** to specify that a parameter has a different meaning:

FUNCTION CALL	MEANING
sum(600)	Sum of all values within 600 seconds
sum(#5)	Sum of the last 5 values

The function **last** uses a different meaning for values when prefixed with the hash mark - it makes it choose the n-th previous value, so given the values 3, 7, 2, 6, 5 (from most recent to least recent), **last(#2)** would return 7 and **last(#5)** would return 5.

A parameter must be given even for those functions which ignore it. Example: last(0)

Several functions support an additional, second `time_shift` parameter. This parameter allows to reference data from a period of time in the past. For example, **avg(1h,1d)** will return the average value for an hour one day ago.

Note:

Triggers only evaluate history information. If history is not available (especially relevant for time shift), trend information is not used, thus history must be kept for at least the period trigger functions expect it.

You can use the supported **unit symbols** in trigger expressions, for example '5m' (minutes) instead of '300' seconds or '1d' (day) instead of '86400' seconds. '1K' will stand for '1024' bytes.

3 Operators

The following operators are supported for triggers (**in descending priority of execution**):

PRIORITY	OPERATOR	DEFINITION
1	/	Division
2	*** Multiplication 3**	- Arithmetical minus
4	+	Arithmetical plus
5	<	Less than. The operator is defined as: $A < B \Leftrightarrow (A \leq B - 0.000001)$
6	>	More than. The operator is defined as: $A > B \Leftrightarrow (A \geq B + 0.000001)$
7	#	Not equal. The operator is defined as: $A \# B \Leftrightarrow (A \leq B - 0.000001) (A \geq B + 0.000001)$
8	=	Is equal. The operator is defined as: $A = B \Leftrightarrow (A > B - 0.000001) \& (A < B + 0.000001)$
9	&	Logical AND
10		Logical OR

4 Examples of triggers

Example 1

Processor load is too high on www.zabbix.com

```
{www.zabbix.com:system.cpu.load[all,avg1].last(0)}>5
```

'www.zabbix.com:system.cpu.load[all,avg1]' gives a short name of the monitored parameter. It specifies that the server is 'www.zabbix.com' and the key being monitored is 'system.cpu.load[all,avg1]'. By using the function 'last()', we are referring to the most recent value. Finally, '>5' means that the trigger is in the PROBLEM state whenever the most recent processor load measurement from www.zabbix.com is greater than 5.

Example 2

www.zabbix.com is overloaded

```
{www.zabbix.com:system.cpu.load[all,avg1].last(0)}>5|{www.zabbix.com:system.cpu.load[all,avg1].min(10m)}>2
```

The expression is true when either the current processor load is more than 5 or the processor load was more than 2 during last 10 minutes.

Example 3

/etc/passwd has been changed

Use of function diff:

```
{www.zabbix.com:vfs.file.cksum[/etc/passwd].diff(0)}=1
```

The expression is true when the previous value of checksum of /etc/passwd differs from the most recent one.

Similar expressions could be useful to monitor changes in important files, such as /etc/passwd, /etc/inetd.conf, /kernel, etc.

Example 4

Someone is downloading a large file from the Internet

Use of function min:

```
{www.zabbix.com:net.if.in[eth0,bytes].min(5m)}>100K
```

The expression is true when number of received bytes on eth0 is more than 100 KB within last 5 minutes.

Example 5

Both nodes of clustered SMTP server are down

Note use of two different hosts in one expression:

```
{smtp1.zabbix.com:net.tcp.service[smtp].last(0)}=0&{smtp2.zabbix.com:net.tcp.service[smtp].last(0)}=0
```

The expression is true when both SMTP servers are down on both smtp1.zabbix.com and smtp2.zabbix.com.

Example 6

Zabbix agent needs to be upgraded

Use of function str():

```
{zabbix.zabbix.com:agent.version.str("beta8")}=1
```

The expression is true if Zabbix agent has version beta8 (presumably 1.0beta8).

Example 7

Server is unreachable

```
{zabbix.zabbix.com:icmping.count(30m,0)}>5
```

The expression is true if host "zabbix.zabbix.com" is unreachable more than 5 times in the last 30 minutes.

Example 8

No heartbeats within last 3 minutes

Use of function nodata():

```
{zabbix.zabbix.com:tick.nodata(3m)}=1
```

To make use of this trigger, 'tick' must be defined as a Zabbix **trapper** item. The host should periodically send data for this item using zabbix_sender. If no data is received within 180 seconds, the trigger value becomes PROBLEM.

Note that 'nodata' can be used for any item type.

Example 9

CPU activity at night time

Use of function time():

```
{zabbix:system.cpu.load[all,avg1].min(5m)}>2&{zabbix:system.cpu.load[all,avg1].time(0)}>000000&{zabbix:sys
```

The trigger may change its status to true, only at night (00:00-06:00) time.

Example 10

Check if client local time is in sync with Zabbix server time

Use of function `fuzzytime()`:

```
{MySQL_DB:system.localtime.fuzzytime(10)}=0
```

The trigger will change to the problem state in case when local time on server MySQL_DB and Zabbix server differs by more than 10 seconds.

Example 11

Comparing average load today with average load of the same time yesterday (using a second `time_shift` parameter).

```
{server:system.cpu.load.avg(1h)}/{server:system.cpu.load.avg(1h,1d)}>2
```

This expression will fire if the average load for the last hour tops the average load of the same hour yesterday more than two times.

5 Hysteresis

Sometimes a trigger must have different conditions for different states. For example, we would like to define a trigger which would become PROBLEM when server room temperature is higher than 20C while it should stay in the state until temperature will not become lower than 15C.

In order to do this, we define the following trigger:

Example 1

Temperature in server room is too high

```
({TRIGGER.VALUE}=0&{server:temp.last(0)}>20) |  
({TRIGGER.VALUE}=1&{server:temp.last(0)}>15)
```

Note the use of a macro `{TRIGGER.VALUE}`. The macro returns current trigger value.

Example 2

Free disk space is too low

Problem: it is less than 10GB for last 5 minutes

Recovery: it is more than 40GB for last 10 minutes

```
({TRIGGER.VALUE}=0&{server:vfs.fs.size[/,free].max(5m)}<10G) |  
({TRIGGER.VALUE}=1&{server:vfs.fs.size[/,free].min(10m)}<40G)
```

Note use of macro `{TRIGGER.VALUE}`. The macro returns current trigger value.

3 Trigger dependencies

Overview

Sometimes the availability of one host depends on another. A server that is behind some router will become unreachable if the router goes down. With triggers configured for both, you might get notifications about two hosts down - while only the router was the guilty party.

This is where some dependency between hosts might be useful. With dependency set notifications of the dependants could be withheld and only the notification for the root problem sent.

While Zabbix does not support dependencies between hosts directly, they may be defined with another, more flexible method - trigger dependencies. A trigger may have one or more triggers it depends on.

So in our simple example we open the server trigger configuration form and set that it depends on the respective trigger of the router. With such dependency the server trigger will not change state as long as the trigger it depends on is in 'Problem' state - and thus no dependant actions will be taken and no notifications sent.

If both the server and the router are down and dependency is there, Zabbix will not execute actions for the dependent trigger.

Also:

- Trigger dependency may be added from any host trigger to any other host trigger, as long as it wouldn't result in a circular dependency.

- Trigger dependency may be added from a template to a template. If a trigger from template A depends on a trigger from template B, template A may only be linked to a host (or another template) together with template B, but template B may be linked to a host (or another template) alone.
- Trigger dependency may be added from template trigger to a host trigger. In this case, linking such a template to a host will create a host trigger that depends on the same trigger template trigger was depending on. This allows to, for example, have a template where some triggers depend on router (host) triggers. All hosts linked to this template will depend on that specific router.
- Trigger dependency from a host trigger to a template trigger may not be added.

Configuration

To define a dependency, open the trigger **configuration form**. Click on *Add* next to 'New dependency' and select one or more triggers that our trigger will depend on.

The screenshot shows the 'The trigger depends on' section of the Zabbix configuration form. It features a list of available triggers, with 'Host New host is unreachable' selected. Below the list is a 'Delete selected' button. In the 'New dependency' section, there is an 'Add' button.

Click Save. Now the trigger has an indication of its dependency in the list.

The screenshot shows a trigger entry in the Zabbix interface. The trigger name is 'Template Linux: Host {HOSTNAME} is unreachable'. Below the name, it says 'Depends on : New host : Host New host is unreachable'.

Example of several dependencies

For example, a Host is behind a Router2 and the Router2 is behind a Router1.

Zabbix - Router1 - Router2 - Host

If Router1 is down, then obviously Host and Router2 are also unreachable yet we don't want to receive three notifications about Host, Router1 and Router2 all being down.

So in this case we define two dependencies:

'Host is down' trigger depends on 'Router2 is down' trigger
 'Router2 is down' trigger depends on 'Router1 is down' trigger

Before changing the status of the 'Host is down' trigger, Zabbix will check for corresponding trigger dependencies. If found, and one of those triggers is in 'Problem' state, then the trigger status will not be changed and thus actions will not be executed and notifications will not be sent.

Zabbix performs this check recursively. If Router1 or Router2 is unreachable, the Host trigger won't be updated.

4 Trigger severity

Trigger severity defines how important a trigger is. Zabbix supports the following trigger severities:

SEVERITY	DEFINITION	COLOUR
Not classified	Unknown severity.	Grey
Information	For information purposes.	Light green
Warning	Be warned.	Yellow
Average	Average problem.	Orange
High	Something important has happened.	Red
Disaster	Disaster. Financial losses, etc.	Bright red

The severities are used for:

- visual representation of triggers. Different colours for different severities.
- audio in global alarms. Different audio for different severities.
- user media. Different media (notification channel) for different severities. For example, SMS - high severity, email - other.
- limiting actions by conditions against trigger severities

It is possible to **customise trigger severity names and colours**.

5 Customising trigger severities

Trigger severity names and colours for severity related GUI elements can be configured in *Administration* → *General* → *Trigger severities*. Colours are shared among all GUI themes.

Translating customised severity names

Attention:

If Zabbix frontend translations are used, custom severity names will override translated names by default.

Default trigger severity names are available for translation in all locales. If a severity name is changed, custom name is used in all locales and additional manual translation is needed.

Custom severity name translation procedure:

- set required custom severity name, for example 'Important'
- edit `<frontend_dir>/locale/<required_locale>/LC_MESSAGES/frontend.po`
- add 2 lines:

```
msgid "Important"
msgstr "<translation string>"
```

and save file.

- create .mo files as described in `<frontend_dir>/locale/README`

Here **msgid** should match the new custom severity name and **msgstr** should be the translation for it in the specific language.

This procedure should be performed after each severity name change.

6 Unit symbols

Overview

Having to use some large numbers, for example '86400' to represent the number of seconds in one day, is both difficult and error-prone. This is why you can use some appropriate unit symbols (or suffixes) to simplify Zabbix trigger expressions and item keys.

Instead of '86400' you can simply enter '1d'. Suffixes function as multipliers.

Trigger expressions

Time and memory size suffixes are supported in trigger **expression** constants and function parameters.

For time you can use:

- **s** - seconds (when used, works the same as the raw value)
- **m** - minutes
- **h** - hours
- **d** - days
- **w** - weeks

Time suffixes are also supported in parameters of the **zabbix[queue,<from>,<to>]** **internal item** and the last parameter of **aggregate checks**.

For memory size you can use:

- **K** - kilobyte
- **M** - megabyte
- **G** - gigabyte
- **T** - terabyte

Other uses

Unit symbols are also used for a human-readable representation of data in the frontend.

In both Zabbix server and frontend these symbols are supported:

- **K** - kilo
- **M** - mega
- **G** - giga
- **T** - tera

When item values in B, Bps are displayed in the frontend, base 2 is applied (1K = 1024). Otherwise a base of 10 is used (1K = 1000).

Additionally the frontend also supports the display of:

- **P** - peta
- **E** - exa
- **Z** - zetta
- **Y** - yotta

Usage examples

By using some appropriate suffixes you can write trigger expressions that are easier to understand and maintain, for example these expressions:

```
{host:zabbix[proxy,zabbix_proxy,lastaccess]}>120
{host:system.uptime[].last(0)}<86400
{host:system.cpu.load.avg(600)}<10
{host:vm.memory.size[available].last(0)}<20971520
```

could be changed to:

```
{host:zabbix[proxy,zabbix_proxy,lastaccess]}>2m
{host:system.uptime.last(0)}<1d
{host:system.cpu.load.avg(10m)}<10
{host:vm.memory.size[available].last(0)}<20M
```

7 Mass update

Overview

With mass update you may change some attribute for a number of triggers at once, saving you the need to open each individual trigger for editing.

Using mass update

To mass-update some triggers, do the following:

- Mark the checkboxes of the triggers to update in the list
- Select *Mass update* below the list and click on *Go*
- Mark the checkboxes of the attributes to update
- Specify new values for the attributes and click on *Update*

Mass update

Severity ☒

Not classified

Information

Warning

Average

High

Disaster

Replace dependencies ☒

Name

New host: Zabbix agent on {HOST.NAME} is unreachable for 5 minutes

Add

Update

Cancel

Replace dependencies will replace existing trigger dependencies (if any) with the ones specified in mass update.

4 Events

Overview

Events in Zabbix are generated by three sources:

- triggers - whenever a trigger changes its status

- discovery - upon detection of hosts or services
- auto registration - when active agents are auto-registered by server

Events are time-stamped and can be the basis of actions such as sending notification e-mail etc.

To view details of events in the frontend, go to *Monitoring | Events*. There you can click on the event date and time to view details of an event.

More information is available on [each event source](#).

1 Event sources

1.1 Trigger events

Change of trigger status is the most frequent and most important source of events.

Each time the trigger changes its state, an event is generated. The event contains details of the trigger state's change - when did it happen and what the new state is.

1.2 Discovery events

Zabbix periodically scans the IP ranges defined in network discovery rules. Frequency of the check is configurable for each rule individually. Once a host or a service is discovered, a discovery event (or several events) are generated.

Zabbix generates the following events:

Event	When generated
Service Up	Every time Zabbix detects active service.
Service Down	Every time Zabbix cannot detect service.
Host Up	If at least one of the services is UP for the IP.
Host Down	If all services are not responding.
Service Discovered	If the service is back after downtime or discovered for the first time.
Service Lost	If the service is lost after being up.
Host Discovered	If host is back after downtime or discovered for the first time.
Host Lost	If host is lost after being up.

1.3 Active agent auto-discovery events

Active agent auto-registration creates events in Zabbix.

If configured, active agent auto-registration can happen when a previously unknown active agent asks for checks. The server adds a new auto-registered host, using the received IP address and port of the agent.

For more information, see the [active agent auto-registration](#) page.

5 Visualisation

1 Graphs

Overview

With lots of data flowing into Zabbix, it becomes much easier for the users if they can look at a visual representation of what is going on rather than only numbers.

This is where graphs come in. Graphs allow to grasp the data flow at a glance, correlate problems, discover when something started or make a presentation of when something might turn into a problem.

Zabbix provides users with built-in [simple graphs](#) as well as with the possibility to create more complex [customised graphs](#).


1 Simple graphs

Overview

Simple graphs are provided for the visualization of data gathered by items.

No configuration effort is required on the user part to view simple graphs. They are freely made available by Zabbix.

Just go to *Monitoring* → *Latest data* and click on the Graph link for the respective item and a graph will be displayed.



Time period selector

Take note of the time period selector above the graph. It allows you to select the desired time period easily.

The slider within the selector can be dragged back and forth, as well as resized, effectively changing the time period displayed. Links on the left hand side allow to choose some often-used predefined periods (above the slider area) and move them back and forth in time (below the slider area). The dates on the right hand side actually work as links, popping up a calendar and allowing to set a specific start/end time.

The **fixed/dynamic** link in the lower right hand corner has the following effects:

- controls whether the time period is kept constant when you change the start/end time in the calendar popup.
- when *fixed*, time moving controls (« 6m 1m 7d 1d 12h 1h | 1h 12h 1d 7d 1m 6m ») will move the slider, while not changing its size, whereas when *dynamic*, the control used will enlarge the slider in the respective direction.
- when *fixed*, pressing the larger < and > buttons will move the slider, while not changing its size, whereas when *dynamic*, < and > will enlarge the slider in the respective direction. The slider will move by the amount of its size, so, for example, if it is one month, it will move by a month; whereas the slider will enlarge by 1 day.

Another way of controlling the displayed time is to highlight an area in the graph with the left mouse button. The graph will zoom into the highlighted area once you release the left mouse button.

Note:


Simple graphs are provided for all numeric items. For textual items, a link to History is available in *Monitoring* → *Latest data*.

Recent data vs longer periods

For very recent data a **single** line is drawn connecting each received value. The single line is drawn as long as there is at least one horizontal pixel available for one value.

For data that show a longer period **three lines** are drawn - a dark green one shows the average, while a light pink and a light green line shows the maximum and minimum values at that point in time. The space between the highs and the lows is filled with yellow background.

Working time (working days) is displayed in graphs as a white background, while non-working time is displayed in grey (with the *Original blue* default frontend theme).



Working time is always displayed in simple graphs, whereas displaying it in **custom graphs** is a user preference.

Working time is not displayed if the graph shows more than 3 months.

Generating from history/trends

Graphs can be drawn based on either item **history** or **trends**. A grey caption at the bottom right of a graph indicates where the data come from.

Two factors influence whether history of trends is used:

- longevity of item history. For example, item history can be kept for 14 days. In that case, any data older than the fourteen days will be coming from trends.
- data congestion in the graph. If the amount of seconds to display in a horizontal graph pixel exceeds 3600/16, trend data are displayed (even if item history is still available for the same period).

Switching to raw values

A dropdown on the upper right above the graph allows to switch from the simple graph to the *Values/500 latest values* listings. This can be useful for viewing the numeric values making up the graph.

The values represented here are raw, i.e. no units or postprocessing of values is used. Value mapping, however, is applied.

2 Custom graphs

Overview

Custom graphs, as the name suggests, offer customisation capabilities.

While simple graphs are good for viewing data of a single item, they do not offer configuration capabilities.

Thus, if you want to change graph style or the way lines are displayed or compare several items, for example incoming and outgoing traffic in a single graph, you need a custom graph.

Custom graphs are configured manually.

They can be created for a host or several hosts or for a single template.

Configuring custom graphs

To create a custom graph, do the following:

- Go to *Configuration* → *Hosts (or Templates)*
- Click on *Graphs* in the row next to the desired host or template
- In the Graphs screen click on *Create graph*
- Edit graph attributes

GraphPreview

NameNetwork utilization

Width900

Height200

Graph typeNormal

Show legend☒

Show working time☒

Show triggers☒

Percentile line (left)☐

Percentile line (right)☐

Y axis MIN valueCalculated

Y axis MAX valueCalculated

Items

	Name	Function	Draw style	Y axis side	Colour	Action
1:	Zabbix server: Outgoing network traffic on eth0	avg	Filled region	Right	00C800	Remove
2:	Zabbix server: Incoming network traffic on eth0	avg	Bold line	Right	C80000	Remove

Add

Save

Cancel

Graph attributes:

Parameter	Description
<i>Name</i>	Unique graph name.
<i>Width</i>	Graph width in pixels (for preview and pie/exploded graphs only).
<i>Height</i>	Graph height in pixels.
<i>Graph type</i>	Graph type: Normal - normal graph, values displayed as lines Stacked - stacked graph, filled areas displayed Pie - pie graph Exploded - "exploded" pie graph, portions displayed as "cut out" of the pie
<i>Show legend</i>	Checking this box will set to display the graph legend.
<i>Show working time</i>	If selected, non-working hours will be shown with gray background.
<i>Show triggers</i>	Not available for pie and exploded pie graphs. If selected, simple triggers will be displayed as red lines. Not available for pie and exploded pie graphs.
<i>Percentile line (left)</i>	Display percentile for left Y axis. If, for example, 95% percentile is set, then the percentile line will be at the level where 95 per cent of the values fall under. Displayed as a bright green line. Only available for normal graphs.
<i>Percentile line (right)</i>	Display percentile for right Y axis. If, for example, 95% percentile is set, then the percentile line will be at the level where 95 per cent of the values fall under. Displayed as a bright red line. Only available for normal graphs.
<i>Y axis MIN value</i>	Minimum value of Y axis: Calculated - Y axis minimum value will be automatically calculated Fixed - fixed minimum value for Y axis. Not available for pie and exploded pie graphs.
<i>Y axis MAX value</i>	Item - last value of the selected item will be the minimum value Maximum value of Y axis: Calculated - Y axis maximum value will be automatically calculated Fixed - fixed maximum value for Y axis. Not available for pie and exploded pie graphs.
<i>3D view</i>	Item - last value of the selected item will be the maximum value Enable 3D style. For pie and exploded pie graphs only.
<i>Items</i>	Items, data of which are to be displayed in this graph.

Configuring graph items

To add items, data of which are to be displayed in the graph, click on *Add* in the *Items* block, select items and then set attributes


for the way item data will be displayed.

Item display attributes:


Parameter	Description
<i>Sort order (0→100)</i>	Draw order. 0 will be processed first. Can be used to draw lines or regions behind (or in front of) another. You can drag and drop items by the arrow in the beginning of line to set the sort order or which item is displayed in front of the other.
<i>Name</i>	Name of item, data of which will be displayed.
<i>Type</i>	Type (only available for pie and exploded pie graphs): Simple - value of the item is represented proportionally on the pie Graph sum - value of the item represents the whole pie Note that colouring of the "graph sum" item will only be visible to the extent that it is not taken up by "proportional" items.
<i>Function</i>	What values will be displayed when more than one value exists for an item: all - all (minimum, average and maximum) min - minimum only avg - average only max - maximum only
<i>Draw style</i>	Draw style (only available for normal graphs; for stacked graphs filled region is always used): Line - draw lines Filled region - draw filled region Bold line - draw bold lines Dot - draw dots Dashed line - draw dashed line
<i>Y axis side</i>	Which Y axis side the element is assigned to.
<i>Colour</i>	RGB colour in HEX notation.

Graph preview

In the *Preview* tab, a preview of the graph is displayed so you can immediately see what you are creating.



Note that the preview will not show any data for template items.



In this example, pay attention to the dashed bold line displaying the trigger level and the trigger information displayed in the legend.

Note:

3 triggers is the hard-coded limit for the number of triggers displayed in the legend.
If graph height is set as less than 120 pixels, no trigger will be displayed in the legend.

2 Network maps

Overview

If you have a network to look after, you may want to have an overview of your infrastructure somewhere. For that purpose you can create maps in Zabbix - of networks and of anything you like.

Proceed to [configuring a network map](#).

1 Configuring a network map

Overview

Configuring a map in Zabbix requires that you first create a map by defining its general parameters and then you start filling the actual map with elements and their links.

You can populate the map with elements that are a host, a host group, a trigger, an image or another map.

Icons are used to represent map elements. You can define the information that will be displayed with the icons and set that recent problems are displayed in a special way. You can link the icons and define information to be displayed on the links.

Maps that are ready can be viewed in *Monitoring* → *Maps*. In the monitoring view you can click on the icons and take advantage of the links to some scripts and URLs.

You can add custom URLs to be accessible by clicking on the icons. Thus you may link a host icon to host properties or a map icon to another map.

Creating a map

To create a map, do the following:

- Go to *Configuration* → *Maps*
- Click on *Create map*
- Edit general map attributes

Map

Name
Network

Width
980

Height
800

Background image
No image

Automatic icon mapping
<manual>
[show icon mappings](#)

Icon highlight
☒

Mark elements on trigger status change
☐

Expand single problem
☒

Advanced labels
☒

Host group label type
Label

Host label type
Label

Trigger label type
Status only

Map label type
Label

Image label type
Nothing

Icon label location
Right

Problem display
All

URLs

Name	URL	Element
Latest data	http://192.168.3.2/zabbix/latest.php	Host Remove
Status of triggers	http://192.168.3.2/zabbix/tr_status.php	Trigger Remove

Add

General map attributes:

Parameter	Description
<i>Name</i>	Unique map name.
<i>Width</i>	Map width in pixels.
<i>Height</i>	Map height in pixels.
<i>Background image</i>	Use background image: No image - no background image (white background) Image - selected image to be used as a background image. No scaling is performed. You may use a geographical map or any other image to enhance your map.
<i>Automatic icon mapping</i>	You can set to use an automatic icon mapping, configured in <i>Administration</i> → <i>General</i> → <i>Icon mapping</i> . Icon mapping allows to map certain icons against certain host inventory fields.
<i>Icon highlighting</i>	If you check this box, icons will receive highlighting. Elements with an active trigger will receive a round background, in the same colour as the highest severity trigger. Moreover, a thick green line will be displayed around the circle, if all problems are acknowledged. Elements with "disabled" or "in maintenance" status will get a square background, gray and orange respectively.
<i>Mark elements on trigger status change</i>	A recent change of trigger status (recent problem or resolution) will be highlighted with markers (inward-pointing red triangles) on the three sides of the element icon that are free of the label. Markers are displayed for 30 minutes.
<i>Expand single problem</i>	If a map element (host, host group or another map) has one single problem, this option controls whether the problem (trigger) name is displayed, or problem count. If marked, problem name is used.
<i>Advanced labels</i>	If you check this box you will be able to define separate label types for separate element types.
<i>Icon label type</i>	Label type used for icons: Label - icon label IP address - IP address Element name - element name (for example, host name) Status only - status only (OK or PROBLEM) Nothing - no labels are displayed

Parameter	Description
<i>Icon label location</i>	Label location in relation to the icon: Bottom - beneath the icon Left - to the left Right - to the right Top - above the icon
<i>Problem display</i>	Display problem count as: All - full problem count will be displayed Separated - unacknowledged problem count will be displayed separated as a number of the total problem count Unacknowledged only - only the unacknowledged problem count will be displayed
<i>URLs</i>	URLs for each element type can be defined (with a label). These will be displayed as links when a user clicks on the element in the monitoring section. Macros that can be used in map URLs: {MAP.ID}, {HOSTGROUP.ID}, {HOST.ID}, {TRIGGER.ID}


When you save this, you have created an empty map with a name, dimensions and certain preferences. Now you need to add some elements. For that, click on the map name in the list to open the editable area.

Adding elements

To add an element, click on the "+" next to Icon. The new element will appear at the top left corner of the map. Drag and drop it wherever you like.

Note that with the Grid option "On", elements will always align to the grid (you can pick various grid sizes from the dropdown, also hide/show the grid). If you want to put elements anywhere without alignment, turn the option to "Off". (Random elements can later again be aligned to the grid with the *Align icons* button.)

Now that you have some elements in place, you may want to start differentiating them by giving names etc. By clicking on the element, a form is displayed and you can set the element type, give a name, choose a different icon etc.



Map element attributes:

Parameter	Description
<i>Type</i>	Type of the element: Host - icon representing status of all triggers of the selected host Map - icon representing status of all elements of a map Trigger - icon representing status of a single trigger Host group - icon representing status of all triggers of all hosts belonging to the selected group Image - an icon, not linked to any resource
<i>Label</i>	Icon label, any string. Macros and multi-line strings can be used in labels.
<i>Label location</i>	Label location in relation to the icon: Default - map's default label location Bottom - beneath the icon Left - to the left Right - to the right Top - above the icon
<i>Host</i>	Select the host, if the element type is 'Host'.
<i>Map</i>	Select the map, if the element type is 'Map'.
<i>Trigger</i>	Select the trigger, if the element type is 'Trigger'.
<i>Host group</i>	Select the host group, if the element type is 'Host group'.
<i>Icon (default)</i>	Icon to be used.
<i>Automatic icon selection</i>	In this case an icon mapping will be used to determine which icon to display.
<i>Icons</i>	You can choose to display different icons for the element in these cases: default, problem, maintenance, disabled.
<i>Coordinate X</i>	X coordinate of the map element.
<i>Coordinate Y</i>	Y coordinate of the map element.
<i>URLs</i>	Element-specific URLs can be set for the element. These will be displayed as links when a user clicks on the element in the monitoring section. If the element has its own URLs and there are map level URLs for its type defined, they will be combined in the same menu. Macros that can be used in map URLs: {MAP.ID}, {HOSTGROUP.ID}, {HOST.ID}, {TRIGGER.ID}

Attention:


Added elements are not automatically saved. If you navigate away from the page, all changes may be lost. Therefore it is a good idea to click on the **Save** button in the top right corner. Once clicked, the changes are saved regardless of what you choose in the following popup.
Selected grid options are also saved with each map.

Selecting elements

To select elements, select one and then hold down *Ctrl* (or *Shift*) to select the others.

You can also select multiple elements by dragging a rectangle in the editable area and selecting all elements in it (option available since Zabbix 2.0).


Once you select more than one element, the element property form shifts to the mass-update mode so you can change attributes of selected elements in one go. To do so, mark the attribute using the checkbox and enter a new value for it. You may use macros here (such as, say, {HOSTNAME} for the element label).



Linking elements

Once you have put some elements on the map, it is time to start linking them. To link two elements you must first select them. With the elements selected, click on the "+" next to Link.

With a link created, the single element form now contains an additional *Edit element links* section. Click on *Edit* before the link to edit its attributes.



Link attributes:

Parameter	Description
<i>Label</i>	Label that will be rendered on top of the link. The <code>{host:key.func(param)}</code> macro is supported in this field, but only with avg, last, min and max trigger functions, with seconds as parameter.
<i>Connect to</i> <i>Type (OK)</i>	The element that the link connects to. Default link style: Line - single line Bold line - bold line Dot - dots Dashed line - dashed line
<i>Colour (OK)</i>	Default link colour.
<i>Link indicators</i>	List of triggers linked to the link. In case a trigger has status PROBLEM, its style is applied to the link.

2 Link indicators

Overview

You can assign some triggers to a **link** between elements in a network map. When these triggers go into a problem state, the link can reflect that.

When you configure a link, you set the default link type and color. When you assign triggers to a link, you can assign different link types and colors with these triggers.

Should any of these triggers go into a problem state, their link style and color will be displayed on the link. So maybe your default link was a green line. Now, with the trigger in problem state, your link may become bold red (if you have defined it so).

Configuration

To assign triggers as link indicators, do the following:

- select a map element

- click on *Edit* in the *Edit element links* section before the appropriate link
- click on *Add* in the *Link indicators* block and select one or more triggers

Edit map element

Type: Host

Label: New element

Label location: Bottom

Host: New host

Icons:

Automatic icon selection: ☐

Default: Server_96 Problem: Server_128 Maintenance: Server_24 Disabled: Default

Coordinate X: 91 Coordinate Y: 388

URLs:

Name: URL: Remove

Add

Apply Remove Close

Edit element links

Element type	Element name	Link indicators
Edit Host	Zabbix server	

Label: 100MBps

Connect to: Zabbix server

Type (OK): Bold line

Colour (OK): 008000

Link indicators:

Triggers	Type	Colour	
New host:Agent is unavailable on New host	Line	DD0000	Remove
New host:Low free disk space on New host volume /	Line	DD0000	Remove

Add


Apply Remove Close

Added triggers can be seen in the *Link indicators* list.

You can set the link type and color for each trigger directly from the list. When done, click on *Apply*, close the form and save the map.

Display

In *Monitoring* → *Maps* the respective color will be displayed on the link if the trigger goes into a problem state.



Note:

If multiple triggers go into a problem state, the one with the highest severity will determine the link style and color. If multiple triggers with the same severity are assigned to the same map link, the one with the lowest ID takes precedence.

3 Screens

Overview

On Zabbix screens you can group information from various sources for a quick overview on a single screen. Building the screens is quite easy and intuitive.

Essentially a screen is a table. You choose how many cells per table and what elements to display in the cells. The following elements can be displayed:

- simple graphs
- user-defined custom graphs
- maps
- other screens
- plain text information
- server information (overview)
- hosts information (overview)
- trigger information (overview)
- status of triggers (by host or hostgroup)
- system status
- data overview
- clock
- history of events
- history of actions
- URL (data taken from another location)

Attention:

Browsers might not load an HTTP page included in a screen (using URL element), if Zabbix frontend is accessed over HTTPS.


Screens that are ready can be viewed in *Monitoring* → *Screens*. They can also be added to the favourites section of the Dashboard.

To configure a screen you must first create it by defining its general properties and then add individual elements in the cells.

Creating a screen

To create a screen, do the following:

- Go to *Configuration* → *Screens*
- Click on *Create Screen*
- Edit general screen attributes




Give your screen a unique name and set the number of columns (vertical cells) and rows (horizontal cells). Click Save.

Now you can click on the screen name in the list to be able to add elements.

Adding elements

On a new screen you probably only see links named *Change*. Clicking those links opens a form whereby you set what to display in each cell.

On an existing screen you click on the existing elements to open the form whereby you set what to display.



Screen element attributes:

Parameter	Description
<i>Resource</i>	Information displayed in the cell: Clock - digital or analog clock displaying current server or local time. <i>Note:</i> To display host time, use the <code>system.localtime[local]</code> item. This item must exist on the host. Data overview - latest data for a group of hosts Graph - single custom graph History of actions - history of recent actions History of events - latest events Hosts info - high level host related information Map - single map Plain text - plain text data Screen - screen (one screen may contain other screens inside) Server info - server high-level information Simple graph - single simple graph Status of hostgroup triggers - status of triggers filtered by the hostgroup Status of host triggers - status of triggers filtered by the host System status - displays system status (similar to the Dashboard) Triggers info - high level trigger related information Triggers overview - status of triggers for a host group URL - include content from an external resource
<i>Horizontal align</i>	Possible values: Center Left Right
<i>Vertical align</i>	Possible values: Middle Top Bottom

Parameter	Description
<i>Column span</i>	Extend cell to a number of columns, same way as HTML column spanning works.
<i>Row span</i>	Extend cell to a number of rows, same way as HTML row spanning works.

Take note of the '+' and '-' controls on each side of the table.

Clicking on '+' above the table will add a column. Clicking on '-' beneath the table will remove a column.

Clicking on '+' on the left side of the table will add a row. Clicking on '-' on the right side of the table will remove a row.

Attention:

If graph height is set as less than 120 pixels, no trigger will be displayed in the legend.

Dynamic elements

For some of the elements there is an extra option called *Dynamic item*. Checking this box at first does not seem to change anything.

However, once you go to *Monitoring* → *Screens*, you may realize that now you have extra dropdowns there for selecting the host. Thus you have a screen where some elements display the same information while others display information depending on the currently selected host.

The benefit of this is that you do not need to create extra screens just because you want to see the same graphs containing data from various hosts.

Dynamic item option is available for several screen elements:

- Graphs (custom graphs)
- Simple graphs
- Plain text

Note:

Clicking on a dynamic graph opens it in full view; although with custom graphs that is currently supported with the default host only (i.e. with host 'not selected' in the dropdown). When selecting another host in the dropdown, the dynamic graph is created using item data of that host and the resulting graph is not clickable.

4 Slide shows

Overview

In a slide show you can configure that a number of **screens** are displayed one after another at set intervals.

Sometimes you might want to switch between some configured screens. While that can be done manually, doing that more than once or twice may become very tedious. This is where the slide show function comes to rescue.

Configuration

To create a slide show, do the following:

- Go to *Configuration* → *Slide shows*
- Click on *Create slide show*
- Edit slide show attributes

Slide

Name

Zabbix administrators

Default delay (in seconds)

30

Slides

	Screen	Delay	Action
⬆ 1	Zabbix server	default	Remove
⬆ 2	New host	15	Remove
Add			

PARAMETER	Description
<i>Name</i>	Name of the slide show.
<i>Default delay (in seconds)</i>	How long one screen is displayed by default, before rotating to the next, in seconds.
<i>Slides</i>	List of screens to be rotated. Click on <i>Add</i> to select screens. The <i>Up/Down</i> arrow before the screen allows to drag a screen up and down in the sort order of display. If you want to display only, say, a single graph in the slide show, create a screen containing just that one graph.
<i>Screen</i>	Screen name.
<i>Delay</i>	A custom value for how long the screen will be displayed, in seconds. If set to 0, the <i>Default delay</i> value will be used.
<i>Action</i>	Click on <i>Remove</i> to remove a screen from the slide show.

The slide show in this example consists of two screens which will be displayed in the following order:

Zabbix server ⇒ Displayed for 30 seconds ⇒ New host ⇒ Displayed for 15 seconds ⇒ Zabbix server ⇒ Displayed for 30 seconds ⇒ New host ⇒ ...

Display

Slide shows that are ready can be viewed in *Monitoring* → *Screens* and then choosing *Slide shows* from the dropdown.

With the Menu option next to the dropdown, you can accelerate or slow down the display by choosing a slide delay multiplier:

Refresh time multiplier
x0.25
x0.5
x1
x1.5
x2
x3
x4
x5

Attention:

If a delay ends up as being less than 5 seconds (either by having entered a delay less than 5 seconds or by using the slide delay multiplier), a 5-second minimum delay will be used.

6 Templates

Overview

A template is a set of entities that can be conveniently applied to multiple hosts.

The entities may be:

- items
- triggers
- graphs
- applications
- screens (*since Zabbix 2.0*)
- low-level discovery rules (*since Zabbix 2.0*)

As many hosts in real life are identical or fairly similar so it naturally follows that the set of entities (items, triggers, graphs,...) you have created for one host, may be useful for many. Of course, you could copy them to each new host, but that would be a lot of manual work. Instead, with templates you can copy them to one template and then apply the template to as many hosts as needed.

When a template is linked to a host, all entities (items, triggers, graphs,...) of the template are added to the host. Templates are assigned to each individual host directly (and not to a host group).

Templates are often used to group entities for particular services or applications (like Apache, MySQL, PostgreSQL, Postfix...) and then applied to hosts running those services.

Another benefit of using templates is when something has to be changed for all the hosts. Changing something on the template level once will propagate the change to all the linked hosts.

Thus, the use of templates is an excellent way of reducing one's workload and streamlining the Zabbix configuration.

Proceed to [creating and configuring a template](#).

7 Notifications upon events

Overview

Assuming that we have configured some items and triggers and now are getting some events happening as a result of triggers changing state, it is time to consider some actions.

To begin with, we would not want to stare at the triggers or events list all the time. It would be much better to receive notification if something significant (such as a problem) has happened. Also, when problems occur, we would like to see that all the people concerned are informed.

That is why sending notifications is one of the primary actions offered by Zabbix. Who and when should be notified upon a certain event can be defined.

To be able to send and receive notifications from Zabbix you have to:

- [define some media](#)
- [configure an action](#) that sends a message to one of the defined media

Actions consist of *conditions* and *operations*. Basically, when conditions are met, operations are carried out. The two principal operations are sending a message (notification) and executing a remote command.

For discovery and auto-registration created events, some additional operations are available. Those include adding or removing a host, linking a template etc.

1 Media types

Overview

Media are the delivery channels used for sending notifications and alerts in Zabbix.

You can configure several media types:

- [E-mail](#)
- [SMS](#)
- [Jabber](#)
- [Ez Texting](#)
- [Custom alertscripts](#)

1 E-mail

Overview

To configure e-mail as the delivery channel for messages, you need to configure e-mail as the media type and assign specific addresses to users.

Configuration

To configure e-mail as the media type:

- Go to *Administration*→*Media types*

- Click on *Create media type* (or click on *E-mail* in the list of pre-defined media types).

Media type attributes:

Parameter	Description
<i>Description</i>	Name of the media type.
<i>Type</i>	Select <i>Email</i> as the type.
<i>SMTP server</i>	Set an SMTP server to handle outgoing messages.
<i>SMTP helo</i>	Set a correct SMTP helo value, normally a domain name.
<i>SMTP email</i>	The address entered here will be used as the From address for the messages sent.

User media

To assign a specific address to the user:

- Go to *Administration→Users*
- Open the user properties form
- In Media tab, click on *Add*

User media attributes:

Parameter	Description
<i>Type</i>	Select <i>Email</i> as the type.
<i>Send to</i>	Specify the e-mail address to send the messages to.
<i>When active</i>	You can limit the time when messages are sent, for example, the working days only (1-5,09:00-18:00). See the Time period specification page for description of the format.

Parameter	Description
<i>Use if severity</i>	Mark the checkboxes of trigger severities that you want to receive notifications for.
<i>Status</i>	Status of the user media. Enabled - is in use. Disabled - is not being used.

2 SMS

Overview

Zabbix supports the sending of SMS messages using a serial GSM modem connected to Zabbix server's serial port.

Make sure that:

- The speed of the serial device (normally /dev/ttyS0 under Linux) matches that of the GSM modem. Zabbix does not set the speed of the serial link. It uses default settings.
- The 'zabbix' user has read/write access to the serial device. Run the command `ls -l /dev/ttyS0` to see current permissions of the serial device.
- The GSM modem has PIN entered and it preserves it after power reset. Alternatively you may disable PIN on the SIM card. PIN can be entered by issuing command `AT+CPIN="NNNN"` (NNNN is your PIN number, the quotes must be present) in a terminal software, such as Unix minicom or Windows HyperTerminal.

Zabbix has been tested with these GSM modems:

- Siemens MC35
- Teltonika ModemCOM/G10

To configure SMS as the delivery channel for messages, you also need to configure SMS as the media type and enter the respective phone numbers for the users.

Configuration

To configure SMS as the media type:

- Go to *Administration*→*Media types*
- Click on *Create media type* (or click on *SMS* in the list of pre-defined media types).

Media type attributes:

Parameter	Description
<i>Description</i>	Name of the media type.
<i>Type</i>	Select <i>SMS</i> as the type.
<i>GSM modem</i>	Set the serial device name of the GSM modem.

User media

To assign a phone number to the user:

- Go to *Administration*→*Users*
- Open the user properties form
- In *Media* tab, click on *Add*

User media attributes:

Parameter	Description
<i>Type</i>	Select <i>SMS</i> as the type.
<i>Send to</i>	Specify the phone number to send messages to.
<i>When active</i>	You can limit the time when messages are sent, for example, the working days only (1-5,09:00-18:00). See the Time period specification page for description of the format.
<i>Use if severity</i>	Mark the checkboxes of trigger severities that you want to receive notifications for.

Parameter	Description
<i>Status</i>	Status of the user media. Enabled - is in use. Disabled - is not being used.

3 Jabber

Overview

Zabbix supports sending Jabber messages.

When sending notifications, Zabbix tries to look up the Jabber SRV record first, and if that fails, it uses an address record for that domain. Among Jabber SRV records, the one with the highest priority and maximum weight is chosen. If it fails, other records are not tried.

To configure Jabber as the delivery channel for messages, you need to configure Jabber as the media type and enter the respective addresses for the users.

Configuration

To configure Jabber as the media type:

- Go to *Administration*→*Media types*
- Click on *Create media type* (or click on *Jabber* in the list of pre-defined media types).

Media type attributes:

Parameter	Description
<i>Description</i>	Name of the media type.
<i>Type</i>	Select <i>Jabber</i> as the type.
<i>Jabber identifier</i>	Enter Jabber identifier.
<i>Password</i>	Enter Jabber password.

User media

To assign a Jabber address to the user:

- Go to *Administration*→*Users*
- Open the user properties form
- In Media tab, click on *Add*

User media attributes:

Parameter	Description
<i>Type</i>	Select <i>Jabber</i> as the type.
<i>Send to</i>	Specify the address to send messages to.
<i>When active</i>	You can limit the time when messages are sent, for example, the working days only (1-5,09:00-18:00). See the Time period specification page for description of the format.
<i>Use if severity</i>	Mark the checkboxes of trigger severities that you want to receive notifications for.
<i>Status</i>	Status of the user media. Enabled - is in use. Disabled - is not being used.

4 Ez Texting

Overview


You can use [Zabbix technological partner](#) Ez Texting for message sending.

To configure Ez Texting as the delivery channel for messages, you need to configure Ez Texting as the media type and assign recipient identification to the users.

Configuration

To configure Ez Texting as the media type:

- Go to *Administration*→*Media types*
- Click on *Create media type*



Media type attributes:

Parameter	Description
<i>Description</i>	Name of the media type.
<i>Type</i>	Select <i>Ez Texting</i> as the type.
<i>Username</i>	Enter the Ez Texting username.
<i>Password</i>	Enter the Ez Texting password.
<i>Message text limit</i>	Select the message text limit. USA (160 characters) Canada (136 characters)

User media

To assign Ez Texting recipient identification to the user:

- Go to *Administration*→*Users*
- Open the user properties form
- In *Media* tab, click on *Add*

User media attributes:

Parameter	Description
<i>Type</i>	Select the Ez Texting media type.
<i>Send to</i>	Specify the recipient to send the messages to.
<i>When active</i>	You can limit the time when messages are sent, for example, the working days only (1-5,09:00-18:00). See the Time period specification page for description of the format.
<i>Use if severity</i>	Mark the checkboxes of trigger severities that you want to receive notifications for.
<i>Status</i>	Status of the user media. Enabled - is in use. Disabled - is not being used.

5 Custom alertscripts

Overview

If you are not satisfied with existing media types for sending alerts there is an alternative way to do that. You can create a script that will handle the notification your way. These scripts are located in the directory defined in the [Zabbix server configuration file](#) **AlertScriptsPath** variable. When alert script is executed it gets 3 command-line variables (as \$1, \$2 and \$3 respectively):

- To
- Subject
- Message

Note:

Alert scripts are executed on the Zabbix server.

The recipient ("To") is specified in **user media properties**. Here is an example alert script:

```
#####!/bin/bash

to=$1
subject=$2
body=$3

cat <<EOF | mail -s "$subject" "$to"
$body
EOF
```

Environment variables are not preserved or created for the script, so they should be handled explicitly.

Configuration

To configure custom alertscripts as the media type:

- Go to *Administration*→*Media types*
- Click on *Create media type*

Media type attributes:

Parameter	Description
<i>Description</i>	Name of the media type.
<i>Type</i>	Select <i>Script</i> as the type.
<i>Script name</i>	Enter the name of the script.

User media

To assign custom alertscripts to the user:

- Go to *Administration*→*Users*
- Open the user properties form
- In *Media* tab, click on *Add*

User media attributes:

Parameter	Description
<i>Type</i>	Select the custom alertscripts media type.
<i>Send to</i>	Specify the recipient to receive the alerts.
<i>When active</i>	You can limit the time when alertscripts are executed, for example, the working days only (1-5,09:00-18:00). See the Time period specification page for description of the format.
<i>Use if severity</i>	Mark the checkboxes of trigger severities that you want to activate the alertscript for.
<i>Status</i>	Status of the user media. Enabled - is in use. Disabled - is not being used.

2 Actions

Overview

If you want some operations taking place as a result of events (for example, notifications sent), you need to configure actions.

Actions can be defined for the events of all three sources:

- Triggers - when trigger status changes
- Discovery - when discovery takes place
- Auto registration - when new active agents auto-register

Configuring an action

To configure an action, do the following:

- Go to *Configuration* → *Actions*
- From the *Event source* dropdown select the required source
- Click on *Create action*
- Set general action attributes
- Choose the **operation** to carry out, in Operations tab
- Choose the **conditions** upon which the operation is carried out, in Conditions tab

General action attributes:

Action	Conditions	Operations
Name	Report problems to Zabbix administrators	
Default operation step duration	300 (minimum 60 seconds)	
Default subject	{TRIGGER.STATUS}: {TRIGGER.NAME}	
Default message	Trigger: {TRIGGER.NAME} Trigger status: {TRIGGER.STATUS} Trigger severity: {TRIGGER.SEVERITY} Trigger expression: {TRIGGER.EXPRESSION} 1. Item value on {HOST.NAME1}: {ITEM.VALUE1} ({ITEM.NAME1}) 2. Item value on {HOST.NAME2}: {ITEM.VALUE2} ({ITEM.NAME2})	
Recovery message	☒	
Recovery subject	{TRIGGER.STATUS}: {TRIGGER.NAME}	
Recovery message	Trigger: {TRIGGER.NAME} Trigger status: {TRIGGER.STATUS} {EVENT.ID} {EVENT.AGE} {EVENT.DATE} {EVENT.TIME} {EVENT.ACK.STATUS} {EVENT.ACK.HISTORY} {EVENT.RECOVERY.ID}	
Enabled	☒	

Parameter	Description
<i>Name</i>	Unique action name.
<i>Default operation step duration</i>	Duration of one operation step by default (minimum 60 seconds). For example, an hour-long step duration means that if an operation is carried out, an hour will pass before the next step.
<i>Default subject</i>	Naming before Zabbix 2.0.2: <i>Default escalation period</i>
<i>Default message</i>	Default message subject. The subject may contain macros . Default message. The message may contain macros .

Parameter	Description
<i>Recovery message</i>	Mark the checkbox to turn on a <i>Recovery message</i> . <i>Recovery message</i> is a special way of getting notified for a resolved problem. If turned on, only a single message with a custom subject/body is sent if trigger value changes to OK. <i>Note:</i> To receive a recovery message, "Trigger value= <i>Problem</i> " must be present in action conditions ; "Trigger value= <i>OK</i> ", however, must not be present. (If "Trigger value= <i>OK</i> " is set, the recovery message will not work; instead you will get a full escalation of defined messages and/or remote commands in the same way as for a problem situation). <i>Recovery message</i> will be sent only to those who received any messages regarding the problem before.
<i>Recovery subject</i>	Recovery message subject. It may contain macros.
<i>Recovery message</i>	Recovery message. It may contain macros. Starting with Zabbix 2.0.4, recovery message can also be left empty.
<i>Enabled</i>	Mark the checkbox to enable the action. Otherwise it will be disabled.

1 Operations

Overview

You can define the following operations for all events:

- send a **message**
- execute a **remote command** (including IPMI)

For discovery events, there are additional operations available:

- add host
- remove host
- enable host
- disable host
- add to group
- delete from group
- link to template
- unlink from template

The additional operations available for auto-registration events are:

- add host
- disable host
- add to group
- link to template

Configuring an operation

To configure an operation, go to *Operations* tab in the action properties form and click on *New*. Edit the operation step and click on *Add* to add to the list of *Action operations*.

Operation attributes:

ActionConditionsOperations

Action operations

Steps	Details	Start in	Duration (sec)	Action
1	Send message to users: Admin Send message to user groups: Zabbix administrators	Immediately	Default	<a>Edit <a>Remove

Operation details

Step

From

2

To

2

(0 - infinitely)

Step duration

3600

(minimum 60 seconds, 0 - use action default)

Operation type

Send message

Send to User groups

User group	Action
IT management	<a>Remove
<a>Add	

Send to Users

User	Action
<a>Add	

Send only to

Email

Default message

☒

Conditions

Label	Name	Action
(A)	Event acknowledged = "Not Ack"	<a>Remove
<a>New		

Add
Cancel

Parameter

Description

Action operations

Action operations are displayed, with these details:

Steps - escalation step(s) to which the operation is assigned

Details - type of operation and its recipient/target.

Start in - how long after an event the operation is performed

Duration (sec) - step duration is displayed. *Default* is displayed if the step uses default duration, and a time is displayed if custom duration is used.

Action - links for editing and removing an operation are displayed.

To configure a new operation, click on *New*.

Operation details

This block is used to configure the details of an operation.

Step

Select the step(s) to assign the operation to in an **escalation** schedule:

From - execute starting with this step

To - execute until this step (0=infinity, execution will not be limited)

Step duration - custom duration for these steps (0=use default step duration).

Several operations can be assigned to the same step. If these operations have different step duration defined, the shortest one is taken into account and applied to the step.

Parameter	Description
Operation type Operation type: send mes- sage Send to user groups Send to users Send only to Default message Subject Message Operation type: re- mote com- mand Target list	Two operation types are available for all events: Send message - send message to user Remote command - execute a remote command More operations are available for discovery and auto-registration based events (see above). Click on <i>Add</i> to select user groups to send the message to. The user group must have at least "read" permissions to the host in order to be notified. Click on <i>Add</i> to select users to send the message to. The user must have at least "read" permissions to the host in order to be notified. Send message to all defined media types or a selected one only. If selected, the default message will be used (see general action attributes). Subject of the custom message. The subject may contain macros. The custom message. The message may contain macros. Select targets to execute the command on: Current host - command is executed on the host of the trigger that caused the problem event. This option will not work if there are multiple hosts in the trigger. Host - select host(s) to execute the command on. Host group - select host group(s) to execute the command on. A command on a host is executed only once, even if the host matches more than once (e.g. from several host groups; individually and from a host group). The target list is meaningless if the command is executed on Zabbix server. Selecting more targets in this case only results in the command being executed on the server more times. Note that for global scripts, the target selection also depends on the <i>Host group</i> setting in global script configuration.

Parameter		Description
	<i>Type</i>	Select the command type: IPMI - execute an IPMI command Custom script - execute a custom set of commands SSH - execute an SSH command Telnet - execute a Telnet command Global script - execute one of the global scripts defined in <i>Administration</i> → <i>Scripts</i> .
	<i>Execute on</i>	Execute a custom script on Zabbix server or Zabbix agent. To execute scripts on the agent, it must be configured to allow remote commands from the server. This field is available if 'Custom script' is selected as <i>Type</i> .
	<i>Commands</i>	Enter the command(s). Supported macros will be resolved based on the trigger expression that caused the event. For example, host macros will resolve to the hosts of the trigger expression (and not of the target list).
	<i>Conditions</i>	Condition for performing the operation: Not ack - only when the event is unacknowledged Ack - only when the event is acknowledged.

1 Sending message

Overview

Sending a message is one of the best ways of notifying people about a problem. That is why it is one of the primary actions offered by Zabbix.

Configuration

To be able to send and receive notifications from Zabbix you have to:

- **define the media** to send a message to
- **configure an action operation** that sends a message to one of the defined media

Attention:

Zabbix sends notifications only to those users that have at least 'read' permissions to the host that generated the event. At least one host of a trigger expression must be accessible.

You can configure custom scenarios for sending messages using **escalations**.

To successfully receive and read e-mails from Zabbix, e-mail servers/clients must support standard 'SMTP/MIME e-mail' format since Zabbix sends UTF-8 data (If the subject contains ASCII characters only, it is not UTF-8 encoded.). The subject and the body of the message are base64-encoded to follow 'SMTP/MIME e-mail' format standard.

Tracking messages

You can view the status of messages sent in *Monitoring* → *Events*.

In the *Actions* column you can see summarized information about actions taken. In there green numbers represent messages sent, red ones - failed messages. *In progress* indicates that an action is initiated. *Failed* informs that no action has executed successfully.

If you click on the event time to view event details, you will also see the *Message actions* block containing details of messages sent (or not sent) due to the event.

In *Administration* → *Audit*, if you select *Actions* from the dropdown, you will see details of all actions taken for those events that have an action configured.

2 Remote commands

Overview

With remote commands you can define that a certain pre-defined command is automatically executed on the monitored host upon some condition.

Thus remote commands are a powerful mechanism for smart pro-active monitoring.

In the most obvious uses of the feature you can try to:

- Automatically restart some application (web server, middleware, CRM) if it does not respond
- Use IPMI 'reboot' command to reboot some remote server if it does not answer requests
- Automatically free disk space (removing older files, cleaning /tmp) if running out of disk space
- Migrate a VM from one physical box to another depending on the CPU load
- Add new nodes to a cloud environment upon insufficient CPU (disk, memory, whatever) resources

Configuring an action for remote commands is similar to that for sending a message, the only difference being that Zabbix will execute a command instead of sending a message.

Attention:

Remote commands are not supported to be executed on Zabbix agents monitored by Zabbix proxy, so for commands from Zabbix server to agent a direct connection is required.

Remote commands are limited to 255 characters. Multiple commands can be executed one after another by placing them on a new line. Remote commands may contain macros. See also the [command execution](#) page.

Remote commands are executed even if the target host is in maintenance.

The following tutorial provides step-by-step instructions on how to set up remote commands.

Configuration

Those remote commands that are executed on Zabbix agent (custom scripts) must be first enabled in the respective `zabbix_agentd.conf`.

Make sure that the **EnableRemoteCommands** parameter is set to **1** and uncommented. Restart agent daemon if changing this parameter.

Attention:

Remote commands do not work with active Zabbix agents.

Then, when configuring a new action in *Configuration→Actions*:

- In the *Operations* tab, select the **Remote command** operation type
- Select the remote command type (IPMI, Custom script, SSH, Telnet, Global script)
- Enter the remote command

For example:

```
sudo /etc/init.d/apache restart
```

In this case, Zabbix will try to restart an Apache process. With this command, make sure that the command is executed on Zabbix agent (mark the respective radio button against *Execute on*).

Attention:

Note the use of **sudo** - Zabbix user does not have permissions to restart system services by default. See below for hints on how to configure **sudo**.

Note:

Zabbix agent should run on the remote host and accept incoming connections. Zabbix agent executes commands in background.

Attention:

Zabbix does not check if a command has been executed successfully.

Remote commands on Zabbix agent are executed without timeout by the `system.run[,nowait]` key. On Zabbix server remote commands are executed with timeout as set in the `TrapperTimeout` parameter of `zabbix_server.conf` file.

- In the *Conditions* tab, define the appropriate conditions. In this example, set that the action is activated upon any disaster problems with one of Apache applications.

ActionConditionsOperations

Type of calculation

AND / OR

(A) and (B) and (C) and (D)

Conditions

(A)

☐

Maintenance status not in "maintenance"

(B)

☐

Application like "Apache"

(C)

☐

Trigger value = "PROBLEM"

(D)

☐

Trigger severity >= "Disaster"

Delete selected

Access permissions

Make sure that the 'zabbix' user has execute permissions for configured commands. One may be interested in using **sudo** to give access to privileged commands. To configure access, execute as root:

visudo

Example lines that could be used in *sudoers* file:

allows 'zabbix' user to run all commands without password.
zabbix ALL=NOPASSWD: ALL

allows 'zabbix' user to restart apache without password.
zabbix ALL=NOPASSWD: /etc/init.d/apache restart

Note:

On some systems *sudoers* file will prevent non-local users from executing commands. To change this, comment out **requiretty** option in */etc/sudoers*.

Remote commands with multiple interfaces

If the target system has multiple interfaces of the selected type (Zabbix agent or IPMI), remote commands will be executed on the default interface.

IPMI remote commands

For IPMI remote commands the following syntax should be used:

<command> [<value>]

where

- <command> - one of IPMI commands without spaces
- <value> - 'on', 'off' or any unsigned integer. <value> is an optional parameter.

Examples

Example 1

Restart of Windows on certain condition.

In order to automatically restart Windows upon a problem detected by Zabbix, define the following actions:

PARAMETER	Description
Operation type	'Remote command'
Type	'Custom script'
Command	c:\windows\system32\shutdown.exe -r -f

Example 2

Restart the host by using IPMI control.

PARAMETER	Description
Operation type	'Remote command'
Type	'IPMI'
Command	reset

Example 3

Power off the host by using IPMI control.

PARAMETER	Description
Operation type	'Remote command'
Type	'IPMI'
Command	power off

3 Additional operations

Overview

For discovery events, there are additional operations available:

- add host
- remove host
- enable host
- disable host
- add to group
- delete from group
- link to template
- unlink from template

The additional operations available for auto-registration events are:

- add host
- disable host
- add to group
- link to template

Adding host

Hosts are added during the discovery process, as soon as a host is discovered, rather than at the end of the discovery process.

Note:

As network discovery can take some time due to many unavailable hosts/services having patience and using reasonable IP ranges is advisable.

When adding a host, its name is decided by the standard **gethostbyname** function. If the host can be resolved, resolved name is used. If not, the IP address is used. Besides, if IPv6 address must be used for a host name, then all ":" (colons) are replaced by "_" (underscores), since colons are not allowed in host names.

Attention:

If performing discovery by a proxy, currently hostname lookup still takes place on Zabbix server.

Attention:

If a host already exists in Zabbix configuration with the same name as a newly discovered one, versions of Zabbix prior to 1.8 would add another host with the same name. Zabbix 1.8.1 and later adds **_N** to the hostname, where **N** is increasing number, starting with 2.

4 Using macros in messages

Overview

In message subjects and message text you can use macros for more efficient problem reporting.

A [full list of macros](#) supported by Zabbix is available.

Examples

Examples here illustrate how you can use macros in messages.

Example 1

Message subject:

```
{TRIGGER.NAME}: {TRIGGER.STATUS}
```

When you receive the message, the message subject will be replaced by something like:

```
Processor load is too high on server zabbix.zabbix.com: PROBLEM
```

Example 2

Message:

```
Processor load is: {zabbix.zabbix.com:system.cpu.load[,avg1].last(0)}
```

When you receive the message, the message will be replaced by something like:

```
Processor load is: 1.45
```

Example 3

Message:

```
Latest value: {{HOST.HOST}}:{{ITEM.KEY}}.last(0)}
```

```
MAX for 15 minutes: {{HOST.HOST}}:{{ITEM.KEY}}.max(900)}
```

```
MIN for 15 minutes: {{HOST.HOST}}:{{ITEM.KEY}}.min(900)}
```

When you receive the message, the message will be replaced by something like:

```
Latest value: 1.45
```

```
MAX for 15 minutes: 2.33
```

```
MIN for 15 minutes: 1.01
```

Example 4

Informing about values from several hosts in a trigger expression.

Message:

```
Trigger: {TRIGGER.NAME}
```

```
Trigger expression: {TRIGGER.EXPRESSION}
```

```
1. Item value on {HOST.NAME1}: {ITEM.VALUE1} ({ITEM.NAME1})
```

```
2. Item value on {HOST.NAME2}: {ITEM.VALUE2} ({ITEM.NAME2})
```

When you receive the message, the message will be replaced by something like:

```
Trigger: Processor load is too high on a local host
```

```
Trigger expression: {Myhost:system.cpu.load[percpu,avg1].last(0)}>5 | {Myotherhost:system.cpu.load[percpu,
```

```
1. Item value on Myhost: 0.83 (Processor load (1 min average per core))
```

```
2. Item value on Myotherhost: 5.125 (Processor load (1 min average per core))
```

2 Conditions

Overview

An action is executed only in case an event matches a defined set of conditions.

Configuration

To set a condition:

- Go to *Conditions* tab in the action properties form
- Select conditions from the *New condition* dropdowns and click on *Add*
- Select the type of calculation (with more than one condition)

The following conditions can be set for trigger-based actions:

Condition type	Supported operators	Description
<i>Application</i>	= like not like	Specify an application or an application to exclude. = - event belongs to a trigger of the item that is linked to the specified application. like - event belongs to a trigger of the item that is linked to an application containing the string. not like - event belongs to a trigger of the item that is linked to an application not containing the string.
<i>Host group</i>	= <>	Specify a host group or a host group to exclude. = - event belongs to this host group. <> - event does not belong to this host group.
<i>Host template</i>	= <>	Specify a host template or a template to exclude. = - event belongs to a trigger inherited from this host template. <> - event does not belong to a trigger inherited from this host template.
<i>Host</i>	= <>	Specify a host or a host to exclude. = - event belongs to this host. <> - event does not belong to this host.
<i>Trigger</i>	= <>	Specify a trigger or a trigger to exclude. = - event is generated by this trigger. <> - event is generated by any other trigger, except this one.

Condition type	Supported operators	Description
<i>Trigger name</i>	like not like	Specify a string in the trigger name or a string to exclude. like - event is generated by a trigger, containing this string in the name. Case sensitive. not like - this string cannot be found in the trigger name. Case sensitive. <i>Note:</i> Entered value will be compared to trigger name with all macros expanded.
<i>Trigger severity</i>	= <> >= <=	Specify trigger severity. = - equal to trigger severity <> - not equal to trigger severity >= - more or equal to trigger severity <= - less or equal to trigger severity
<i>Trigger value</i>	=	Specify a trigger value. = - equal to trigger value (OK or PROBLEM)
<i>Time period</i>	in not in	Specify a time period or a time period to exclude. in - event time is within the time period. not in - event time is not within the time period. See Time period specification page for description of the format.
<i>Maintenance status</i>	in not in	Specify a host in maintenance or not in maintenance. in - host is in maintenance mode. not in - host is not in maintenance mode. <i>Note:</i> If several hosts are involved in the trigger expression, the condition matches if at least one of the hosts is/is not in maintenance mode.

Trigger value:

- if a trigger changes status from OK to PROBLEM, trigger value is PROBLEM
- if a trigger changes status from PROBLEM to OK, trigger value is OK

Note:

When a new action for triggers is created, it gets two automatic conditions (both can be removed by the user):

- "Trigger value = *PROBLEM*" - so that notifications are sent for problems only. This means that if you configure an action without any more specific conditions, messages will be received for any problem. Having this condition by default is also important if you want to receive a single **recovery message**.
- "Maintenance status = not in *maintenance*" - so that notifications are not sent for hosts in maintenance.

The following conditions can be set for discovery-based events:

Condition type	Supported operators	Description
<i>Host IP</i>	= <>	Specify an IP address range or a range to exclude for a discovered host. = - host IP is in the range. <> - host IP is not in the range.
<i>Service type</i>	= <>	Specify a service type of a discovered service or a service type to exclude. = - matches the discovered service. <> - does not match the discovered service.
<i>Service port</i>	= <>	Specify a TCP port range of a discovered service or a range to exclude. = - service port is in the range. <> - service port is not in the range.
<i>Discovery rule</i>	= <>	Specify a discovery rule or a discovery rule to exclude. = - using this discovery rule. <> - using any other discovery rule, except this one.
<i>Discovery check</i>	= <>	Specify a discovery check or a discovery check to exclude. = - using this discovery check. <> - using any other discovery check, except this one.
<i>Discovery object</i>	=	Specify the discovered object. = - equal to discovered object (a device or a service).
<i>Discovery status</i>	=	Up - matches 'Host Up' and 'Service Up' events Down - matches 'Host Down' and 'Service Down' events Discovered - matches 'Host Discovered' and 'Service Discovered' events Lost - matches 'Host Lost' and 'Service Lost' events
<i>Uptime/Downtime</i>	>= <=	Uptime for 'Host Up' and 'Service Up' events. Downtime for 'Host Down' and 'Service Down' events. >= - is more or equal to. Parameter is given in seconds. <= - is less or equal to. Parameter is given in seconds.

Condition type	Supported operators	Description
<i>Received value</i>	=	Specify the value received from an agent (Zabbix, SNMP). Case sensitive string comparison. If multiple Zabbix agent or SNMP checks are configured for a rule, received values for all of them are checked (each check generates a new event which is matched against all conditions). = - equal to the value. <> - not equal to the value. >= - more or equal to the value. <= - less or equal to the value. like - contains the substring. Parameter is given as a string. not like - does not contain the substring. Parameter is given as a string.
	<>	
	>=	
	<=	
	like	
	not like	
<i>Proxy</i>	=	Specify a proxy or a proxy to exclude. = - using this proxy. <> - using any other proxy except this one.
	<>	

Note:

Service checks in a discovery rule, which result in discovery events, do not take place simultaneously. Therefore, if **multiple** values are configured for *Service type*, *Service port* or *Received value* conditions in the action, they will be compared to one discovery event at a time, but **not** to several events simultaneously. As a result, actions with multiple values for the same check types may not be executed correctly.

Type of calculation

The following options of calculating conditions are available:

- AND - all conditions must be met

Note that "AND" calculation should not be used between several triggers when they are selected as a *Trigger=* condition. Actions can only be executed based on the event of one trigger.

- OR - enough if one condition is met
- AND/OR - combination of the two: AND with different condition types and OR with the same condition type, for example:

Host group = Oracle servers

Host group = MySQL servers

Trigger name like 'Database is down'

Trigger name like 'Database is unavailable'

is evaluated as

(*Host group* = Oracle servers **or** *Host group* = MySQL servers) **and** (*Trigger name* like 'Database is down' **or** *Trigger name* like 'Database is unavailable')

Actions disabled due to deleted objects

Since Zabbix 2.0.6 if a certain object (host, template, trigger, etc) used in an action condition/operation is deleted, the condition/operation is removed and the action is disabled to avoid incorrect execution of the action. The action can be re-enabled by the user.

This behavior takes place when deleting:

- host groups ("host group" condition, "remote command" operation on a specific host group);
- hosts ("host" condition, "remote command" operation on a specific host);
- templates ("template" condition, "link to template" and "unlink from template" operations);
- triggers ("trigger" condition);
- discovery rules (when using "discovery rule" and "discovery check" conditions);
- proxies ("proxy" condition).

Before 2.0.6 the missing object was displayed as *unknown* and the condition remained in place.

Note: If a remote command has many target hosts, and we delete one of them, only this host will be removed from the target list, the operation itself will remain. But, if it's the only host, the operation will be removed, too. The same goes for "link to template" and "unlink from template" operations.

Actions are not disabled when deleting a user or user group used in a "send message" operation.

3 Escalations

Overview

With escalations you can create custom scenarios for sending notifications or executing remote commands.

In practical terms it means that:

- Users can be informed about new problems immediately
- Notifications can be repeated until the problem is resolved
- Sending a notification can be delayed
- Notifications can be escalated to another "higher" user group
- Remote commands can be executed immediately or when a problem is not resolved for a lengthy period
- Recovery messages can be sent

Actions are escalated based on the **escalation step**. Each step has a duration in time.

You can define both the default duration and a custom duration of an individual step. The minimum duration of one escalation step is 60 seconds.

You can start actions, such as sending notifications or executing commands, from any step. Step one is for immediate actions. If you want to delay an action, you can assign it to a later step. For each step, several actions can be defined.

The number of escalation steps is not limited.

Escalations are defined when **configuring an operation**.

Miscellaneous aspects of escalation behaviour

Let's consider what happens in different circumstances if an action contains several escalation steps.

Situation	Behaviour
<i>The host in question goes into maintenance after the initial problem notification is sent</i>	All remaining escalation steps are executed. A maintenance cannot stop operations; maintenance has effect with regard to when actions are started/not started, not operations.
<i>The time period defined in the Time period action condition ends after the initial notification is sent</i>	All remaining escalation steps are executed. The <i>Time period</i> condition cannot stop operations; it has effect with regard to when actions are started/not started, not operations.
<i>A problem starts during maintenance and continues (is not resolved) after maintenance ends</i>	All escalation steps are executed starting from the moment maintenance ends.
<i>A problem starts during a no-data maintenance and continues (is not resolved) after maintenance ends</i>	It must wait for the trigger to fire, before all escalation steps are executed.
<i>Different escalations follow in close succession and overlap</i>	The execution of each new escalation supersedes the previous escalation, but for at least one escalation step that is always executed on the previous escalation. This behavior is relevant in actions upon events that are created with EVERY problem evaluation of the trigger.

Situation	Behaviour
An action is disabled during an escalation in progress (like a message being sent)	The message in progress will be sent and then one more message on the escalation will be sent. The follow-up message will have the following text at the beginning of the message body: <i>NOTE: Escalation cancelled: action '<Action name>' disabled.</i> This way the recipient is informed that the escalation is cancelled and no more steps will be executed. This message is sent to the recipients defined in the following escalation step.

Escalation examples

Example 1

Sending a repeated notification once every 30 minutes (5 times in total) to a 'MySQL Administrators' group. To configure:

- Set the *Default operation step duration* to '1800' seconds (30 minutes) in general action attributes
- in Operations tab, set the escalation steps to be *From '1' To '5'*
- Select the 'MySQL Administrators' group as recipients of the message

Action

Conditions

Operations

Action operations

Steps	Details	Start in	Duration (sec)	Action
1 - 5	Send message to user groups: MySQL Administrators	Immediately	Default	Edit Remove
New				

Notifications will be sent at 0:00, 0:30, 1:00, 1:30, 2:00 hours after the problem starts (unless, of course, the problem is resolved sooner).

If the problem is resolved and a recovery message is configured, it will be sent to those who received at least one problem message within this escalation scenario.

Note:

If the trigger that generated an active escalation is disabled, Zabbix sends an informative message about it to all those that have already received notifications.

Example 2

Sending a delayed notification about a long-standing problem. To configure:

- Set the *Default operation step duration* to '36000' seconds (10 hours) in general action attributes
- In Operations tab, set the escalation steps to be *From '2' To '2'*

Action

Conditions

Operations

Action operations

Steps	Details	Start in	Duration (sec)	Action
2	Send message to users: Dpt_Head	10:00:00	Default	Edit Remove
New				

A notification will only be sent at Step 2 of the escalation scenario, or 10 hours after the problem starts.

You can customize the message text to something like 'The problem is more than 10 hours old'.

Example 3

Escalating the problem to the Boss.

In the first example above we configured periodical sending of messages to MySQL administrators. In this case, the administrators will get four messages before the problem will be escalated to the Database manager. Note that the manager will get a message only in case the problem is not acknowledged yet, supposedly no one is working on it.

Action	Conditions	Operations																			
Action operations <table border="1"> <thead> <tr> <th>Steps</th> <th>Details</th> <th>Start in</th> <th>Duration (sec)</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>1 - 0</td> <td>Send message to user groups: MySQL Administrators</td> <td>Immediately</td> <td>Default</td> <td>Edit Remove</td> </tr> <tr> <td>5</td> <td>Send message to user groups: Database manager</td> <td>02:00:00</td> <td>Default</td> <td>Edit Remove</td> </tr> </tbody> </table>			Steps	Details	Start in	Duration (sec)	Action	1 - 0	Send message to user groups: MySQL Administrators	Immediately	Default	Edit Remove	5	Send message to user groups: Database manager	02:00:00	Default	Edit Remove				
Steps	Details	Start in	Duration (sec)	Action																	
1 - 0	Send message to user groups: MySQL Administrators	Immediately	Default	Edit Remove																	
5	Send message to user groups: Database manager	02:00:00	Default	Edit Remove																	
Operation details Step From 5 To 5 (0 - infinitely) Step duration 0 (minimum 60 seconds, 0 - use action default) Operation type Send message Send to User groups <table border="1"> <thead> <tr> <th>User group</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>Database manager</td> <td>Remove</td> </tr> <tr> <td>Add</td> <td></td> </tr> </tbody> </table> Send to Users <table border="1"> <thead> <tr> <th>User</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>Add</td> <td></td> </tr> </tbody> </table> Send only to Email Default message ☐ Subject Unacknowledged problem Message Trigger: {TRIGGER.NAME} Trigger status: {TRIGGER.STATUS} Trigger severity: {TRIGGER.SEVERITY} Escalation history: {ESC.HISTORY} Conditions <table border="1"> <thead> <tr> <th>Label</th> <th>Name</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>(A)</td> <td>Event acknowledged = "Not Ack"</td> <td>Remove</td> </tr> <tr> <td>New</td> <td></td> <td></td> </tr> </tbody> </table> Update Cancel			User group	Action	Database manager	Remove	Add		User	Action	Add		Label	Name	Action	(A)	Event acknowledged = "Not Ack"	Remove	New		
User group	Action																				
Database manager	Remove																				
Add																					
User	Action																				
Add																					
Label	Name	Action																			
(A)	Event acknowledged = "Not Ack"	Remove																			
New																					

Note the use of {ESC.HISTORY} macro in the message. The macro will contain information about all previously executed steps on this escalation, such as notifications sent and commands executed.

Example 4

A more complex scenario. After multiple messages to MySQL administrators and escalation to the manager, Zabbix will try to restart the MySQL database. It will happen if the problem exists for 2:30 hours and it hasn't been acknowledged.

If the problem still exists, after another 30 minutes Zabbix will send a message to all guest users.

If this does not help, after another hour Zabbix will reboot server with the MySQL database (second remote command) using IPMI commands.

Action	Conditions	Operations																																			
Action operations <table border="1"> <thead> <tr> <th>Steps</th> <th>Details</th> <th>Start in</th> <th>Duration (sec)</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>1 - 0</td> <td>Send message to user groups: MySQL Administrators</td> <td>Immediately</td> <td>Default</td> <td>Edit Remove</td> </tr> <tr> <td>5</td> <td>Send message to user groups: Database manager</td> <td>02:00:00</td> <td>Default</td> <td>Edit Remove</td> </tr> <tr> <td>6</td> <td>Run remote commands on current host</td> <td>02:30:00</td> <td>Default</td> <td>Edit Remove</td> </tr> <tr> <td>7</td> <td>Send message to user groups: Guests</td> <td>03:00:00</td> <td>Default</td> <td>Edit Remove</td> </tr> <tr> <td>9</td> <td>Run remote commands on current host</td> <td>04:00:00</td> <td>Default</td> <td>Edit Remove</td> </tr> <tr> <td>New</td> <td></td> <td></td> <td></td> <td></td> </tr> </tbody> </table>			Steps	Details	Start in	Duration (sec)	Action	1 - 0	Send message to user groups: MySQL Administrators	Immediately	Default	Edit Remove	5	Send message to user groups: Database manager	02:00:00	Default	Edit Remove	6	Run remote commands on current host	02:30:00	Default	Edit Remove	7	Send message to user groups: Guests	03:00:00	Default	Edit Remove	9	Run remote commands on current host	04:00:00	Default	Edit Remove	New				
Steps	Details	Start in	Duration (sec)	Action																																	
1 - 0	Send message to user groups: MySQL Administrators	Immediately	Default	Edit Remove																																	
5	Send message to user groups: Database manager	02:00:00	Default	Edit Remove																																	
6	Run remote commands on current host	02:30:00	Default	Edit Remove																																	
7	Send message to user groups: Guests	03:00:00	Default	Edit Remove																																	
9	Run remote commands on current host	04:00:00	Default	Edit Remove																																	
New																																					

Example 5

An escalation with several operations assigned to one step and custom intervals used. The default operation step duration is 30 minutes.

Action	Conditions	Operations		
Action operations				
Steps	Details	Start in	Duration (sec)	Action
1 - 4	Send message to user groups: MySQL Administrators	Immediately	Default	Edit Remove
5 - 6	Send message to user groups: Database manager	02:00:00	3600	Edit Remove
5 - 7	Send message to user groups: Zabbix administrators	02:00:00	600	Edit Remove
11	Send message to user groups: Guests	04:00:00	Default	Edit Remove
New				

Notifications will be sent as follows:

- to MySQL administrators at 0:00, 0:30, 1:00, 1:30 after the problem starts
- to Database manager at 2:00 and 2:10 (and not at 3:00; seeing that steps 5 and 6 overlap, the shorter custom step duration of 600 seconds in the next operation overrides the longer custom step duration of 3600 seconds tried to set here)
- to Zabbix administrators at 2:00, 2:10, 2:20 after the problem starts (the custom step duration of 600 seconds working)
- to guest users at 4:00 hours after the problem start (the default step duration of 30 minutes returning between steps 8 and 11)

8 Macros

Overview

Zabbix supports a number of macros which may be used in various situations. Macros are variables, identified by a **{MACRO}** syntax, and resolve to a specific value depending on the context.

Effective use of macros allows to save time and make Zabbix configuration more transparent.

In one of typical uses, a macro may be used in a template. Thus a trigger on a template may be named "Processor load is too high on {HOST.NAME}". When the template is applied to the host, such as Zabbix server, the name will resolve to "Processor load is too high on Zabbix server" when the trigger is displayed in the Monitoring section.

Macros may be used in item key parameters. A macro may be used for only a part of the parameter, for example `item.key[server_{HOST.HOST}_local]`. Double-quoting the parameter is not necessary as Zabbix will take care of any ambiguous special symbols, if present in the resolved macro.

See a full list of [supported macros](#) by location.

You can also configure your own [user macros](#).

1 User macros

Overview

For greater flexibility, Zabbix supports user macros, which can be defined on global, template and host level. These macros have a special syntax: **{\$MACRO}**.

The macros can be used in:

- item names
- item key parameters
- trigger names
- trigger expression parameters and constants (see examples)
- several other [locations](#)

The following characters are allowed in the macro names: **A-Z**, **0-9**, **_**, **.**

Zabbix substitutes macros according to the following precedence:

1. host level macros (checked first)
2. macros defined for first level templates of the host (i.e., templates linked directly to the host), sorted by template ID
3. macros defined for second level templates of the host, sorted by template ID
4. macros defined for third level templates of the host, sorted by template ID
5. ...
6. global macros (checked last)

In other words, if a macro does not exist for a host, Zabbix will try to find it in the host templates of increasing depth. If still not found, a global macro will be used, if exists.

If Zabbix is unable to find a macro, the macro will not be substituted.

To define user macros, go to the corresponding locations in the frontend:

- for global macros, visit *Administration* → *General* → *Macros*
- for host and template level macros, open host or template properties and look for the *Macros* tab

Note:

If a user macro is used in items or triggers in a template, it is suggested to add that macro to the template even if it is defined on a global level. That way, exporting the template to XML and importing it in another system will still allow it to work as expected.

Most common use cases of global and host macros:

1. taking advantage of templates with host specific attributes: passwords, port numbers, file names, regular expressions, etc
2. global macros for global one-click configuration changes and fine tuning

Examples

Example 1

Use of host-level macro in the "Status of SSH daemon" item key:

net.tcp.service[ssh,{\$SSH_PORT}]

This item can be assigned to multiple hosts, providing that the value of **{\$SSH_PORT}** is defined on those hosts.

Example 2

Use of host-level macro in the "CPU load is too high" trigger:

{ca_001:system.cpu.load[avg1].last(0)}>{\$MAX_CPULOAD}

Such a trigger would be created on the template, not edited in individual hosts.

Note:

If you want to use amount of values as the function parameter (for example, **max(#3)**), include hash mark in the macro definition like this: **SOME_PERIOD => #3**

Example 3

Use of two macros in the "CPU load is too high" trigger:

{ca_001:system.cpu.load[avg1].min({\$CPULOAD_PERIOD})}>{\$MAX_CPULOAD}

Note that a macro can be used as a parameter of trigger function, in this example function **min()**.

Attention:

In trigger expressions user macros will expand if referencing a parameter or constant. They will NOT expand if referencing the host, item key, function, operator or another trigger expression.

9 Users and user groups

Overview

All users in Zabbix access the Zabbix application through the web-based frontend. Each user is assigned a unique login name and a password.

All user passwords are encrypted and stored in the Zabbix database. Users cannot use their user id and password to log directly into the UNIX server unless they have also been set up accordingly to UNIX. Communication between the web server and the user browser can be protected using SSL.

With a flexible **user permission schema** you can restrict and differentiate access to:

- administrative Zabbix frontend functions
- monitored hosts in hostgroups

The initial Zabbix installation has two predefined users - 'Admin' and 'guest'. The 'guest' user is used for unauthenticated users. Before you log in as 'Admin', you are 'guest'. Proceed to [configuring a user](#) in Zabbix.

1 Configuring a user

Overview

To configure a user:

- Go to *Administration* → *Users*
- Select *Users* from the dropdown to the right
- Click on *Create user* (or on the user name to edit an existing user)
- Edit user attributes in the form

General attributes

The *User* tab contains general user attributes:

UserMediaPermissions

AliasAdmin

NameZabbix

SurnameAdministrator

Password

Change password

Groups

Zabbix administrators

AddDelete selected

Language

English (en_GB)

Theme

System default

Auto-login☒

Auto-logout (min 90 seconds)☐

90

Refresh (in seconds)

30

Rows per page

50

URL (after login)

Parameter	Description
Alias	Unique username, used as the login name.
Name	User first name (required).
Surname	User second name (required).

Parameter	Description
<i>Password</i>	Two fields for entering the user password. With an existing password, contains a <i>Password</i> button, clicking on which opens the password fields.
<i>Groups</i>	List of all user groups the user belongs to. Adherence to user groups determines what host groups and hosts the user will have access to . Click on <i>Add</i> to add groups.
<i>Language</i>	Language of Zabbix GUI.
<i>Theme</i>	Defines how the GUI looks like: System Default - use default system settings Original Blue - standard blue theme Black & Blue - alternative theme Dark orange - alternative theme
<i>Auto-login</i>	Enable if you want Zabbix to remember you and log you in automatically for 30 days. Browser cookies are used for this.
<i>Auto-logout (min 90 seconds)</i>	Mark the checkbox to enable automatic user logout after the set seconds of inactivity (minimum value = 90 seconds).
<i>Refresh (in seconds)</i>	Set the refresh rate used for graphs, screens, plain text data, etc. Can be set to 0 to disable.
<i>Rows per page</i>	You can determine how many rows per page will be displayed in lists.
<i>URL (after login)</i>	You can make Zabbix to transfer you to a specific URL after successful login, for example, the status of triggers page.

User media

The *Media* tab contains a listing of all media defined for the user. Media are used for sending notifications. Click on *Add* to assign media to the user.

See the **Media types** section for details on configuring media types.

Permissions

The *Permissions* tab contains information on:

- the user type (Zabbix User, Zabbix Admin, Zabbix Super Admin). Users cannot change their own type.
- host groups and hosts the user has access to. 'Zabbix User' and 'Zabbix Admin' users do not have access to any host groups and hosts by default. To get access they need to be included in user groups that have access to respective host groups and hosts.

See the **User permissions** page for details.

2 Permissions

Overview

You can differentiate user permissions in Zabbix by defining the respective user type and then by including the unprivileged users in user groups that have access to host group data.

User type

The user type defines the level of access to administrative menus and the default access to host group data.

User type	Description
<i>Zabbix User</i>	The user has access to the Monitoring menu. The user has no access to any resources by default. Any permissions to host groups must be explicitly assigned.
<i>Zabbix Admin</i>	The user has access to the Monitoring and Configuration menus. The user has no access to any host groups by default. Any permissions to host groups must be explicitly given.
<i>Zabbix Super Admin</i>	The user has access to everything: Monitoring, Configuration and Administration menus. The user has a read-write access to all host groups. Permissions cannot be revoked by denying access to specific host groups.

Permissions to host groups

Access to any host data in Zabbix are granted to user groups on host group level only.

That means that an individual user cannot be directly granted access to a host (or host group). It can only be granted access to a host by being part of a user group that is granted access to the host group that contains the host.

3 User groups

Overview

User groups allow to group users both for organizational purposes and for assigning permissions to data. Permissions to monitoring data of host groups are assigned to user groups, not individual users.

It may often make sense to separate what information is available for one group of users and what - for another. This can be accomplished by grouping users and then assigning varied permissions to host groups.

A user can belong to any amount of groups.

Configuration

To configure a user group:

- Go to *Administration* → *Users*
- Select *User groups* from the dropdown to the right
- Click on *Create group* (or on the group name to edit an existing group)
- Edit group attributes in the form

The *User group* tab contains general group attributes:

Parameter	Description
<i>Group name</i> <i>Users</i>	Unique group name. The In group block contains a listing of the members of this group. To add users to the group select them in the <i>Other groups</i> block and click on «.
<i>Frontend access</i>	How the users of the group are authenticated. System default - use default authentication Internal - use Zabbix authentication. Ignored if HTTP authentication is set Disabled - access to Zabbix GUI is forbidden
<i>Users status</i>	Status of group members: Enabled - users are active Disabled - users are disabled
<i>Debug mode</i>	Mark this checkbox to activate debug mode for the users.

The *Permissions* tab allows you to specify user group access to host group (and thereby host) data:

<i>Composing permissions</i>	Click on <i>Add</i> beneath the respective list to specify the host groups that the user group will have access to on the level of: Read-write - read-write access to a host group Read only - read-only access to a host group Deny - access to a host group denied
<i>Calculated permissions</i>	Depending on the permissions set above, <i>Calculated permissions</i> will display all host groups and all hosts that the user group has access to on the level of: Read-write - host groups with read-write access Read only - host groups with read-only access Deny - host groups with access denied

7. IT services

Overview IT services are intended for those who want to get a high-level (business) view of monitored infrastructure. In many cases, we are not interested in low-level details, like the lack of disk space, high processor load, etc. What we are interested in is the availability of service provided by our IT department. We can also be interested in identifying weak places of IT infrastructure, SLA of various IT services, the structure of existing IT infrastructure, and other information of a higher level.

Zabbix IT services provide answers to all mentioned questions.

IT services is a hierarchy representation of monitored data.

A very simple IT service structure may look like:

```
IT Service
|
|-Workstations
| |
| |-Workstation1
| |
| |-Workstation2
|
|-Servers
```

Each node of the structure has attribute status. The status is calculated and propagated to upper levels according to the selected algorithm. At the lowest level of IT services are triggers. The status of individual nodes is affected by the status of their triggers.

Note:
Note that triggers with a *Not classified* or *Information* severity do not impact SLA calculation.

Configuration To configure IT services, go to: *Configuration* → *IT services*.

On this screen you can build a hierarchy of your monitored infrastructure. The highest-level parent service is 'root'. You can build your hierarchy downward by adding lower-level parent services and then individual nodes to them.

Service	Status calculation	Trigger
root		
SLA by service	Problem, if all children have problems	None
Server 1	Problem, if at least one child has a problem	None
Server 2	Problem, if at least one child has a problem	None
Server 3	Problem, if at least one child has a problem	None
Server 4	Problem, if at least one child has a problem	None
Add Service	Problem, if at least one child has a problem	None
Edit Service		
Delete Service		

Copyright 2001-2011 by Zabbix SIA

Connected as 'Admin'

Click on a service to add services to it or edit the service. A form is displayed where you can edit the service attributes.

Configuring an IT service

The **Service** tab contains general service attributes:

ServiceDependenciesTime

Name

Zabbix server

Parent service

root

Change

Status calculation algorithm

Problem, if at least one child has a problem

Calculate SLA, acceptable SLA (in %)

☐ 99.05

Trigger

Zabbix server:Agent is unavailable on Zabbix server

Select

Sort order (0->999)

0

Parameter	Description
Name	Service name.

Parameter	Description
<i>Parent service</i>	Parent service the service belongs to.
<i>Status calculation algorithm</i>	Method of calculating service status: Do not calculate - do not calculate service status Problem, if at least one child has a problem - problem status, if at least one child service has a problem Problem, if all children have problems - problem status, if all child services are having problems
<i>Calculate SLA</i>	Enable SLA calculation and display.
<i>Acceptable SLA (in %)</i>	SLA percentage that is acceptable for this service. Used for reporting.
<i>Trigger</i>	Linkage to trigger: None - no linkage trigger name - linked to the trigger, thus depends on the trigger status Services of the lowest level must be linked to triggers. (Otherwise their state will not be represented accurately.) When triggers are linked, their state prior to linking is not counted.
<i>Sort order</i>	Sort order for display, lowest comes first.

The **Dependencies** tab contains services the service depends on. Click on *Add* to add a service from those that are configured.

ServiceDependenciesTime

Depends on

Services	Soft	Trigger	Action
Local applications	☐	-	Remove
VMs	☒	-	Remove
New host	☒	Zabbix agent on New host is unreachable for 5 minutes	Remove
Add			

Hard and soft dependency

Availability of a service may depend on several other services, not just one. The first option is to add all those directly as child services.

However, if some service is already added somewhere else in the services tree, it cannot be simply moved out of there to a child service here. How to create a dependency on it? The answer is "soft" linking. Add the service and mark the *Soft* check box. That way the service can remain in its original location in the tree, yet be depended upon from several other services. Services that are "soft-linked" are displayed in grey in the tree. Additionally, if a service has only "soft" dependencies, it can be deleted directly, without deleting child services first.

The **Time** tab contains the service time specification.

ServiceDependenciesTime

Service times

Type	Interval	Note	Action
No times defined. Work 24x7.			

New service time

Uptime

From

Sunday

Time

hh

:

mm

Till

Sunday

Time

hh

:

mm

[Add](#)

Parameter	Description
<i>Service times</i>	By default, all services are expected to operate 24x7x365. If exceptions needed, add new service times.

Parameter	Description
<i>New service time</i>	Service times: Uptime - service uptime Downtime - service state within this period does not affect SLA. One-time downtime - a single downtime. Service state within this period does not affect SLA. Add the respective hours. <i>Note:</i> Service times affect only the service they are configured for. Thus, a parent service will not take into account the service time configured on a child service (unless a corresponding service time is configured on the parent service as well). Service times are taken into account when calculating IT service status and SLA by the frontend. However, information on service availability is being inserted into database continuously, regardless of service times.

Display To monitor IT services, go to *Monitoring → IT services*.

8. Web monitoring

Overview With Zabbix you can check several availability aspects of web sites.

Attention:

To perform web monitoring Zabbix server must be initially **configured** with cURL (libcurl) support.

To activate web monitoring you need to define web scenarios. A web scenario consists of one or several HTTP requests or "steps". The steps are periodically executed by Zabbix server in a pre-defined order.

The following information is collected in any web scenario:

- average download speed per second for all steps of whole scenario
- failed step number
- last error message

The following information is collected in any web scenario step:

- download speed per second
- response time
- response code

For more details, see *web monitoring items*.

Data collected from executing web scenarios is kept in the database. The data is automatically used for graphs, triggers and notifications.

Zabbix can also check if a retrieved HTML page contains a pre-defined string. It can execute a simulated login and follow a path of simulated mouse clicks on the page.

Zabbix web monitoring supports both HTTP and HTTPS. When running a web scenario, Zabbix always follows redirects. All cookies are preserved during the execution of a single scenario.

Note:

To use HTTP proxy, set the **http_proxy** environment variable for Zabbix server user. For example, `//http_proxy=http:%%/%%proxy_ip:proxy_port//`.

To perform HTTPS checks with the proxy, set the additional **HTTPS_PROXY** environment variable. For example, `//HTTPS_PROXY=http:%%/%%proxy_ip:proxy_port//`.

More details are available by running a shell command `# man curl`.

Configuring a scenario The first step of creating a web scenario is **creating a suitable application** (unless you already have one). Web scenarios are linked to applications for grouping.

Then, to configure a web scenario:

- Go to: *Configuration* → *Web*
- Click on *Create scenario* to the right (or on the scenario name to edit an existing scenario)
- Enter parameters of the scenario in the form

The **Scenario** tab allows you to configure the general parameters of a web scenario.

General parameters:

Parameter	Description
<i>Application</i>	Select an application the scenario will belong to. The application must exist.
<i>Name</i>	Unique scenario name.
<i>Authentication</i>	Authentication options. None - no authentication used. Basic authentication - basic authentication is used. NTLM authentication - NTLM (Windows NT LAN Manager) authentication is used. Selecting an authentication method will provide two additional fields for entering a user name and password.
<i>Update interval (in sec)</i>	How often the scenario will be executed, in seconds.
<i>Agent</i>	Select a client agent. Zabbix will pretend to be the selected browser. This is useful when a website returns different content for different browsers.
<i>Variables</i>	List of variables (macros) for use in scenario steps (URL, Post variables). They have the following format: {macro1} =value1 {macro2} =value2 For example: {username}=Alexei {password}=kj3h5kj34bd The macros can then be referenced in the steps as {username} and {password}. Zabbix will automatically replace them with actual values. <i>Note:</i> Variables are not URL-encoded.

Parameter	Description
<i>Active</i>	The scenario is active if this box is checked, otherwise - disabled.

The **Steps** tab allows you to configure the web scenario steps.

	Name	Timeout	URL	Required	Status codes	Action
1:	Home	15 sec	http://www.google.com		200	Remove
2:	About	15 sec	http://www.google.com/intl/en/about		200	Remove

[Add](#)

Configuring steps To add a web scenario step, click on *Add* in the *Steps* tab of a scenario form.

Step of scenario

Name:

URL:

Post:

Timeout:

Required string:

Required status codes:

Step parameters:

Parameter	Description
<i>Name</i>	Unique step name.
<i>URL</i>	URL to connect to and retrieve data. For example: http://www.zabbix.com https://www.google.com GET variables can be passed in the URL parameter.
<i>Post</i>	HTTP POST variables, if any. For example: id=2345&userid={user} If {user} is defined as a macro of the web scenario, it will be replaced by its value when the step is executed.
<i>Timeout</i>	The information will be sent as is, variables are not URL-encoded. Zabbix will not spend more than the set amount of seconds on processing the URL. Actually this parameter defines maximum time for making connection to the URL and maximum time for performing an HTTP request. Therefore, Zabbix will not spend more than 2 x Timeout seconds on the step. For example: 15
<i>Required</i>	Required regular expressions pattern. Unless retrieved content (HTML) matches required pattern the step will fail. If empty, no check is performed. For example: Homepage of Zabbix Welcome.*admin <i>Note:</i> Referencing regular expressions created in the Zabbix frontend is not supported in this field.


Parameter	Description
Status codes	List of expected HTTP status codes. If Zabbix gets a code which is not in the list, the step will fail. If empty, no check is performed. For example: 200,201,210-299

Note:

Any changes in web scenario steps will only be saved when the whole scenario is saved.

See also a [real-life example](#) of how web monitoring steps can be configured.

Display To view detailed data of defined web scenarios, go to *Monitoring* → *Web* or *Latest data*. Click on the scenario name to see more detailed statistics.



An overview of web monitoring scenarios can be viewed in *Monitoring* → *Dashboard*.

1 Web monitoring items

Overview

Some new items are automatically added for monitoring when web scenarios are created.

Scenario items

As soon as a scenario is created, Zabbix automatically adds the following items for monitoring, linking them to the selected application.

Item	Description
<i>Download speed for scenario</i> <Scenario>	This item will collect information about the download speed (bytes per second) of the whole scenario, i.e. average for all steps. Item key: web.test.in[Scenario,,bps] Type: <i>Numeric(float)</i>
<i>Failed step of scenario</i> <Scenario>	This item will display the number of the step that failed on the scenario. If all steps are executed successfully, 0 is returned. Item key: web.test.fail[Scenario] Type: <i>Numeric(unsigned)</i>
<i>Last error message of scenario</i> <Scenario>	This item returns the last error message text of the scenario. A new value is stored only if the scenario has a failed step. If all steps are ok, no new value is collected. Item key: web.test.error[Scenario] Type: <i>Character</i>

The actual scenario name will be used instead of "Scenario".

Note:

Web monitoring items are added with a 30 day history and a 90 day trend retention period.

These items can be used to create triggers and define notification conditions.

Example 1

To create a "Web scenario failed" trigger, you can define a trigger expression:

```
{host:web.test.fail[Scenario].last(0)}#0
```

Make sure to replace 'Scenario' with the real name of your scenario.

Example 2

To create a "Web scenario failed" trigger with a useful problem description in the trigger name, you can define a trigger with name:

Web scenario "Scenario" failed: {ITEM.VALUE}

and trigger expression:

```
{host:web.test.error[Scenario].strlen()}>0 and {host:web.test.fail[Scenario].last()}>0
```

Make sure to replace 'Scenario' with the real name of your scenario.

Example 3

To create a "Web application is slow" trigger, you can define a trigger expression:

```
{host:web.test.in[Scenario,,bps].last(0)}<10000
```

Make sure to replace 'Scenario' with the real name of your scenario.

Scenario step items

As soon as a step is created, Zabbix automatically adds the following items for monitoring, linking them to the selected application.

Item	Description
<i>Download speed for step</i> <Step> of scenario <Scenario>	This item will collect information about the download speed (bytes per second) of the step. Item key: web.test.in[Scenario,Step,bps] Type: <i>Numeric(float)</i>
<i>Response time for step</i> <Step> <i>of scenario</i> <Scenario>	This item will collect information about the response time of the step in seconds. Response time is counted from the beginning of the request until all information has been transferred. Item key: web.test.time[Scenario,Step,resp] Type: <i>Numeric(float)</i>
<i>Response code for step</i> <Step> <i>of scenario</i> <Scenario>	This item will collect response codes of the step. Item key: web.test.rspcode[Scenario,Step] Type: <i>Numeric(unsigned)</i>

Actual scenario and step names will be used instead of "Scenario" and "Step" respectively.

Note:

Web monitoring items are added with a 30 day history and a 90 day trend retention period.

These items can be used to create triggers and define notification conditions. For example, to create a "Zabbix GUI login is too slow" trigger, you can define a trigger expression:

```
{zabbix:web.test.time[ZABBIX GUI,Login,resp].last(0)}>3
```

2 Real life scenario

Overview

This section presents a step-by-step real-life example of how web monitoring can be used.

Let's use Zabbix Web monitoring to monitor the web interface of Zabbix. We want to know if it is available, provides the right content and how quickly it works. To do that we also must log in with our user name and password.

Scenario

Step 1

Add a new host application.

Go to *Configuration* → *Hosts*, then click on *Applications* next to the host you want to use for web monitoring. In the application section click on *Create application*.



The screenshot shows the 'Application' configuration form in Zabbix. It has a blue header bar with the title 'Application'. Below the header, there are two input fields: 'Host' with the value 'Zabbix server' and a 'Select' button to its right, and 'Name' with the value 'Zabbix frontend'.


This step is not required if you already have a suitable application. You may also want to create a host if one does not exist.

Step 2

Add a new web scenario.

We will add a scenario to monitor the web interface of Zabbix. The scenario will execute a number of steps.

Go to *Configuration* → *Web*, select the host in the dropdown, then click on *Create scenario*.



The screenshot shows the 'Scenario' configuration form in Zabbix. It has a blue header bar with two tabs: 'Scenario' and 'Steps'. The 'Scenario' tab is active. Below the header, there are several fields: 'Application' with the value 'Zabbix frontend' and a 'Select' button; 'Name' with the value 'Zabbix frontend'; 'Authentication' with a dropdown menu showing 'None'; 'Update interval (in sec)' with a text input field containing '60'; 'Agent' with a dropdown menu showing 'Mozilla Firefox 8.0'; and 'Variables' with a text area containing '{user}=Admin' and '{password}=zabbix'. At the bottom, there is an 'Active' checkbox which is checked.

In the new scenario form, click on *Select* next to the Application field to choose the application we just created.

Note that we also create two macros, {user} and {password}.

Step 3

Define steps for the scenario.

Click on *Add* button in the *Steps* tab to add individual steps.

Web scenario step 1

We start by checking that the first page responds correctly, returns with HTTP response code 200 and contains text "SIA Zabbix".

Step of scenario	
Name	First page
URL	http://localhost/zabbix/index.php
Post	
Timeout	15
Required string	Zabbix SIA
Required status codes	200

When done configuring the step, click *Add*.

Web scenario step 2

We continue by logging in to the Zabbix frontend, and we do so by reusing the macros (variables) we defined on the scenario level, {user} and {password}.

Step of scenario	
Name	Logging in
URL	http://localhost/zabbix/index.php
Post	name={user}&password={password}&enter=Sign in
Timeout	15
Required string	
Required status codes	200

Attention:

Note that Zabbix frontend uses JavaScript redirect when logging in, thus first we must log in, and only in further steps we may check for logged-in features. Additionally, the login step must use full URL to **index.php** file.

All the post variables must be on a single line and concatenated with **&** symbol. Example string for logging into Zabbix frontend:

name=Admin&password=zabbix&enter=Sign in

If using the macros as in this example, login string becomes:

name={user}&password={password}&enter=Sign in

Web scenario step 3

Being logged in, we should now verify the fact. To do so, we check for a string that is only visible when logged in - for example, **Profile** link appears in the upper right corner.

Step of scenario

Name

Login check

URL

http://localhost/zabbix/index.php

Post

Timeout

15

Required string

Profile

Required status codes

200

Web scenario step 4

Now that we have verified that frontend is accessible and we can log in and retrieve logged-in content, we should also log out - otherwise Zabbix database will become polluted with lots and lots of open session records.

Step of scenario

Name

Logging out

URL

http://localhost/zabbix/index.php?reconnect=1

Post

Timeout

15

Required string

Required status codes

200

Complete configuration of steps

A complete configuration of web scenario steps should look like this:

Scenario

Steps

Steps

	Name	Timeout	URL	Required	Status codes	Action
↕ 1:	First page	15 sec	http://localhost/zabbix/index.php	Zabbix SIA	200	Remove
↕ 2:	Logging in	15 sec	http://localhost/zabbix/index.php		200	Remove
↕ 3:	Login check	15 sec	http://localhost/zabbix/index.php	Profile	200	Remove
↕ 4:	Logging out	15 sec	http://localhost/zabbix/index.php?reconnect=1		200	Remove
Add						


Step 4

Save the finished web monitoring scenario.

The list of applications and linked scenarios will appear in *Monitoring* → *Web*:

WEB CHECKS				
		Group	Zabbix servers	Host Zabbix server
Name	Number of steps	State	Last check	Status
Zabbix frontend (1 scenario)				
Zabbix frontend	4	Idle till 29 Mar 2012 09:49:07	29 Mar 2012 09:48:07	OK

Click on the scenario name to see more detailed statistics:



9. Maintenance

Overview You can define maintenance periods for hosts and host groups in Zabbix. There are two maintenance types - with data collection and with no data collection.

To skip receiving notifications during host maintenance, actions should be configured by retaining the default **action condition** 'Maintenance status = not in "maintenance"' - then you should not get problem notifications during maintenance. If a problem happens during the maintenance and is not resolved, a notification is generated after the maintenance period ends.

To receive a notification for unresolved problems during the maintenance you have to remove the default action condition about not taking actions during maintenance.

Note:

If at least one host (used in the trigger expression) is not in maintenance mode, Zabbix will send a problem notification.

Zabbix server must be running during maintenance. Timer processes are responsible for switching host status to/from maintenance at 0 seconds of every minute. A proxy will always collect data regardless of the maintenance type (including "no data" maintenance). The data is later ignored by the server if 'no data collection' is set.

When "no data" maintenance ends, triggers using `nodata()` function will not fire before the next check during the period they are checking.

If a log item is added while a host is in maintenance and the maintenance ends, only new logfile entries since the end of the maintenance will be gathered.

Attention:

To ensure predictable behaviour of recurring maintenance periods (daily, weekly, monthly), it is required to use a common timezone for all parts of Zabbix.

Configuration To configure a maintenance period:

- Go to: *Configuration* → *Maintenance*
- Click on *Create maintenance period* (or on the name of an existing maintenance period)

The **Maintenance** tab contains general maintenance period attributes:

Parameter	Description
<i>Name</i>	Name of the maintenance period.
<i>Maintenance type</i>	Two types of maintenance can be set: With data collection - data will be collected by the server during maintenance No data collection - data will not be collected by the server during maintenance
<i>Active since</i>	The date and time when executing maintenance periods becomes active. <i>Note:</i> Setting this time alone does not activate a maintenance period; for that go to the <i>Periods</i> tab.
<i>Active till</i>	The date and time when executing maintenance periods stops being active.
<i>Description</i>	Description of maintenance period.

The **Periods** tab allows you to define the exact days and hours when the maintenance takes place. Clicking on *New* will give you a flexible *New maintenance period* form where you can define the times - for daily, weekly, monthly or one-time maintenance.

Daily and weekly periods have an *Every day/Every week* parameter, which defaults to 1. Setting it to 2 would make the maintenance take place every second day or every second week and so on. The starting day or week depends on the *Active since* time and a maintenance period will start the second day/second week after the day/week of *Active since* time.

For example, having *Active since* set to 2013-09-06 12:00 and an hour long daily recurrent period every 2nd day at 01:00 will result in the first maintenance period starting on 2013-09-07 at 01:00, while the second maintenance period will start on 2013-09-09 at 01:00. Or, with the same *Active since* time and an hour long daily recurrent period every 2nd day at 23:00, the first maintenance period will start on 2013-09-07 at 23:00, and the second maintenance period on 2013-09-09 at 23:00.

The **Hosts & Groups** tab allows you to select the hosts and host groups for maintenance.

Display Hosts in maintenance are displayed accordingly in:

- *Monitoring* → *Dashboard* and *Monitoring* → *Triggers* - a round **orange icon** with a white wrench (since Zabbix 2.0.3 version) indicates that the host is in maintenance. Maintenance details are displayed when the mouse pointer is positioned over the icon. (Previously, before Zabbix 2.0.3, in the dashboard the names of hosts in maintenance were displayed in orange and, in *Monitoring* → *Triggers*, with an orange string displayed next to host name indicating the maintenance type.)

Note:

The display of hosts in maintenance in Dashboard can be unset altogether with the dashboard filtering function.

- in the host list - the *Status* column displays 'In maintenance'
- in network maps - hosts in maintenance have orange background

10. Regular expressions

Overview [POSIX extended regular expressions](#) are supported in Zabbix.

There are two ways of using regular expressions in Zabbix:

- manually entering a regular expression
- using a global regular expression created in Zabbix

Regular expressions You may manually enter a regular expression in supported places. Note that the expression may not start with @ because that symbol is used in Zabbix for referencing global regular expressions.

Global regular expressions There is an advanced editor for creating and testing complex regular expressions in Zabbix frontend.

Once a regular expression has been created this way, it can be used in several places in the frontend by referring to its name, prefixed with @, for example, @mycustomregex.

To create a global regular expression:

- Go to: *Administration* → *General*
- Select *Regular expressions* from the dropdown
- Click on *New regular expression*

Regular expression		
Name	Network interfaces for discovery	
Test string	eth0	
Result	Expression	Expected result
	^lo\$	Result is FALSE
	^Software Loopback Interface	Result is FALSE
	Combined result	TRUE

Expressions			
☐	Expression	Expected result	Case sensitive
☐	^lo\$	Result is FALSE	Yes
☐	^Software Loopback Interface	Result is FALSE	Yes

Set the regular expression name. Any Unicode characters are allowed in the name.

Attention:

A custom regular expression name in Zabbix may contain commas, spaces, etc. In those cases where that may lead to misinterpretation when referencing (for example, a comma in the parameter of an item key) the whole reference may be put in quotes like this: "@My custom regexp for purpose1, purpose2".

Regular expression names must not be quoted in other locations (for example, in LLD rule properties).

To add a new subexpression, click on *New* in the Expressions block.

A custom regular expression may consist of multiple subexpressions, and it can be tested by providing a test string. Results show status of each subexpression and total custom expression status.

11. Event acknowledgment

Overview Problem events in Zabbix can be acknowledged by users.

If a user gets notified about of a problem event, they can go to Zabbix frontend, navigate from events to the acknowledgment screen and acknowledge the problem. When acknowledging, they can enter their comment for it, saying that they are working on it or whatever else they may feel like saying about it.

This way, if another system user spots the same problem, they immediately see if it has been acknowledged and the comments so far.

This way the workflow of resolving problems with more than one system user can take place in a more coordinated way.


Acknowledgment status is also used when defining **action operations**. You can define, for example, that a notification is sent to a higher level manager only if an event is not acknowledged for some time.

To acknowledge events, a user must have at least read permission to the corresponding trigger.

Acknowledgment screen The acknowledgment status of problems is displayed in *Monitoring → Events*.

The *Ack* column contains either a 'Yes' or a 'No', indicating an acknowledged or an unacknowledged problem respectively. A 'Yes' may also have a number with it in brackets, indicating the number of comments for the problem so far.

Both 'Yes' and 'No' are links. Clicking them will take you to the acknowledgment screen.



To acknowledge a problem, enter your comment and click on *Acknowledge and return* or simply *Acknowledge*. 'Acknowledge and return' will take you back to the event screen.

Any previous comments for the problem are displayed above the comment area.

Display Acknowledgment information is fully displayed in the event details accessible by clicking the time of event in *Monitoring → Events*.

Acknowledgment status is displayed in the *Last 20 issues* block of *Monitoring → Dashboard*.

Based on acknowledgment information it is possible to configure how the problem count is displayed in the dashboard or maps. To do that, you have to make selections in the *Problem display* option, available in both map configuration and the dashboard filter. It is possible to display all problem count, unacknowledged problem count as separated from the total or unacknowledged problem count only.

Acknowledgment status is displayed in *Monitoring → Triggers*. There, acknowledgment status is also used with the trigger filtering options. You can filter by unacknowledged triggers or triggers with the last event unacknowledged.

12. Configuration export/import

Overview Zabbix export/import functionality makes it possible to exchange various configuration entities between one Zabbix system and another.

Typical use cases for this functionality:

- sharing of templates or network maps - Zabbix users may share their configuration parameters

- integration with third-party tools - the universal XML format makes integration and data import/export possible with third party tools and applications.

What can be exported/imported

Objects that can be exported/imported are:

- host groups (*through Zabbix API only*)
- templates (including all directly attached items, triggers, graphs, screens, discovery rules and template linkage)
- hosts (including all directly attached items, triggers, graphs, discovery rules and template linkage)
- network maps (including all related images; map export/import is supported since *Zabbix 1.8.2*)
- images
- screens

Export format

Data can be exported using the Zabbix web frontend or [Zabbix API](#). Supported export formats are:

- XML - in the frontend
- XML or JSON - in Zabbix API

Details about export

- All supported elements are exported in one file.
- Host and template entities (items, triggers, graphs, discovery rules) that are inherited from linked templates are not exported. Any changes made to those entities on a host level (such as changed item interval, modified regular expression or added prototypes to the low-level discovery rule) will be lost when exporting; when importing, all entities from linked templates are re-created as on the original linked template.
- Entities created by low-level discovery and any entities depending on them are not exported. For example, a trigger created for an LLD-rule generated item will not be exported.
- Triggers and graphs that use web items are not exported.

Details about import

- Import stops at the first error.
- When updating existing images during image import, "imagetype" field is ignored, i.e. it is impossible to change image type via import.
- Empty tags for items, triggers, graphs, host/template applications, discoveryRules, itemPrototypes, triggerPrototypes, graph-Prototypes are meaningless i.e. it's the same as if it was missing. Other tags, for example, item applications, are meaningful i.e. empty tag means no applications for item, missing tag means don't update applications.
- Import supports both XML and JSON, the import file must have a correct file extension: .xml for XML and .json for JSON.
- See [compatibility information](#) about supported XML versions.

Attention:

Before 2.0.1 import processed different empty tag syntax differently.

```
<?xml version="1.0" encoding="UTF-8"?>
<zabbix_export>
  <version>2.0</version>
  <date>2012-04-18T14:07:36Z</date>
</zabbix_export>
```

XML base format

```
<?xml version="1.0" encoding="UTF-8"?>
```

Default header for XML documents.

```
<zabbix_export>
```

Root element for Zabbix XML export.

```
<version>2.0</version>
```

Export version.

```
<date>2012-04-18T14:07:36Z</date>
```

Date when export was created in ISO 8601 long format.

Other tags are dependent on exported objects.

Groups

Frontend can export groups only with hosts or templates. When host or template is exported all groups it belongs to are exported with it automatically.

API allows to export groups independently from hosts or templates.

```
<groups>
  <group>
    <name>Zabbix servers</name>
  </group>
</groups>
```

groups/group

Parameter	Type	Description	Details
name	<i>string</i>	Group name.	

Hosts

Hosts are exported with many related objects and object relations.

Host export contains:

- hosts data
- host inventory data
- groups relations
- templates relations
- interfaces
- macros
- applications
- items
- discovery rules with all prototypes

When host is imported and updated, it can only be linked to additional templates and never be unlinked from any.

```
<hosts>
  <host>
    <host>Zabbix server</host>
    <name>Zabbix server</name>
    <proxy/>
    <status>0</status>
    <ipmi_authtype>-1</ipmi_authtype>
    <ipmi_privilege>2</ipmi_privilege>
    <ipmi_username/>
    <ipmi_password/>
    <templates/>
    <groups>
      <group>
        <name>Zabbix servers</name>
      </group>
    </groups>
    <interfaces>
      <interface>
        <default>1</default>
        <type>1</type>
        <useip>1</useip>
```

```

        <ip>127.0.0.1</ip>
        <dns/>
        <port>20001</port>
        <interface_ref>if1</interface_ref>
    </interface>
</interfaces>
<applications>
    <application>
        <name>Memory</name>
    </application>
    <application>
        <name>Zabbix agent</name>
    </application>
</applications>
<items>
    <item>
        <name>Agent ping</name>
        <type>0</type>
        <snmp_community/>
        <multiplier>0</multiplier>
        <snmp_oid/>
        <key>agent.ping</key>
        <delay>60</delay>
        <history>7</history>
        <trends>365</trends>
        <status>0</status>
        <value_type>3</value_type>
        <allowed_hosts/>
        <units/>
        <delta>0</delta>
        <snmpv3_securityname/>
        <snmpv3_securitylevel>0</snmpv3_securitylevel>
        <snmpv3_authpassphrase/>
        <snmpv3_privpassphrase/>
        <formula>1</formula>
        <delay_flex/>
        <params/>
        <ipmi_sensor/>
        <data_type>0</data_type>
        <authtype>0</authtype>
        <username/>
        <password/>
        <publickey/>
        <privatekey/>
        <port/>
        <description>The agent always returns 1 for this item. It could be used in combination
        <inventory_link>0</inventory_link>
        <applications>
            <application>
                <name>Zabbix agent</name>
            </application>
        </applications>
        <valuemap>
            <name>Zabbix agent ping status</name>
        </valuemap>
        <interface_ref>if1</interface_ref>
    </item>
    <item>
        <name>Available memory</name>
        <type>0</type>
        <snmp_community/>
        <multiplier>0</multiplier>

```

```

        <snmp_oid/>
        <key>vm.memory.size[available]</key>
        <delay>60</delay>
        <history>7</history>
        <trends>365</trends>
        <status>0</status>
        <value_type>3</value_type>
        <allowed_hosts/>
        <units>B</units>
        <delta>0</delta>
        <snmpv3_securityname/>
        <snmpv3_securitylevel>0</snmpv3_securitylevel>
        <snmpv3_authpassphrase/>
        <snmpv3_privpassphrase/>
        <formula>1</formula>
        <delay_flex/>
        <params/>
        <ipmi_sensor/>
        <data_type>0</data_type>
        <authtype>0</authtype>
        <username/>
        <password/>
        <publickey/>
        <privatekey/>
        <port/>
        <description>Available memory is defined as free+cached+buffers memory.</description>
        <inventory_link>0</inventory_link>
        <applications>
            <application>
                <name>Memory</name>
            </application>
        </applications>
        <valuemap/>
        <interface_ref>if1</interface_ref>
    </item>
</items>
<discovery_rules>
    <discovery_rule>
        <name>Mounted filesystem discovery</name>
        <type>0</type>
        <snmp_community/>
        <snmp_oid/>
        <key>vfs.fs.discovery</key>
        <delay>3600</delay>
        <status>0</status>
        <allowed_hosts/>
        <snmpv3_securityname/>
        <snmpv3_securitylevel>0</snmpv3_securitylevel>
        <snmpv3_authpassphrase/>
        <snmpv3_privpassphrase/>
        <delay_flex/>
        <params/>
        <ipmi_sensor/>
        <authtype>0</authtype>
        <username/>
        <password/>
        <publickey/>
        <privatekey/>
        <port/>
        <filter>{#FSTYPE}:@File systems for discovery</filter>
        <lifetime>30</lifetime>
        <description>Discovery of file systems of different types as defined in global regular

```

```

<item_prototypes>
  <item_prototype>
    <name>Free disk space on $1</name>
    <type>0</type>
    <snmp_community/>
    <multiplier>0</multiplier>
    <snmp_oid/>
    <key>vfs.fs.size[{#FSNAME},free]</key>
    <delay>60</delay>
    <history>7</history>
    <trends>365</trends>
    <status>0</status>
    <value_type>3</value_type>
    <allowed_hosts/>
    <units>B</units>
    <delta>0</delta>
    <snmpv3_securityname/>
    <snmpv3_securitylevel>0</snmpv3_securitylevel>
    <snmpv3_authpassphrase/>
    <snmpv3_privpassphrase/>
    <formula>1</formula>
    <delay_flex/>
    <params/>
    <ipmi_sensor/>
    <data_type>0</data_type>
    <authtype>0</authtype>
    <username/>
    <password/>
    <publickey/>
    <privatekey/>
    <port/>
    <description/>
    <inventory_link>0</inventory_link>
    <applications>
      <application>
        <name>Filesystems</name>
      </application>
    </applications>
    <valuemap/>
    <interface_ref>if1</interface_ref>
  </item_prototype>
</item_prototypes>
<trigger_prototypes>
  <trigger_prototype>
    <expression>{Zabbix server 2:vfs.fs.size[{#FSNAME},pfree].last(0)}<20</expression>
    <name>Free disk space is less than 20% on volume {#FSNAME}</name>
    <url/>
    <status>0</status>
    <priority>2</priority>
    <description/>
    <type>0</type>
  </trigger_prototype>
</trigger_prototypes>
<graph_prototypes>
  <graph_prototype>
    <name>Disk space usage {#FSNAME}</name>
    <width>600</width>
    <height>340</height>
    <yaxismin>0.0000</yaxismin>
    <yaxismax>0.0000</yaxismax>
    <show_work_period>0</show_work_period>
    <show_triggers>0</show_triggers>
  </graph_prototype>
</graph_prototypes>

```

```

        <type>2</type>
        <show_legend>1</show_legend>
        <show_3d>1</show_3d>
        <percent_left>0.0000</percent_left>
        <percent_right>0.0000</percent_right>
        <ymin_type_1>0</ymin_type_1>
        <ymax_type_1>0</ymax_type_1>
        <ymin_item_1>0</ymin_item_1>
        <ymax_item_1>0</ymax_item_1>
        <graph_items>
            <graph_item>
                <sortorder>0</sortorder>
                <drawtype>0</drawtype>
                <color>C80000</color>
                <yaxisside>0</yaxisside>
                <calc_fnc>2</calc_fnc>
                <type>2</type>
                <item>
                    <host>Zabbix server 2</host>
                    <key>vfs.fs.size[{#FSNAME},total]</key>
                </item>
            </graph_item>
            <graph_item>
                <sortorder>1</sortorder>
                <drawtype>0</drawtype>
                <color>00C800</color>
                <yaxisside>0</yaxisside>
                <calc_fnc>2</calc_fnc>
                <type>0</type>
                <item>
                    <host>Zabbix server 2</host>
                    <key>vfs.fs.size[{#FSNAME},free]</key>
                </item>
            </graph_item>
        </graph_items>
    </graph_prototype>
</graph_prototypes>
    <interface_ref>if1</interface_ref>
</discovery_rule>
</discovery_rules>
<macros>
    <macro>
        <macro>{$M1}</macro>
        <value>m1</value>
    </macro>
    <macro>
        <macro>{$M2}</macro>
        <value>m2</value>
    </macro>
</macros>
<inventory/>
</host>
</hosts>

```

hosts/host

Parameter	Type	Description	Details
host	<i>string</i>	Host name.	
name	<i>string</i>	Visible host name.	
status	<i>int</i>	Host Status.	
proxy	<i>int</i>	Proxy name.	
ipmi_authtype	<i>int</i>	IPMI authentication type.	

Parameter	Type	Description	Details
ipmi_privilege	<i>int</i>	IPMI privilege.	
ipmi_username	<i>string</i>	IPMI username.	
ipmi_password	<i>string</i>	IPMI password.	

hosts/host/groups/group

Parameter	Type	Description	Details
name	<i>string</i>	Group name.	

hosts/host/templates/template

Parameter	Type	Description	Details
name	<i>string</i>	Template technical name.	

hosts/host/interfaces/interface

Column name	Type	Description
default	integer	Interface status: 0 - Not default interface 1 - Default interface
type	integer	Interface type: 1 - agent 2 - SNMP 3 - IPMI 4 - JMX
useip	integer	How to connect to the host: 0 - connect to the host using DNS name 1 - connect to the host using IP address
ip	varchar	IP address, can be either IPv4 or IPv6.
dns	varchar	DNS name.
port	varchar	Port number.
interface_ref	varchar	Interface reference name to be used in items.

hosts/host/applications/application

Parameter	Type	Description	Details
name	<i>string</i>	Application name.	

hosts/host/items/item

Parameter	Type	Description
type	<i>int</i>	Item type: 0 - Zabbix agent 1 - SNMPv1 2 - Trapper 3 - Simple check 4 - SNMPv2 5 - Internal 6 - SNMPv3 7 - Active check 8 - Aggregate 9 - HTTP test (web monitoring scenario step) 10 - External 11 - Database monitor

Parameter	Type	Description
		12 - IPMI 13 - SSH 14 - telnet 15 - Calculated 16 - JMX 17 - SNMP trap
snmp_community	string	SNMP Community name
snmp_oid	string	SNMP OID
port	int	Item custom port
name	string	Item name
key	string	Item key
delay	int	Check interval
history	int	How long to keep item history (days)
trends	int	How long to keep item trends (days)
status	int	Item status
value_type	int	Value type
trapper_hosts	string	
units	string	Value units
multiplier	int	Value multiplier
delta	int	Store values as delta
snmpv3_securityname	string	SNMPv3 security name
snmpv3_securitylevel	int	SNMPv3 security level
snmpv3_authpassphrase	string	SNMPv3 authentication phrase
snmpv3_privpassphrase	string	SNMPv3 private phrase
formula	string	
delay_flex	string	Flexible delay
params	string	
ipmi_sensor	string	IPMI sensor
data_type	int	
authtype	int	
username	string	
password	string	
publickey	string	
privatekey	string	
interface_ref	varchar	Reference to host interface
description	string	Item description
inventory_link	int	Host inventory field number, that will be updated with the value returned by the item
applications		Item applications
valuemap		Value map assigned to item

hosts/host/items/item/applications/application

Parameter	Type	Description	Details
name	string	Application name.	

13. Discovery

Please use the sidebar to access content in the Discovery section.

1 Network discovery

Overview

Zabbix offers automatic network discovery functionality that is effective and very flexible.

With network discovery properly set up you can:

- speed up Zabbix deployment
- simplify administration
- use Zabbix in rapidly changing environments without excessive administration

Zabbix network discovery is based on the following information:

- IP ranges
- Availability of external services (FTP, SSH, WEB, POP3, IMAP, TCP, etc)
- Information received from Zabbix agent
- Information received from SNMP agent

It does NOT provide:

- Discovery of network topology

Network discovery basically consists of two phases: discovery and actions.

Discovery

Zabbix periodically scans the IP ranges defined in **network discovery rules**. The frequency of the check is configurable for each rule individually.

Note that one discovery rule will always be processed by a single discoverer process. The IP range will not be split between multiple discoverer processes.

Each rule has a set of service checks defined to be performed for the IP range.

Note:

Discovery checks are processed independently from the other checks. If any checks do not find a service (or fail), other checks will still be processed.

Every check of a service and a host (IP) performed by the network discovery module generates a discovery event.

Event	Check of service result
<i>Service Discovered</i>	The service is 'up' after it was 'down' or when discovered for the first time.
<i>Service Up</i>	The service is 'up', consecutively.
<i>Service Lost</i>	The service is 'down' after it was 'up'.
<i>Service Down</i>	The service is 'down', consecutively.
<i>Host Discovered</i>	At least one service of a host is 'up' after all services of that host were 'down'.
<i>Host Up</i>	At least one service of a host is 'up', consecutively.
<i>Host Lost</i>	All services of a host are 'down' after at least one was 'up'.
<i>Host Down</i>	All services of a host are 'down', consecutively.

Actions

Discovery events can be the basis of relevant **actions**, such as:

- Sending notifications
- Adding/removing hosts
- Enabling/disabling hosts
- Adding hosts to a group
- Removing hosts from a group
- Linking hosts to/unlinking from a template
- Executing remote scripts

These actions can be configured with respect to the device type, IP, status, uptime/downtime, etc. For full details on configuring actions for network-discovery based events, see action **operation** and **conditions** pages.

Host creation

A host is added if the *Add host* operation is selected. A host is also added, even if the *Add host* operation is missing, if you select operations resulting in actions on a host. Such operations are:

- enable host
- disable host
- add host to a host group
- link template to a host

When adding hosts, a host name is the result of reverse DNS lookup or IP address if reverse lookup fails. Lookup is performed from the Zabbix server or Zabbix proxy, depending on which is doing the discovery. If lookup fails on the proxy, it is not retried on the server. If the host with such a name already exists, the next host would get `_2` appended to the name, then `_3` and so on.

Created hosts are added to the *Discovered hosts* group (by default, configurable in *Administration* → *General* → *Other*). If you wish hosts to be added to another group, add a *Remove from host groups* operation (specifying "Discovered hosts") and also add an *Add to host groups* operation (specifying another host group), because a host must belong to a host group.

If a host already exists with the discovered IP address, a new host is not created. However, if the discovery action contains operations (link template, add to host group, etc), they are performed on the existing host.

Interface creation when adding hosts

When hosts are added as a result of network discovery, they get interfaces created according to these rules:

- the services detected - for example, if an SNMP check succeeded, an SNMP interface will be created
- if a host responded both to Zabbix agent and SNMP requests, both types of interfaces will be created
- if uniqueness criteria are Zabbix agent or SNMP-returned data, the first interface found for a host will be created as the default one. Other IP addresses will be added as additional interfaces.
- if a host responded to agent checks only, it will be created with an agent interface only. If it would start responding to SNMP later, additional SNMP interfaces would be added.
- if 3 separate hosts were initially created, having been discovered by the "IP" uniqueness criteria, and then the discovery rule is modified so that hosts A, B and C have identical uniqueness criteria result, B and C are created as additional interfaces for A, the first host. The individual hosts B and C remain. In *Monitoring* → *Discovery* the added interfaces will be displayed in the "Discovered device" column, in black font and indented, but the "Monitored host" column will only display A, the first created host. "Uptime/Downtime" is not measured for IPs that are considered to be additional interfaces.

Configuring a network discovery rule

Overview

To configure a network discovery rule used by Zabbix to discover hosts and services:

- Go to *Configuration* → *Discovery*
- Click on *Create rule* (or on the rule name to edit an existing one)
- Edit the discovery rule attributes

Rule attributes

Discovery rule

Name

Local network

Discovery by proxy

No proxy

IP range

192.168.1.1-100

Delay (in sec)

3600

Checks

ICMP ping

Remove

SNMPv2 agent ".1.3.6.1.2.1.1.0"

Remove

Zabbix agent "system.uname"

Remove

New

Device uniqueness criteria

☒ IP address

☐ SNMPv2 agent ".1.3.6.1.2.1.1.0"

☐ Zabbix agent "system.uname"

Enabled

☒

Save

Clone

Delete

Cancel

Parameter	Description
<i>Name</i>	Unique name of the rule. For example, "Local network".
<i>Discovery by proxy</i>	What performs discovery: no proxy - Zabbix server is doing discovery <proxy name> - this proxy performs discovery
<i>IP range</i>	The range of IP addresses for discovery. It may have the following formats: Single IP: 192.168.1.33 Range of IP addresses: 192.168.1.1-255 IP mask: 192.168.4.0/24 supported IP masks: /16 - /30 for IPv4 addresses /112 - /128 for IPv6 addresses List: 192.168.1.1-255,192.168.2.1-100,192.168.2.200,192.168.4.0/24 Note: Each IP address should be included only once; having multiple rules for a single IP address can have unexpected behavior such as having deadlocks and/or duplicate hosts in the database. The same could happen if two hosts having the same DNS name are included in separate discovery rules.
<i>Delay (seconds)</i>	This parameter defines how often Zabbix will execute the rule. Delay is measured after the execution of previous discovery instance ends so there is no overlap.

Parameter	Description
Checks	Zabbix will use this list of checks for discovery. Supported checks: SSH, LDAP, SMTP, FTP, HTTP, POP, NNTP, IMAP, TCP, Zabbix agent, SNMPv1 agent, SNMPv2 agent, SNMPv3 agent, ICMP ping. A protocol-based discovery uses the net.tcp.service[] functionality to test each host, except for SNMP which queries an SNMP OID. Zabbix agent is tested by querying an item. Please see agent items for more details. The 'Ports' parameter may be one of following: Single port: 22 Range of ports: 22-45 List: 22-45,55,60-70
Device uniqueness criteria	Uniqueness criteria may be: IP address - no processing of multiple single-IP devices. If a device with the same IP already exists it will be considered already discovered and a new host will not be added. Type of discovery check - either SNMP or Zabbix agent check.
Status	Active - the rule is active and will be executed by Zabbix server Disabled - the rule is not active. It won't be executed.

A real life scenario

In this example we would like to set up network discovery for the local network having an IP range of 192.168.1.1-192.168.1.255.

In our scenario we want to:

- discover those hosts that have Zabbix agent running
- run discovery every 10 minutes
- add a host to monitoring if the host uptime is more than 1 hour
- remove hosts if the host downtime is more than 24 hours
- add Linux hosts to the "Linux servers" group
- add Windows hosts to the "Windows servers" group
- use *Template_Linux* for Linux hosts
- use *Template_Windows* for Windows hosts

Step 1

Defining a network discovery rule for our IP range.

Discovery rule

Name

Local network

Discovery by proxy

No proxy

IP range

192.168.1.1-255

Delay (seconds)

300

Checks

Zabbix agent "system.uname"

Remove

New

Device uniqueness criteria

☒ IP address

☐ Zabbix agent "system.uname"

Status

Enabled

Zabbix will try to discover hosts in the IP range of 192.168.1.1-192.168.1.255 by connecting to Zabbix agents and getting the value from **system.uname** key. The value received from the agent can be used to apply different actions for different operating systems. For example, link Windows servers to *Template_Windows*, Linux servers to *Template_Linux*.

The rule will be executed every 10 minutes (600 seconds).

With this rule is added, Zabbix will automatically start the discovery and generating discovery-based events for further processing.

Step 2

Defining an **action** for adding the discovered Linux servers to the respective group/template.

Action	Conditions	Operations															
Type of calculation AND / OR (A) and (B) and (C) and (D)																	
Conditions	<table border="1"><thead><tr><th>Label</th><th>Name</th><th>Action</th></tr></thead><tbody><tr><td>(A)</td><td>Received value like "Linux"</td><td>Remove</td></tr><tr><td>(B)</td><td>Uptime/Downtime >= "3600"</td><td>Remove</td></tr><tr><td>(C)</td><td>Discovery status = "Up"</td><td>Remove</td></tr><tr><td>(D)</td><td>Service type = "Zabbix agent"</td><td>Remove</td></tr></tbody></table>		Label	Name	Action	(A)	Received value like "Linux"	Remove	(B)	Uptime/Downtime >= "3600"	Remove	(C)	Discovery status = "Up"	Remove	(D)	Service type = "Zabbix agent"	Remove
Label	Name	Action															
(A)	Received value like "Linux"	Remove															
(B)	Uptime/Downtime >= "3600"	Remove															
(C)	Discovery status = "Up"	Remove															
(D)	Service type = "Zabbix agent"	Remove															
New condition	<table border="1"><tr><td>Service type</td><td>=</td><td>Zabbix agent</td></tr><tr><td colspan="3">Add</td></tr></table>		Service type	=	Zabbix agent	Add											
Service type	=	Zabbix agent															
Add																	

The action will be activated if:

- the "Zabbix agent" service is "up"
- the value of system.uname (the Zabbix agent key we used in rule definition) contains "Linux"
- Uptime is more than 1 hour (3600 seconds)

Action	Conditions	Operations								
Action operations	<table border="1"><thead><tr><th>Details</th><th>Action</th></tr></thead><tbody><tr><td>Add to host groups: Linux servers</td><td>Edit Remove</td></tr><tr><td>Link to templates: Template OS Linux</td><td>Edit Remove</td></tr><tr><td colspan="2">New</td></tr></tbody></table>		Details	Action	Add to host groups: Linux servers	Edit Remove	Link to templates: Template OS Linux	Edit Remove	New	
Details	Action									
Add to host groups: Linux servers	Edit Remove									
Link to templates: Template OS Linux	Edit Remove									
New										

The action will execute the following operations:

- add the discovered host to the "Linux servers" group (and also add host if it wasn't added previously)
- link host to the "Template_Linux" template. Zabbix will automatically start monitoring the host using items and triggers from "Template_Linux".

Step 3

Defining an action for adding the discovered Windows servers to the respective group/template.

Action	Conditions	Operations															
Type of calculation AND / OR (A) and (B) and (C) and (D)																	
Conditions	<table border="1"><thead><tr><th>Label</th><th>Name</th><th>Action</th></tr></thead><tbody><tr><td>(A)</td><td>Received value like "Windows"</td><td>Remove</td></tr><tr><td>(B)</td><td>Uptime/Downtime >= "3600"</td><td>Remove</td></tr><tr><td>(C)</td><td>Discovery status = "Up"</td><td>Remove</td></tr><tr><td>(D)</td><td>Service type = "Zabbix agent"</td><td>Remove</td></tr></tbody></table>		Label	Name	Action	(A)	Received value like "Windows"	Remove	(B)	Uptime/Downtime >= "3600"	Remove	(C)	Discovery status = "Up"	Remove	(D)	Service type = "Zabbix agent"	Remove
Label	Name	Action															
(A)	Received value like "Windows"	Remove															
(B)	Uptime/Downtime >= "3600"	Remove															
(C)	Discovery status = "Up"	Remove															
(D)	Service type = "Zabbix agent"	Remove															
New condition	<table border="1"><tr><td>Service type</td><td>=</td><td>Zabbix agent</td></tr><tr><td colspan="3">Add</td></tr></table>		Service type	=	Zabbix agent	Add											
Service type	=	Zabbix agent															
Add																	

Action	Conditions	Operations
Action operations		
Details Add to host groups: Windows servers Link to templates: Template OS Windows New		Action Edit Remove Edit Remove

Step 4

Defining an action for removing lost servers.

Action	Conditions	Operations												
Type of calculation AND / OR (A) and (B) and (C)														
Conditions														
	<table border="1"> <thead> <tr> <th>Label</th> <th>Name</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>(A)</td> <td>Uptime/Downtime >= "86400"</td> <td>Remove</td> </tr> <tr> <td>(B)</td> <td>Discovery status = "Down"</td> <td>Remove</td> </tr> <tr> <td>(C)</td> <td>Service type = "Zabbix agent"</td> <td>Remove</td> </tr> </tbody> </table>	Label	Name	Action	(A)	Uptime/Downtime >= "86400"	Remove	(B)	Discovery status = "Down"	Remove	(C)	Service type = "Zabbix agent"	Remove	
Label	Name	Action												
(A)	Uptime/Downtime >= "86400"	Remove												
(B)	Discovery status = "Down"	Remove												
(C)	Service type = "Zabbix agent"	Remove												
New condition														
	Uptime/Downtime <input <input="" type="text" value="600"/>													
Add														

Action	Conditions	Operations
Action operations		
Details Remove host New		Action Edit Remove

A server will be removed if "Zabbix agent" service is 'down' for more than 24 hours (86400 seconds).

2 Active agent auto-registration

Overview

It is possible to allow active Zabbix agent auto-registration, after which the server can start monitoring them. This way new hosts can be added for monitoring without configuring them manually on the server.

Auto registration can happen when a previously unknown active agent asks for checks.

The feature might be very handy for automatic monitoring of new Cloud nodes. As soon as you have a new node in the Cloud Zabbix will automatically start the collection of performance and availability data of the host.

Active agent auto-registration also supports the monitoring of added hosts with passive checks. When the active agent asks for checks, providing it has the 'ListenIP' or 'ListenPort' configuration parameters defined in the configuration file, these are sent along to the server. (If multiple IP addresses are specified, the first one is sent to the server.)

Server, when adding the new auto-registered host, uses the received IP address and port to configure the agent. If no IP address value is received, the one used for the incoming connection is used. If no port value is received, 10050 is used.

Configuration

Configuring active agent auto-registration requires that you set up an **action** for agent auto-registration and have required parameters set in the agent configuration file.

Note:

Setting up **network discovery** is not required to have active agents auto-register.

Action for active agent auto-registration

Go to *Configuration* → *Actions*, select *Auto registration* as the event source and click on *Create action*:

- In the Action tab, give your action a name
- In the Conditions tab, no conditions are required
- In the Operations tab, add relevant operations, such as - 'Add host', 'Add to host groups' (for example, *Discovered hosts*), 'Link to templates', etc.

Note:

If the hosts that will be auto-registering are likely to be supported for active monitoring only (such as hosts that are firewalled from your Zabbix server) then you might want to create a specific template like *Template_Linux-active* to link to.

Agent configuration file

Make sure that you have the Zabbix server identified in **the agent configuration file** - `zabbix_agentd.conf`

`ServerActive=10.0.0.1`

Unless you specifically define a *Hostname* in `zabbix_agentd.conf`, the system hostname of agent location will be used for naming the host. The system hostname in Linux can be obtained by running the 'hostname' command.

Restart the agent after making any changes to the configuration file.

3 Low-level discovery

Overview

Low-level discovery provides a way to automatically create items, triggers, and graphs for different entities on a computer. For instance, Zabbix can automatically start monitoring file systems or network interfaces on your machine, without the need to create items for each file system or network interface manually. Additionally it is possible to configure Zabbix to remove unneeded entities automatically based on actual results of periodically performed discovery.

In Zabbix 2.0, three types of item discovery are supported out of the box:

- discovery of file systems;
- discovery of network interfaces;
- discovery of SNMP OIDs.

A user can define their own types of discovery, provided they follow a particular JSON protocol.

The general architecture of the discovery process is as follows.

First, a user creates a discovery rule in "Configuration" → "Templates" → "Discovery" column. A discovery rule consists of (1) an item that discovers the necessary entities (for instance, file systems or network interfaces) and (2) prototypes of items, triggers, and graphs that should be created based on the value of that item.

An item that discovers the necessary entities is like a regular item seen elsewhere: the server asks a Zabbix agent (or whatever the type of the item is set to) for a value of that item, the agent responds with a textual value. The difference is that the value the agent responds with should contain a list of discovered entities in a specific JSON format. While the details of this format are only important for implementers of custom discovery checks, it is necessary to know that the returned value contains a list of macro → value pairs. For instance, item "net.if.discovery" might return two pairs: "{#IFNAME}" → "lo" and "{#IFNAME}" → "eth0".

Note:

Low-level discovery items "vfs.fs.discovery" and "net.if.discovery" are supported since Zabbix agent version 2.0. Discovery of SNMP OIDs is supported since Zabbix server and proxy version 2.0.

Note:

Return values of a low-level discovery rule are limited to 2048 bytes on a Zabbix proxy run with IBM DB2 database and, before Zabbix proxy 2.0.4, were limited to 4000 characters with Oracle DB. These limits do not apply to Zabbix server as return values are processed without being stored in a database.

These macros are then used in names, keys, and other prototype fields that are basis for creating real items, triggers, and graphs for each discovered entity. These macros can be used:

- for item prototypes in
 - names
 - keys
 - SNMP OIDs
 - calculated item formulas
 - SSH and Telnet scripts
 - database monitor item parameters
- for trigger prototypes in
 - names
 - expressions (insofar as when referencing an item key prototype)
- for graph prototypes in
 - names

When the server receives a value for a discovery item, it looks at the macro → value pairs and for each pair generates real items, triggers, and graphs, based on their prototypes. In the example with "net.if.discovery" above, the server would generate one set of items, triggers, and graphs for the loopback interface "lo", and another set for interface "eth0".

The following sections illustrate the process described above in detail and serve as a how-to for performing discovery of file systems, network interfaces, and SNMP OIDs. The last section describes the JSON format for discovery items and gives an example of how to implement your own file system discoverer as a Perl script.

3.1 Discovery of file systems

To configure the discovery of file systems, do the following:

- Go to: *Configuration → Templates*
- Click on *Discovery* in the row of an appropriate template

TEMPLATES

Displaying 1 to 44 of 44 found

Group

Templates

Create

Import

☐	Templates	Applications	Items	Triggers	Graphs	Screens	Discovery	Linked templates	Linked to
☐	C Template Linux	Applications (12)	Items (102)	Triggers (44)	Graphs (0)	Screens (0)	Discovery (0)	-	-

- Click on *Create discovery rule* in the upper right corner of the screen
- Fill in the form with the following details

Discovery rule

Name

Mounted filesystem discovery

Type

Zabbix agent

Key

vfs.fs.discovery

Select

Update interval (in sec)

3600

Flexible intervals

Interval	Period	Action
No flexible intervals defined.		

New flexible interval

Interval (in sec)

50

Period

1-7,00:00-24:00

Add

Keep lost resources period (in days)

30

Filter

Macro

{#FSTYPE}

Regexp

@File systems for discovery

Description

Discovery of file systems of different types as defined in global regular expression "File systems for discovery".

Status

Enabled

Save

Cancel

Parameter	Description
<i>Name</i>	Name of discovery rule.
<i>Type</i>	The type of check to perform discovery; should be <i>Zabbix agent</i> or <i>Zabbix agent (active)</i> for file system discovery.
<i>Key</i>	An item with "vfs.fs.discovery" key is built into the Zabbix agent on many platforms (see supported item key list for details), and will return a JSON with the list of file systems present on the computer and their types.
<i>Update interval (in sec)</i>	This field specifies how often Zabbix performs discovery. In the beginning, when you are just setting up file system discovery, you might wish to set it to a small interval, but once you know it works you can set it to 30 minutes or more, because file systems usually do not change very often. <i>Note:</i> If set to '0', the item will not be polled. However, if a flexible interval also exists with a non-zero value, the item will be polled during the flexible interval duration.
<i>Flexible intervals</i>	You can create exceptions to <i>Update interval</i>. For example: Interval: 0, Period: 6-7,00:00-24:00 - will disable the polling at the weekend. Otherwise default update interval will be used. Up to seven flexible intervals can be defined. If multiple flexible intervals overlap, the smallest <i>Interval</i> value is used for the overlapping period. See Time period specification page for description of the <i>Period</i> format. <i>Note:</i> If set to '0', the item will not be polled during the flexible interval duration and will resume polling according to the <i>Update interval</i> once the flexible interval period is over.

Parameter	Description
<i>Keep lost resources period (in days)</i>	This field allows you to specify for how many days the discovered entity will be retained (won't be deleted) once its discovery status becomes "Not discovered anymore" (max 3650 days). <i>Note:</i> If set to "0", entities will be deleted immediately. Using "0" is not recommended, since just wrongly editing the filter may end up in the entity being deleted with all the historical data.
<i>Filter</i>	The filter can be used to only generate real items, triggers, and graphs for certain file systems. It expects POSIX Extended Regular Expression . For instance, if you are only interested in C:, D:, and E: file systems, you could put {#FSNAME} into "Macro" and "^C ^D ^E" regular expression into "Regexp" text fields. Filtering is also possible by file system types using {#FSTYPE} macro (e. g. "^ext ^reiserfs"). You can enter a regular expression or reference a global regular expression in "Regexp" field. In order to test the regular expression you can use "grep -E", for example: <pre>for f in ext2 nfs reiserfs smbfs; do echo \$f \\\ grep -E '^ext'</pre> that if some macro from the filter is missing in the response, the found entity will be ignored.
<i>Description</i>	Enter a description.
<i>Status</i>	Enabled - the rule will be processed. Disabled - the rule will not be processed. Not supported - the item is not supported. This item will not be processed, however Zabbix may try to periodically set the status of the item to <i>Enabled</i> according to the interval set for refreshing unsupported items .

Attention:

Zabbix database in MySQL must be created as case-sensitive if file system names that differ only by case are to be discovered correctly.

Note:

Discovery rule history is not preserved.

Once a rule is created, go to the items for that rule and press "Create prototype" to create an item prototype. Note how macro {#FSNAME} is used where a file system name is required. When the discovery rule is processed, this macro will be substituted with the discovered file system.

Item : Free disk space on \$1 (percentage)

Name

Free disk space on \$1 (percentage)

Type

Zabbix agent

Key

vfs.fs.size[#{FSNAME},pfree]

Select

Type of information

Numeric (float)

Units

%

Use custom multiplier

☐

1

Update interval (in sec)

60

Flexible intervals

Interval	Period	Action
No flexible intervals defined.		

New flexible interval

Interval (in sec)

50

Period

1-7,00:00-24:00

Add

Keep history (in days)

7

Keep trends (in days)

365

Store value

As is

Show value

As is

show value mappings

New application

Applications

-None-

CPU

Filesystems

General

Memory

Network interfaces

Description

Enabled

☒

Save

Cancel

Note:

If an item prototype is created with a *Disabled* status, it will be added to a discovered entity, but in a disabled state.

We can create several item prototypes for each file system metric we are interested in:

Item prototypes of Mounted filesystem discovery

Displaying 1 to 5 of 5 found

[« Template list](#)
[Template: Template OS Linux](#)
[« Discovery list](#)
[Discovery: Mounted filesystem discovery](#)
[Item prototypes \(5\)](#)

[Trigger prototypes \(2\)](#)
[Graph prototypes \(1\)](#)

☐	Name	Key	Interval	History	Trends	Type	Status	Applications
☐	Free disk space on {FSNAME}	vfs.fs.size[#{FSNAME},free]	60	7	365	Zabbix agent	Enabled	Filesystems
☐	Free disk space on {FSNAME} (percentage)	vfs.fs.size[#{FSNAME},pfree]	60	7	365	Zabbix agent	Enabled	Filesystems
☐	Free inodes on {FSNAME} (percentage)	vfs.fs.inode[#{FSNAME},pfree]	60	7	365	Zabbix agent	Enabled	Filesystems
☐	Total disk space on {FSNAME}	vfs.fs.size[#{FSNAME},total]	3600	7	365	Zabbix agent	Enabled	Filesystems
☐	Used disk space on {FSNAME}	vfs.fs.size[#{FSNAME},used]	60	7	365	Zabbix agent	Enabled	Filesystems

Then, we create trigger prototypes in a similar way:

Trigger

Name

Free disk space is less than 20% on volume {#FSNAME}

Expression

{Template OS Linux:vfs.fs.size[{#FSNAME}],pfree].last(0)}<20

Add

Expression constructor

Multiple PROBLEM events

☐

generation

Description

URL

Severity

Not classified

Information

Warning

Average

High

Disaster

Enabled

☒

Save

Clone

Delete

Cancel

Trigger prototypes of Mounted filesystem discovery

Displaying 1 to 2 of 2 found

[Hide disabled triggers]

« Template list

Template: Template OS Linux

« Discovery list

Discovery: Mounted filesystem discovery

Item prototypes (5)

Trigger prototypes (2)

Graph prototypes (1)

☐	Severity	Name	Expression	Status
☐	Warning	Free disk space is less than 20% on volume {#FSNAME}	{Template OS Linux:vfs.fs.size[{#FSNAME}],pfree].last(0)}<20	Enabled
☐	Warning	Free inodes is less than 20% on volume {#FSNAME}	{Template OS Linux:vfs.fs.inode[{#FSNAME}],pfree].last(0)}<20	Enabled

And graph prototypes too:

Graph

Preview

Name

Disk space usage {#FSNAME}

Width

600

Height

340

Graph type

Pie

Show legend

☒

3D view

☒

Items

Name	Type	Function	Colour	Action
1: Template OS Linux: Total disk space on {#FSNAME}	Graph sum	avg	C80000	Remove
2: Template OS Linux: Free disk space on {#FSNAME}	Simple	avg	00C800	Remove

Add

Add prototype

Save

Clone

Delete

Cancel

Graph prototypes of Mounted filesystem discovery

Displaying 1 to 1 of 1 found

« Template list

Template: Template OS Linux

« Discovery list

Discovery: Mounted filesystem discovery

Item prototypes (5)

Trigger prototypes (2)

Graph prototypes (1)

☐	Name	Width	Height	Graph type
☐	Disk space usage {#FSNAME}	600	340	Pie

Finally, we have created a discovery rule that looks like shown below. It has five item prototypes, two trigger prototypes, and one

graph prototype.

Discovery rules									
Displaying 1 to 2 of 2 found									
« Template list Template: Template OS Linux Applications (10) Items (32) Triggers (15) Graphs (4) Screens (1) Discovery rules (2)									
☐	Name 	Items	Triggers	Graphs	Key	Interval	Type	Status	Error
☐	Mounted filesystem discovery	Item prototypes (5)	Trigger prototypes (2)	Graph prototypes (1)	vfs.fs.discovery	3600	Zabbix agent	Enabled	

The screenshots below illustrate how discovered items, triggers, and graphs look like in the host's configuration. Discovered entities are prefixed with a golden link to a discovery rule they come from.

Items

Displaying 33 to 64 of 67 found

Filter

« Host list

Host: [Zabbix server](#)

[Monitored](#)

Availability: [Available](#)

[Applications](#) (11)

Items (67)

[Triggers](#) (41)

[Graphs](#) (10)

[Discovery rules](#) (2)

< Previous

1

2

3

Next >

☐	Wizard	Name	Triggers	Key 	Interval	History	Trends	Type	Applications	Status	Error
☐		Template OS Linux: Number of logged in users		system.users.num	60	7	365	Zabbix agent	OS, Security	Enabled	
☐		Template OS Linux: Checksum of /etc/passwd	Triggers (1)	vfs.file.cksum[/etc/passwd]	3600	7	365	Zabbix agent	Security	Enabled	
☐		Mounted filesystem discovery: Free inodes on / (percentage)	Triggers (1)	vfs.fs.inode[/,pfree]	60	7	365	Zabbix agent	Filesystems	Enabled	
☐		Mounted filesystem discovery: Free disk space on /		vfs.fs.size[/,free]	60	7	365	Zabbix agent	Filesystems	Enabled	
☐		Mounted filesystem discovery: Free disk space on / (percentage)	Triggers (1)	vfs.fs.size[/,pfree]	60	7	365	Zabbix agent	Filesystems	Enabled	
☐		Mounted filesystem discovery: Total disk space on /		vfs.fs.size[/,total]	3600	7	365	Zabbix agent	Filesystems	Enabled	
☐		Mounted filesystem discovery: Used disk space on /		vfs.fs.size[/,used]	60	7	365	Zabbix agent	Filesystems	Enabled	

Note that discovered entities will not be created in case there are already existing entities with the same uniqueness criteria, for example, an item with the same key or graph with the same name.

Items (similarly, triggers and graphs) created by a low-level discovery rule cannot be manually deleted. However, they will be deleted automatically if a discovered entity (file system, interface, etc) stops being discovered (or does not pass the filter anymore). In this case the items, triggers and graphs will be deleted after the days defined in the *Keep lost resources period* field pass. Note that triggers (until Zabbix 2.0.11) and graphs (until Zabbix 2.0.12) are deleted immediately.

When discovered entities become 'Not discovered anymore', an orange lifetime indicator is displayed in the item list. Move your mouse pointer over it and a message will be displayed indicating how many days are left until the item is deleted.

Type	Applications	Status	Error
Zabbix agent		Active	 
Close			
The item is not discovered anymore and will be deleted in 3h 22m 3s (on 10 Jan 2012 at 15:25:03).			

If entities were marked for deletion, but were not deleted at the expected time (disabled discovery rule or item host), they will be deleted the next time the discovery rule is processed.

Triggers			Group	Zabbix servers
Displaying 1 to 30 of 41 found				
« Host list Host: Zabbix server Monitored Availability: Available Applications (11) Items (65) Triggers (41) Graphs (10) Discovery rules (2)				
1 2 Next >				
☐	Severity	Name 	Expression	
☐	Warning	Template OS Linux: /etc/passwd has been changed on Zabbix server	[Zabbix server:vfs.file.cksum[/etc/passwd].diff(0)]>0	
☐	Information	Template OS Linux: Configured max number of opened files is too low on Zabbix server	[Zabbix server:kern.el.maxfiles.last(0)]<1024	
☐	Information	Template OS Linux: Configured max number of processes is too low on Zabbix server	[Zabbix server:kern.el.maxproc.last(0)]<256	
☐	Warning	Template OS Linux: Disk I/O is overloaded on Zabbix server	[Zabbix server:system.cpu.util[,iowait].last(0)]>20	
☐	Warning	Mounted filesystem discovery: Free disk space is less than 20% on volume /	[Zabbix server:vfs.fs.size[/,pfree].last(0)]<20	
☐	Warning	Mounted filesystem discovery: Free inodes is less than 20% on volume /	[Zabbix server:vfs.fs.inode[/,pfree].last(0)]<20	

Graphs				
Group Zabbix servers Host Zabbix server				
Displaying 1 to 10 of 10 found				
« Host list Host: Zabbix server Monitored Availability: Available Applications (11) Items (65) Triggers (41) Graphs (10)				
Discovery rules (2)				
☐ Name	Width	Height	Graph type	
☐ Template OS Linux: CPU jumps	900	200	Normal	
☐ Template OS Linux: CPU load	900	200	Normal	
☐ Template OS Linux: CPU utilization	900	200	Stacked	
☐ Mounted filesystem discovery: Disk space usage /	600	340	Pie	

3.2 Discovery of network interfaces

Discovery of network interfaces is done in exactly the same way as discovery of file systems, except that you use the discovery rule key "net.if.discovery" instead of "vfs.fs.discovery" and use macro {#IFNAME} instead of {#FSNAME} in filter and item/trigger/graph prototypes.

Examples of item prototypes that you might wish to create based on "net.if.discovery": "net.if.in[{#IFNAME},bytes]", "net.if.out[{#IFNAME},bytes]".

See above for more information about the filter.

3.3 Discovery of SNMP OIDs

In this example, we will perform SNMP discovery on a switch. First, go to "Configuration" → "Templates".

Templates								
Group Templates								
Displaying 1 to 25 of 25 found								
☐ Templates	Applications	Items	Triggers	Graphs	Screens	Discovery	Linked templates	Linked to
☐ Template App Agentless	Applications (1)	Items (12)	Triggers (12)	Graphs (0)	Screens (0)	Discovery (0)	-	-
☐ Template App MySQL	Applications (1)	Items (14)	Triggers (1)	Graphs (2)	Screens (1)	Discovery (0)	-	-
☐ Template App Zabbix Agent	Applications (1)	Items (3)	Triggers (3)	Graphs (0)	Screens (0)	Discovery (0)	-	Template OS AIX , Template OS FreeBSD , Template OS HP-UX , Template OS Linux , Template OS Mac OS X , Template OS OpenBSD , Template OS Solaris , Template OS Windows
☐ Template App Zabbix Server	Applications (1)	Items (26)	Triggers (24)	Graphs (4)	Screens (1)	Discovery (0)	-	Zabbix server
☐ Template HP Procurve	Applications (0)	Items (0)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (0)	-	-
☐ Template HP Procurve2	Applications (8)	Items (0)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (1)	-	ProCurve J4900B Switch 2626
☐ Template IPMI Intel SR1530	Applications (3)	Items (8)	Triggers (11)	Graphs (2)	Screens (0)	Discovery (0)	-	-

To edit discovery rules for a template, click on the link in the "Discovery" column.

Then, press "Create rule" and fill the form with the details in the screenshot below.

Unlike file system and network interface discovery, the item does not necessarily have to have "snmp.discovery" key - item type of SNMP agent is sufficient.

Also, unlike the previous examples, this discovery item will generate two macros for each discovered entity: {#SNMPINDEX} and {#SNMPVALUE}. In case you would like to filter out loopback interfaces from returned values you could put "{#SNMPVALUE}" into filter "Macro" and "^(^|)|\$|^o?" regular expression into "Regexp" text fields. See above for more information about the filter.

In "SNMP OID" field, we have to put an OID that is capable of generating meaningful values for these macros.

To understand what we mean, let us perform snmpwalk on our switch:

```
$ snmpwalk -v 2c -c public 192.168.1.1 IF-MIB::ifDescr
IF-MIB::ifDescr.1 = STRING: WAN
IF-MIB::ifDescr.2 = STRING: LAN1
IF-MIB::ifDescr.3 = STRING: LAN2
```

Macro {#SNMPINDEX} takes its value from the part of the OID that is after ifDescr (in this example: 1, 2, 3). Macro {#SNMPVALUE} comes from the value of the corresponding OID (here: WAN, LAN1, LAN2). Thus, our "snmp.discovery" item would return three sets of macro → value pairs:

```
{#SNMPINDEX} → 1    {#SNMPVALUE} → WAN
{#SNMPINDEX} → 2    {#SNMPVALUE} → LAN1
{#SNMPINDEX} → 3    {#SNMPVALUE} → LAN2
```

Discovery rule

Name

Interfaces

Type

SNMPv2 agent

Key

snmp.discovery

Select

SNMP OID

ifDescr

SNMP community

public

Port

161

Update interval (in sec)

30

Flexible intervals

Interval	Period	Action
No flexible intervals defined.		

New flexible interval

Interval (in sec)

50

Period

1-7,00:00-24:00

Add

Keep lost resources period (in days)

30

Filter

Macro

Regexp

Description

Status

Enabled

Save

Cancel

The following screenshot illustrates how we can use these macros in item prototypes:

Item :

Name

Type

Key

SNMP OID

SNMP community

Port

Type of information

Units

Use custom multiplier ☐

Update interval (in sec)

Flexible intervals

Interval	Period	Action
No flexible intervals defined.		

New flexible interval

Interval (in sec)	Period	Action
50	1-7,00:00-24:00	Add

Keep history (in days)

Keep trends (in days)

Store value

Show value [show value mappings](#)

New application

Applications

Description

Enabled ☒

Again, creating as many item prototypes as needed:

Item prototypes of Interfaces

Displaying 1 to 8 of 8 found

« [Template list](#) **Template:** [Template HP Procurve](#) « [Discovery list](#) **Discovery:** [Interfaces](#) Item prototypes (8) [Trigger prototypes](#) (0)

[Graph prototypes](#) (0)

☐	Name	Key	Interval	History	Trends	Type	Status	Applications
☐	ifDescr.{#SNMPINDEX}	ifDescr.['{#SNMPINDEX}']	30	7		SNMPv2 agent	Enabled	ifDescr
☐	ifInDiscards.{#SNMPINDEX}	ifInDiscards.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifInDiscards
☐	ifInErrors.{#SNMPINDEX}	ifInErrors.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifInErrors
☐	ifInOctets.{#SNMPINDEX}	ifInOctets.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifInOctets
☐	ifOperStatus.{#SNMPINDEX}	ifOperStatus.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifOperStatus
☐	ifOutDiscards.{#SNMPINDEX}	ifOutDiscards.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifOutDiscards
☐	ifOutErrors.{#SNMPINDEX}	ifOutErrors.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifOutErrors
☐	ifOutOctets.{#SNMPINDEX}	ifOutOctets.['{#SNMPINDEX}']	30	7	365	SNMPv2 agent	Enabled	ifOutOctets

As well as trigger prototypes:

Trigger

Name

ifOperStatus.{#SNMPINDEX} on {HOST.HOST} has changed

Expression

{Template HP Procurve:ifOperStatus["{#SNMPINDEX}"].diff()=1}

Add

Expression constructor

Multiple PROBLEM events generation

Description

URL

Severity

Not classified

Information

Warning

Average

High

Disaster

Enabled

☒

Save

Cancel

Trigger prototypes of Interfaces

Displaying 1 to 2 of 2 found

[Hide disabled triggers]

« Template list

Template: Template HP Procurve

« Discovery list

Discovery: Interfaces

Item prototypes (8)

Trigger prototypes (2)

Graph prototypes (0)

☐	Severity	Name	Expression	Status
☐	Information	#Descr.{#SNMPINDEX} on {HOST.HOST} has changed	{Template HP Procurve:ifDescr["{#SNMPINDEX}"].diff()=1}	Enabled
☐	Warning	ifOperStatus.{#SNMPINDEX} on {HOST.HOST} has changed	{Template HP Procurve:ifOperStatus["{#SNMPINDEX}"].diff()=1}	Enabled

And graph prototypes:

Graph

Preview

Name

Utilization of interface {#SNMPINDEX}

Width

900

Height

100

Graph type

Normal

Show legend

☐

Show working time

☒

Show triggers

☒

Percentile line (left)

☐

Percentile line (right)

☐

Y axis MIN value

Calculated

Y axis MAX value

Calculated

Items

Name	Function	Draw style	Y axis side	Colour	Action
1: Template HP Procurve: ifInOctets.\$1	avg	Line	Left	C80000	Remove
2: Template HP Procurve: ifOutOctets.\$1	avg	Line	Left	00C800	Remove
Add	Add prototype				

Save

Cancel

Graph prototypes of Interfaces			
Displaying 1 to 1 of 1 found			
« Template list Template: Template HP Procurve « Discovery list Discovery: Interfaces Item prototypes (8) Trigger prototypes (2)			
Graph prototypes (1)			
☐	Name ↑	Width	Height
☐	Utilization of interface {#SNMPINDEX}	900	100
			Graph type
			Normal

A summary of our discovery rule:

Discovery rules									
Displaying 1 to 1 of 1 found									
« Template list Template: Template HP Procurve Applications (8) Items (0) Triggers (0) Graphs (0) Screens (0)									
Discovery rules (1)									
☐	Name ↑	Items	Triggers	Graphs	Key	Interval	Type	Status	Error
☐	Interfaces	Item prototypes (8)	Trigger prototypes (2)	Graph prototypes (1)	snmp.discovery	30	SNMPv2 agent	Enabled	✓

When server runs, it will create real items, triggers, and graphs, based on the values “snmp.discovery” returns. In host’s configuration they will be prefixed with a golden link to a discovery rule they come from.

Items											
Displaying 113 to 140 of 241 found											
Filter											
« Host list Host: ProCurve J4900B Switch 2626 Monitored Availability: Unknown Applications (8) Items (241) Triggers (60) Graphs (30)											
Discovery rules (1)											
< Previous 1 2 3 4 5 6 7 8 9 Next >											
☐	Wizard	Name	Triggers	Key ↑	Interval	History	Trends	Type	Applications	Status	Error
☐		Interfaces : ifInOctets.23		ifInOctets.["23"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.24		ifInOctets.["24"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.25		ifInOctets.["25"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.26		ifInOctets.["26"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.63		ifInOctets.["63"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.67		ifInOctets.["67"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.69		ifInOctets.["69"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifInOctets.4158		ifInOctets.["4158"]	30	7	365	SNMPv2 agent	ifInOctets	Enabled	✓
☐		Interfaces : ifOperStatus.1	Triggers (1)	ifOperStatus.["1"]	30	7	365	SNMPv2 agent	ifOperStatus	Enabled	✓
☐		Interfaces : ifOperStatus.2	Triggers (1)	ifOperStatus.["2"]	30	7	365	SNMPv2 agent	ifOperStatus	Enabled	✓
☐		Interfaces : ifOperStatus.3	Triggers (1)	ifOperStatus.["3"]	30	7	365	SNMPv2 agent	ifOperStatus	Enabled	✓
☐		Interfaces : ifOperStatus.4	Triggers (1)	ifOperStatus.["4"]	30	7	365	SNMPv2 agent	ifOperStatus	Enabled	✓
☐		Interfaces : ifOperStatus.5	Triggers (1)	ifOperStatus.["5"]	30	7	365	SNMPv2 agent	ifOperStatus	Enabled	✓

Triggers
Group Switches Host ProCurve J4900B Switch 2626
[\[Hide disabled triggers \]](#)

Displaying 29 to 56 of 60 found

« [Host list](#) **Host:** [ProCurve J4900B Switch 2626](#) **Monitored** **Availability:** Unknown [Applications](#) (8) [Items](#) (241) **Triggers** (60) [Graphs](#) (30)

[Discovery rules](#) (1)

< Previous | 1 | 2 | 3 | Next >

☐	Severity	Name	Expression	Status	Error
☐	Not classified	Interfaces : ifDescr.8 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifDescr["8"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifDescr.9 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifDescr["9"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.1 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["1"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.2 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["2"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.3 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["3"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.4 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["4"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.5 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["5"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.6 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["6"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.10 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["10"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.11 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["11"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.12 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["12"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.13 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["13"].diff()=1	Enabled	✓
☐	Not classified	Interfaces : ifOperStatus.14 on ProCurve J4900B Switch 2626 has changed	{ProCurve J4900B Switch 2626:ifOperStatus["14"].diff()=1	Enabled	✓

Graphs
Group Switches Host ProCurve J4900B Switch 2626

Displaying 1 to 28 of 30 found

« [Host list](#) **Host:** [ProCurve J4900B Switch 2626](#) **Monitored** **Availability:** Unknown [Applications](#) (8) [Items](#) (241) [Triggers](#) (60) **Graphs** (30)

[Discovery rules](#) (1)

1 | 2 | Next >

☐	Name	Width	Height	Graph type
☐	Interfaces : Utilization of interface 1	900	100	Normal
☐	Interfaces : Utilization of interface 2	900	100	Normal
☐	Interfaces : Utilization of interface 3	900	100	Normal
☐	Interfaces : Utilization of interface 4	900	100	Normal
☐	Interfaces : Utilization of interface 5	900	100	Normal
☐	Interfaces : Utilization of interface 6	900	100	Normal
☐	Interfaces : Utilization of interface 7	900	100	Normal
☐	Interfaces : Utilization of interface 8	900	100	Normal
☐	Interfaces : Utilization of interface 9	900	100	Normal
☐	Interfaces : Utilization of interface 10	900	100	Normal
☐	Interfaces : Utilization of interface 11	900	100	Normal
☐	Interfaces : Utilization of interface 12	900	100	Normal
☐	Interfaces : Utilization of interface 13	900	100	Normal

3.4 Creating custom LLD rules

It is also possible to create a completely custom LLD rule, discovering any type of entities - for example, databases on a database server.

To do so, a custom item should be created that returns JSON, specifying found objects and optionally - some properties of them. The amount of macros per entity is not limited - while the built-in discovery rules return either one or two macros (for example, two for filesystem discovery), it is possible to return more.

The required JSON format is best illustrated with an example. Suppose we are running an old Zabbix 1.8 agent (one that does not support "vfs.fs.discovery"), but we still need to discover file systems. Here is a simple Perl script for Linux that discovers mounted file systems and outputs JSON, which includes both file system name and type. One way to use it would be as a **UserParameter** with key "vfs.fs.discovery_perl":

```
#!/usr/bin/perl

$first = 1;

print "{\n";
print "\t\"data\": [\n\n";
```

```

for (`cat /proc/mounts`)
{
    ($fsname, $fstype) = m/\S+ (\S+) (\S+)/;
    $fsname =~ s!/!\!/!g;

    print "\t,\n" if not $first;
    $first = 0;

    print "\t{\n";
    print "\t\t\"{#FSNAME}\" : \"$fsname\" ,\n";
    print "\t\t\"{#FSTYPE}\" : \"$fstype\" \n";
    print "\t}\n";
}

print "\n\t]\n";
print "}\n";

```

Attention:

Allowed symbols for LLD macro names are **0-9** , **A-Z** , **_** , **.**

Lowercase letters are not supported in the names.

An example of its output (reformatted for clarity) is shown below. JSON for custom discovery checks has to follow the same format.

```

{
  "data": [

    { "{#FSNAME}": "\/", "{#FSTYPE}": "rootfs" },
    { "{#FSNAME}": "\/sys", "{#FSTYPE}": "sysfs" },
    { "{#FSNAME}": "\/proc", "{#FSTYPE}": "proc" },
    { "{#FSNAME}": "\/dev", "{#FSTYPE}": "devtmpfs" },
    { "{#FSNAME}": "\/dev/pts", "{#FSTYPE}": "devpts" },
    { "{#FSNAME}": "\/", "{#FSTYPE}": "ext3" },
    { "{#FSNAME}": "\/lib/init/rw", "{#FSTYPE}": "tmpfs" },
    { "{#FSNAME}": "\/dev/shm", "{#FSTYPE}": "tmpfs" },
    { "{#FSNAME}": "\/home", "{#FSTYPE}": "ext3" },
    { "{#FSNAME}": "\/tmp", "{#FSTYPE}": "ext3" },
    { "{#FSNAME}": "\/usr", "{#FSTYPE}": "ext3" },
    { "{#FSNAME}": "\/var", "{#FSTYPE}": "ext3" },
    { "{#FSNAME}": "\/sys/fs/fuse/connections", "{#FSTYPE}": "fusectl" }

  ]
}

```

Then, in the discovery rule's "Filter" field, we could specify "{#FSTYPE}" as a macro and "rootfs|ext3" as a regular expression.

Note:

You don't have to use macro names FSNAME/FSTYPE with custom LLD rules, you are free to use whatever names you like.

14. Distributed monitoring

Overview Zabbix provides effective and reliable ways of monitoring distributed IT infrastructure. The two main solutions for large environments provided by Zabbix are:

- use of **proxies**
- use of **nodes**

Proxies can be used to collect data locally on behalf of a centralized Zabbix server and then report the data to the server. Nodes are full Zabbix servers that can be set up in a hierarchy of distributed monitoring.

Proxy vs. node

When making a choice between using a proxy or a node, several considerations must be taken into account.

	Proxy	Node
<i>Lightweight</i>	Yes	No
<i>GUI</i>	No	Yes
<i>Works independently</i>	Yes	Yes
<i>Easy maintenance</i>	Yes	No
<i>Automatic DB creation¹</i>	Yes	No
<i>Local administration</i>	No	Yes
<i>Ready for embedded hardware</i>	Yes	No
<i>One way TCP connections</i>	Yes	Yes
<i>Centralised configuration</i>	Yes	No
<i>Generates notifications</i>	No	Yes

Note:

[1] Automatic DB creation feature only works with SQLite. Other databases require a **manual setup**.

1 Proxies


Overview

A Zabbix proxy can collect performance and availability data on behalf of the Zabbix server. This way, a proxy can take on itself some of the load of collecting data and offload the Zabbix server.

Also, using a proxy is the easiest way of implementing centralized and distributed monitoring, when all agents and proxies report to one Zabbix server and all data is collected centrally.

A Zabbix proxy can be used to:

- Monitor remote locations
- Monitor locations having unreliable communications
- Offload the Zabbix server when monitoring thousands of devices
- Simplify the maintenance of distributed monitoring



The proxy requires only one TCP connection to the Zabbix server. This way it is easier to get around a firewall as you only need to configure one firewall rule.

Attention:

Zabbix proxy must use a separate database. Pointing it to the Zabbix server database will break the configuration.

All data collected by the proxy is stored locally before transmitting it over to the server. This way no data is lost due to any temporary communication problems with the server. The *ProxyLocalBuffer* and *ProxyOfflineBuffer* parameters in the **proxy configuration file** control for how long the data are kept locally.

Attention:

It may happen that a proxy, which receives the latest configuration changes directly from Zabbix server database, has a more up-to-date configuration than Zabbix server whose configuration may not be updated as fast due to the value of **CacheUpdateFrequency**. As a result, proxy may start gathering data and send them to Zabbix server that ignores these data.

Zabbix proxy is a data collector. It does not calculate triggers, process events or send alerts. For an overview of what proxy functionality is, review the following table:

Function	Supported by proxy
Items	
<i>Zabbix agent checks</i>	Yes
<i>Zabbix agent checks (active)</i>	Yes ¹
<i>Simple checks</i>	Yes
<i>Trapper items</i>	Yes
<i>SNMP checks</i>	Yes
<i>SNMP traps</i>	Yes
<i>IPMI checks</i>	Yes
<i>JMX checks</i>	Yes
<i>Log file monitoring</i>	Yes
<i>Internal checks</i>	<i>No</i>
<i>SSH checks</i>	Yes
<i>Telnet checks</i>	Yes
<i>External checks</i>	Yes
Built-in web monitoring	Yes
Network discovery	Yes
Low-level discovery	Yes
Calculating triggers	<i>No</i>
Processing events	<i>No</i>
Sending alerts	<i>No</i>
Remote commands	<i>No</i>

Note:

[1] To make sure that an agent asks the proxy (and not the server) for active checks, the proxy must be listed in the **ServerActive** parameter in the agent configuration file.

Configuration

Once you have **installed** and **configured** a proxy, it is time to configure it in the Zabbix frontend.

Adding proxies

To configure a proxy in Zabbix front end:

- Go to: *Administration* → *DM*
- Click on *Create proxy*

Proxy

Proxy name: W1 proxy

Proxy mode: Active

Hosts: Proxy hosts

New host

srv_01

srv_02

srv_03

srv_04

srv_05

Other hosts

Baseline Switch 2250-SFP Plus

Host with Notepad

ProCurve J4900B Switch 2626


Zabbix server

Save Clone Delete Cancel

Parameter	Description
<i>Proxy name</i>	Enter the proxy name. It must be the same name as in the <i>Hostname</i> parameter in the proxy configuration file.
<i>Proxy mode</i>	Select the proxy mode. Active - the proxy will connect to the Zabbix server and request configuration data Passive - Zabbix server connects to the proxy
<i>Hosts</i>	Add hosts to be monitored by the proxy.

Host configuration

You can specify that an individual host should be monitored by a proxy in the **host configuration** form, using the *Monitored by proxy* field.



Host **mass update** is another way of specifying that hosts should be monitored by a proxy.

2 Nodes

Overview

You can use nodes to build a hierarchy of distributed monitoring.

Each node is a full Zabbix server and is responsible for monitoring its own location. Zabbix supports up to a thousand nodes in a distributed setup.

The benefits of using a node setup:

- building a multi-level hierarchy of monitoring in a large network involving several geographical locations. A node in the hierarchy reports to its master node only.
- a node can be configured locally or through its master node, which has a copy of configuration data of all child nodes.
- data gathering becomes more immune to possible communication problems. If communication between a master and a child node breaks down, nodes can keep operating. Historical information and events are stored locally. When communication is back, a child node will optionally send the data to the master node.
- the nodes can split the work of a single Zabbix server having to monitor thousands of hosts
- attaching and detaching new nodes does not affect the functionality of the existing setup. No restart of any node required.

Platform independence

A node may use its own platform (OS, hardware) and database engine independently of other nodes. Also child nodes can be installed without Zabbix frontend.

The nodes of higher levels should use a combination of better hardware with MySQL InnoDB, Oracle or PostgreSQL backend.

Attention:

A distributed monitoring setup will not work with an SQLite backend database.

Configuration

Node configuration

A Zabbix server installed by following the **standard installation procedure** is not configured as a node for a distributed setup.

To configure it as a node:

- Specify a unique **NodeID** in the server configuration file (zabbix_server.conf). Available values: 1-999 ('0' being the default value of a standalone server)
- Stop zabbix_server, make sure that it is NOT running

- Convert database data for a distributed setup, by running:

```
zabbix_server -n <node id>
```

Warning:

Run this command only **once**. Running it twice will corrupt the database, so make sure that you run it with the correct node id.

Warning:

It is strongly recommended to stop Apache web server before the conversion step.

For example, you may run (if NodeID is '1'):

```
cd bin
./zabbix_server -n 1 -c /usr/local/etc/zabbix_server.conf
```

Note:


Running zabbix_server with the **-n** option does not start the server process.

In a very simple setup, we may envisage this node (with NodeID=1) as the master, and go on to configure another Zabbix server as a child node, using the same procedure, only using a different node identifier, say, '2'. With two nodes configured, it is time to add them in the Zabbix front-end, in a very simple master-child relationship.

Front-end configuration (master node)

To configure the master node, open its Zabbix frontend:

- Go to: *Administration* → *DM*
- Make sure that *Nodes* are selected in the dropdown to the right
- Click on *Local node* to review its parameters



The screenshot shows the 'Node "Local node"' configuration window in the Zabbix frontend. It contains the following fields:

- Name:** Local node
- ID:** 1
- Type:** Local
- IP:** 192.168.1.1
- Port:** 10051

At the bottom right, there are 'Save' and 'Cancel' buttons. A help icon (?) is visible in the top right corner of the window.

Node attributes:

Parameter	Description
<i>Name</i>	Unique node name.
<i>Id</i>	Unique node ID. This is the value of NodeID from the configuration file.
<i>Type</i>	Local - the local node
<i>IP</i>	IP address of the local node. Zabbix trapper must be listening on this IP address.
<i>Port</i>	Port number of the local node. Zabbix trapper must be listening on this port number. Default is 10051.

Then add the child node:

- Click on *New node* in *Administration* → *DM*

Node

?

Name

Child node

ID

2

Type

Child

Master node

Local node

IP

192.168.1.5

Port

10051

Save

Cancel

Node attributes:

<i>Name</i>	Unique name of the child node.
<i>Id</i>	Unique node ID. This is the NodeID from the child node configuration file.
<i>Type</i>	Select the first of the two available values: Child - a child node Master - a master node
<i>Master node</i>	Select the master node for this child node.
<i>IP</i>	IP address of the child node. Zabbix trapper must be listening on this IP address.
<i>Port</i>	Port number of the child node. Zabbix trapper must be listening on this port number. Default is 10051.

Front-end configuration (child node)

To configure the child node, open its Zabbix frontend:

- Go to: *Administration* → *DM*
- Make sure that *Nodes* are selected in the dropdown to the right
- Click on *Local node* to review its parameters (see above for how to configure the local node)

Then add the master node:

- Click on *New node* in *Administration* → *DM*

Node

?

Name

Master node

ID

1

Type

Master

IP

192.168.1.1

Port

10051

Save

Cancel

Node attributes:

<i>Name</i>	Unique name of the master node.
<i>Id</i>	Unique master node ID. This is the NodeID from the master node configuration file.
<i>Type</i>	Select the second of the two available values: Child - a child node Master - a master node


<i>IP</i>	IP address of the master node. Zabbix trapper must be listening on this IP address on the master node.
<i>Port</i>	Port number of the master node. Zabbix trapper must be listening on this port number. Default is 10051.

Starting server daemons

To finish configuring a simple distributed setup, start the zabbix_server daemons, beginning with the master node daemon.

Display

As soon as nodes are defined, a dropdown for selecting one or several nodes appears in the Zabbix front-end.




Once selected, all information displayed in the front-end will come from the selected node(s).

More complex configurations

You can use the principles outlined to build more complex, multi-level monitoring hierarchies.

In this example, Riga (Node 4) will collect events and history from all the child nodes.



15. Web interface

Overview For an easy access to Zabbix from anywhere and from any platform, the web-based interface is provided.

Note:

Trying to access two Zabbix frontend installations on the same host, on different ports, simultaneously will fail. Logging into the second one will terminate the session on the first one and so on.

1 Frontend sections

1 Monitoring


Overview

The Monitoring menu is all about displaying data. Whatever information Zabbix is configured to gather, visualize and act upon, it will be displayed in the various sections of the Monitoring menu.

1 Dashboard

Overview


The *Monitoring* → *Dashboard* section, similar to the dashboard on your car, displays a summary of all the important information.



Favourites

There are some widgets for favourites where you can create quick shortcuts to the most needed graphs, custom graphs, screens, slide shows and maps.

Just click on the Menu button in the widget, select to add, for example, some screen and then select from the configured screens. The selected screens will be displayed as shortcuts in the favourites widget.




Status widgets

A number of status widgets - *Status of Zabbix*, *System status*, *Host status*, *Last 20 issues*, *Web monitoring*, *Discovery status* each display a summary of the respective data.

As you may have noticed from the screenshot, in Zabbix 2.0 the widgets can be arranged in up to three columns. Additionally, all widgets can be freely moved around. Just grab a widget by its title bar, drag and drop wherever you would like it.

Dashboard filter

Clicking on the first icon on the right of the *Personal dashboard* title bar allows you to access the dashboard filter.



The screenshot shows the 'Filter' dialog box in Zabbix. At the top, 'Dashboard filter' is set to 'Enabled'. Below this, there are two dropdown menus: 'Host groups' and 'Groups', both currently set to 'Selected'. A mouse cursor is hovering over the 'Groups' dropdown. Below these dropdowns is a large, empty list box. Underneath the list box are two buttons: 'Add' and 'Delete selected'. Further down, there are several checkboxes: 'Show hosts in maintenance' (checked), 'Not classified' (checked), 'Information' (checked), 'Warning' (checked), 'Average' (checked), 'High' (checked), and 'Disaster' (checked). At the bottom, there is a 'Problem display' dropdown menu set to 'All'.

If you enable the filter by clicking on the very first link there, you can then limit what hosts, triggers and problems are displayed in the dashboard.

2 Overview

Overview

The *Monitoring* → *Overview* section allows you to overview trigger states and also compare data for various hosts at once.

The following display options are available:

- you can select horizontal or vertical display of information in the *Hosts location* dropdown
- you can select all hosts or specific groups in the *Group* dropdown
- you can choose what information to display (triggers or data) in the *Type* dropdown

In the screenshot below information is displayed horizontally and trigger states are displayed, as colored blocks (the color depending on the state of the trigger):

OVERVIEW											
Hosts location		all		Type		Triggers					
Host											

Note that recent trigger changes (within the last 30 minutes) will be displayed as blinking blocks.

Trigger names, displayed vertically, is vertical text since Zabbix 2.0.4. Before that, it used to be vertical images.

In the next screenshot information is displayed vertically and data is displayed:

OVERVIEW		
Hosts location Top		
Group all Type Data		
Items	Zabbix server	New host
Buffers memory	130.35 MB	77.11 MB
Cached memory	162.64 MB	338.51 MB
Checksum of /etc/passwd	2494012136	3965433518
CPU idle time (avg1)	78.36	76.54
CPU nice time (avg1)	0.12	4.89
CPU system time (avg1)	11.41	8.46
CPU iowait time (avg1)	0.65	0.87
CPU user time (avg1)	6.57	10.22
Download speed for scenario "Availability of google".	-	208.3 KBps
Download speed for scenario "Zabbix frontend".	18.24 KBps	-
Download speed for step "Home" of scenario "Availability of google".	-	306.43 KBps
Download speed for step "About" of scenario "Availability of google".	-	110.17 KBps
Download speed for step "First page" of scenario "Zabbix frontend".	10.12 KBps	-
Download speed for step "Logging in" of scenario "Zabbix frontend".	15.47 KBps	-
Download speed for step "Login check" of scenario "Zabbix frontend".	23.22 KBps	-
Download speed for step "Logging out" of scenario "Zabbix frontend".	24.17 KBps	-
Failed step of scenario "Availability of google".	-	0
Failed step of scenario "Zabbix frontend".	0	-
Free disk space on /	1.25 GB	205.26 GB
Free disk space on / in %	32.42 %	95.05 %
Free memory	69.73 MB	57.16 MB
Host boot time	2012.05.07 09:35:21	2012.05.07 09:33:40
Host information	Linux linux-w7x1 2.6 ...	Linux martins-v 2.6. ...
Host name	linux-w7x1	martins-v
Host status	Up (1)	Down (0)
Host uptime (in sec)	1 day, 01:57:41	1.07 Mbps
Incoming traffic on interface eth0	2.01 KBps	1.24 KBps

Both trigger blocks and data items are displayed as links. If you click on a trigger block you can access trigger events and a simple graph/latest values list. If you click on a piece of data you can access some graphs for predefined periods or 500 latest values:

Zabbix server	135.51 MB	194.31 MB	2494012136	63.26	0.1	13.59
New host	54.23 MB	206.36				47
Graphs Last hour graph Last week graph Last month graph Values 500 latest values						

3 Web


Overview

In the *Monitoring* → *Web* section current information about **web scenarios** is displayed.

WEB CHECKS					
			Group	all	Host
Host	Name	Number of steps	State	Last check	Status
Zabbix server					
Zabbix frontend (1 scenario)					
	Zabbix frontend	4	Idle till 07 May 2012 08:40:11	07 May 2012 08:39:11	OK
New host					
Web scenarios (1 scenario)					
	Availability of google	2	Idle till 07 May 2012 08:40:05	07 May 2012 08:39:05	OK

The *Name* column contains first of all the applications (in bold type) under which the scenarios are grouped. Expanding applications by clicking on the '+' reveals the existing scenarios and their status overview.

The scenario name is link to more detailed statistics about it:



4 Latest data

Overview

The *Monitoring* → *Latest data* section displays the latest values gathered by items.

Just click on '+' before a host and the relevant application, and the items of that host and application will be displayed with their latest values.

LATEST DATA					
ITEMS					
		Filter		Group	all
				Host	all
Host	Name	Last check	Last value	Change	History
+	Zabbix server	Availability (10 Items)			
+	Zabbix server	CPU (8 Items)			
	Processor load15	11 May 2012 11:06:27	0.23	-	Graph
	Processor load	11 May 2012 11:06:39	0.17	-0.02	Graph
	Processor load5	11 May 2012 11:06:39	0.22	-0.01	Graph
	CPU idle time (avg1)	11 May 2012 11:06:40	86.48	+1.68	Graph
	CPU nice time (avg1)	11 May 2012 11:06:41	0.16	+0.03	Graph
	CPU system time (avg1)	11 May 2012 11:06:42	8.54	-1.2	Graph
	CPU user time (avg1)	11 May 2012 11:06:33	3.15	-0.46	Graph
	CPU iowait time (avg1)	11 May 2012 11:06:34	0.43	+0.02	Graph
+	Zabbix server	Filesystem (6 Items)			
+	Zabbix server	General (3 Items)			
+	Zabbix server	Integrity (1 Item)			
+	Zabbix server	Memory (5 Items)			
+	Zabbix server	Network (3 Items)			
+	Zabbix server	OS (4 Items)			
+	Zabbix server	Performance (2 Items)			
+	Zabbix server	Processes (4 Items)			
+	Zabbix server	Zabbix frontend (14 Items)			
+	New host	Availability (5 Items)			
+	New host	CPU (8 Items)			
+	New host	Filesystem (5 Items)			

You can expand all hosts and all applications, thus revealing all items by clicking on '+' in the header row.

Items are displayed with their name, last check time, last value, change amount and a link to a simple graph/history of item values.

Using filter

You can use the filter to display only the items you are interested in. The filter link is located above the table in the middle. You can use it to filter items by a string in the name; you can also select to display items that have no data gathered.

ITEMS					
		Filter		Group	all
				Host	all
		Show items with name like interface			
		Show items without data ☒			
		Filter Reset			
Host	Name	Last check	Last value	Change	History
+	Zabbix server	Network (3 Items)			
	Incoming traffic on interface eth0	11 May 2012 11:13:30	183.46 Bps	-422.83 Bps	Graph
	Outgoing traffic on interface eth0	11 May 2012 11:13:28	34.26 Bps	-523.66 Bps	Graph
	Network interfaces	11 May 2012 11:08:16	0	-	History
+	New host	Network (3 Items)			

Links to value history/simple graph

The last column in the latest value list offers:

- a **History** link (for all textual items) - leading to listings (*Values/500 latest values*) displaying the history of previous item values.
- a **Graph** link (for all numeric items) - leading to a **simple graph**. However, once the graph is displayed, a dropdown on the upper right offers a possibility to switch to *Values/500 latest values* as well.

Zabbix server: Processor load (1 min average per core)		Values	As plain text
		Filter	
Zoom: 1h 2h 3h 6h 12h 1d 7d 14d 1m 3m 6m All		19 Mar 2013 17:48 - 22 Mar 2013 17:48 (now)	
« 6m 1m 7d 1d 12h 1h 1h 12h 1d 7d 1m 6m »»		3d (fixed)	
Timestamp	Value		
2013.Mar.22 17:48:16	0.27		
2013.Mar.22 17:47:16	0.57		
2013.Mar.22 17:46:16	0.48		
2013.Mar.22 17:45:16	0.46		
2013.Mar.22 17:44:16	0.44		
2013.Mar.22 17:43:16	0.51		
2013.Mar.22 17:42:16	0.23		
2013.Mar.22 17:41:16	0.33		
2013.Mar.22 17:40:17	0.14		
2013.Mar.22 17:39:16	0.16		
2013.Mar.22 17:38:16	0.16		
2013.Mar.22 17:37:16	0.24		

The values displayed in this list are "raw", that is, no postprocessing is applied.

Note:

The total amount of values displayed is defined by the value of *Search/Filter elements limit* parameter, set in **Administration** → **General**.

5 Triggers

Overview

The *Monitoring* → *Triggers* section displays the status of triggers.

Triggers									
Displaying 1 to 3 of 3 found									
Filter									
Triggers status		Problem							
Acknowledge status		Any							
Events		Hide all							
Min severity		Warning							
Age less than		☒ 14 days							
Show details		☐							
Filter by name									
		Filter Reset							
☐	Severity	Status	Info	Last change	Age	Acknowledged	Host	Name	Comments
☐	Warning	OK		11 Jun 2012 13:53:52	1m 48s	Acknowledge (3)	New host	Processor load is too high on New host	Add
☐	Average	PROBLEM		11 Jun 2012 12:18:30	1h 37m 10s	Acknowledge (3)	Zabbix server	Zabbix agent on Zabbix server is unreachable for 5 minutes	Add
☐	Warning	PROBLEM		29 May 2012 12:10:30	13d 1h 45m	Acknowledged	Zabbix server	Lack of free swap space on Zabbix server	Show
Agent is unavailable.									
Bulk acknowledge Go (0)									

Column

Description

Severity

The trigger severity is displayed.
The color of the severity is used as cell background for problem triggers. For OK triggers, green background is used.

Column	Description
<i>Status</i>	The trigger status is displayed - OK or PROBLEM. By default, it will be blinking for 30 minutes for triggers that have recently changed their state. Additionally, triggers that have recently changed their state to OK, will be displayed for 30 minutes even if the filter is set to show only problems. Text color and blinking options can be configured in <i>Administration</i> → <i>General</i> → <i>Trigger displaying options</i>.
<i>Info</i>	A grey icon with a question mark indicates that there is some relevant information to be displayed. If you move your mouse pointer over it, the message will be displayed.
<i>Last change</i>	The date and time of last trigger status change is displayed.
<i>Age</i>	The age of last trigger status change is displayed.
<i>Acknowledged</i>	The acknowledgment status of the trigger is displayed: Acknowledged - green text indicating that the trigger is acknowledged. A trigger is considered to be acknowledged if all problem events for it are acknowledged (or if there have been only OK events). Acknowledge - a red link indicating unacknowledged problem events (and their number in parenthesis). If you click on the link you will be taken to a bulk acknowledgement screen where all events for this trigger can be acknowledged at once. <i>Note:</i> If you wish to acknowledge a single event only, go to <i>Monitoring</i> → <i>Events</i>. No events - if there have been no problem events for the trigger. Displaying this string is supported since Zabbix 2.0.4; prior to that these triggers were displayed as <i>Acknowledged</i>.
<i>Host</i>	The host of the trigger is displayed. It is also a link to the defined custom scripts, latest host data, host screens, and, if host inventory is enabled, to inventory information.
<i>Name</i>	The name of the trigger is displayed. It is also a link to the trigger event list and the trigger configuration page, as well as to a simple graph of item data. The link list may also contain a custom trigger URL, if one is defined in trigger configuration.
<i>Comments</i>	A link to comments about the trigger.

Using filter

You can use the filter to display only the triggers you are interested in. The filter link is the blue bar located above the table.

By default the filter is set to display triggers in problem status, also including the triggers that have only very recently changed from 'Problem' status to 'OK'. You can switch to display triggers in 'Any' status, however, the next time you return to this page, problem triggers will again be selected by default.

6 Events

Overview

The *Monitoring* → *Events* section displays latest **events**.

Events

Group

all

 Host

all

 Source

Trigger

Displaying 1 to 32 of 88 found

Filter

1 | 2 | 3 | Next >

Time	Host	Description	Status	Severity	Duration	Ack	Actions
11 Jun 2012 15:43:30	Zabbix server	Zabbix agent on Zabbix server is unreachable for 5 minutes	PROBLEM	Average	12s	No	In progress
11 Jun 2012 15:40:52	New host	Processor load is too high on New host	PROBLEM	Warning	2m 50s	No	Ok
11 Jun 2012 15:18:33	Zabbix server	Zabbix server has just been restarted	OK	Information	25m 9s	No	-
11 Jun 2012 15:16:32	ProCurve J4900B Switch 2626	#OperStatus.16 on ProCurve J4900B Switch 2626 has changed	OK	Not classified	27m 10s	No	-
11 Jun 2012 15:15:32	ProCurve J4900B Switch 2626	#OperStatus.16 on ProCurve J4900B Switch 2626 has changed	PROBLEM	Not classified	1m	No	-
11 Jun 2012 15:08:33	Zabbix server	Zabbix server has just been restarted	PROBLEM	Information	10m	No	-
11 Jun 2012 15:02:07	Zabbix server	Zabbix agent on Zabbix server is unreachable for 5 minutes	OK	Average	41m 23s	No	-
11 Jun 2012 14:40:56	ProCurve J4900B Switch 2626	#OperStatus.10 on ProCurve J4900B Switch 2626 has changed	OK	Not classified	1h 2m 46s	No	-
11 Jun 2012 14:39:26	ProCurve J4900B Switch 2626	#OperStatus.10 on ProCurve J4900B Switch 2626 has changed	PROBLEM	Not classified	1m 30s	No	Failed

In the last dropdown to the right you can select trigger or discovery based events.

Currently displayed events can be exported to a CSV file. Look for the *Export to CSV* button to the right on the title bar.

Export to CSV

Host all Source Trigger

Clicking on the timestamp in the first column of trigger events will take you to event details.

Event source details

Host

Zabbix server

Trigger

Zabbix agent on Zabbix server is unreachable for 5 minutes

Severity

Average

Expression

[Zabbix server.agent.ping.nodata(5m)]=1

Event generation

Normal

Disabled

No

Event details

Event

Zabbix agent on Zabbix server is unreachable for 5 minutes

Time

11 Jun 2012 15:43:30

Acknowledged

Yes (1)

Acknowledges

Time

User

Comments

11 Jun 2012 15:47:56

Admin

Agent restarted.

Message actions

Time

Type

Status

Retries left

Recipient(s)

Message

Error

Step: 1

Email

sent

Martins.Valkovskis@zabbix.com

Subject:

PROBLEM: Zabbix agent on Zabbix server is unreachable for 5 minutes

Time:

11 Jun 2012 15:43:31

Message:

Trigger: Zabbix agent on Zabbix server is unreachable for 5 minutes
 Trigger status: PROBLEM
 Trigger severity: Average
 Trigger URL:

Item values:

1. Agent ping (Zabbix server.agent.ping): Up (1)
 2. *UNKNOWN*
 (*UNKNOWN*:*UNKNOWN*): *UNKNOWN*
 3. *UNKNOWN*
 (*UNKNOWN*:*UNKNOWN*): *UNKNOWN*

Command actions

Time

Status

Command

Error

No actions found.

Event list [previous 20]

Time

Status

Duration

Age

Ack

Actions

11 Jun 2012 15:43:30

PROBLEM

5m 14s

5m 15s

Yes (1)

1

11 Jun 2012 15:02:07

OK

41m 23s

46m 38s

No

-

11 Jun 2012 12:18:30

PROBLEM

2h 43m 37s

3h 30m 15s

No


1

Details about the event, its source, acknowledgments, actions taken (messages, remote commands) and previous similar events are displayed.

7 Graphs

Overview

In the *Monitoring* → *Graphs* section any **custom graph** that has been configured can be displayed.



To display a graph, select the host group, host and then the graph from the dropdowns to the right.

Time period selector

The filter section above the graph contains a time period selector. It allows you to select the desired time period easily.

The slider within the selector can be dragged back and forth, as well as resized, effectively changing the time period displayed. Links on the left hand side allow to choose some often-used predefined periods (above the slider area) and move them back and forth in time (below the slider area). The dates on the right hand side actually work as links, popping up a calendar and allowing to set a specific start/end time.

The **fixed/dynamic** link in the lower right hand corner has the following effects:

- controls whether the time period is kept constant when you change the start/end time in the calendar popup.
- when *fixed*, time moving controls (« 6m 1m 7d 1d 12h 1h | 1h 12h 1d 7d 1m 6m ») will move the slider, while not changing its size, whereas when *dynamic*, the control used will enlarge the slider in the respective direction.
- when *fixed*, pressing the larger < and > buttons will move the slider, while not changing its size, whereas when *dynamic*, < and > will enlarge the slider in the respective direction. The slider will move by the amount of its size, so, for example, if it is one month, it will move by a month; whereas the slider will enlarge by 1 day.

Another way of controlling the displayed time is to highlight an area in the graph with the left mouse button. The graph will zoom into the highlighted area once you release the left mouse button.

Controls

Three control buttons are available in the title bar:

-  - add graph to the favourites widget in the **Dashboard**
-  - reset graph display to the original setting of displaying the last hour data
-  - use the full browser window to display the graph

8 Screens

Overview

In the *Monitoring* → *Screens* section any configured **screen** or **slide show** can be displayed.



Use the dropdown in the title bar to switch between screens and slide shows.

Time period selector

The filter section above the screen/slide show contains a time period selector. It allows you to select the desired time period easily, affecting the data displayed in graphs etc.

Controls

Three control buttons are available in the title bar:

- add screen/slide show to the favourites widget in the **Dashboard**
- use the full browser window to display the screen/slide show
- slow down or speed up a slide show

Referencing a screen

Screens can be referenced by both `elementid` and `screenname` GET parameters. For example,

`http://zabbix/zabbix/screens.php?screenname=Zabbix%20server`


will open the screen with that name (Zabbix server).

If both `elementid` (screen ID) and `screenname` (screen name) are specified, `screenname` has higher priority.

9 Maps

Overview

In the *Monitoring* → *Maps* section any configured **network map** can be viewed.



You can use the dropdown in the map title bar to switch between different configured maps.

Icon highlighting

If a map element is in problem status, it is highlighted with a round circle. The fill colour of the circle corresponds to the severity colour of the problem trigger. If all problems are acknowledged, a thick green border around the circle is displayed.

Additionally, a host in **maintenance** is highlighted with an orange, filled square and a disabled (not-monitored) host is highlighted with a grey, filled square. Highlighting is displayed if the *Icon highlighting* check-box is marked in map **configuration**.

Recent change markers

Inward pointing red triangles around an element indicate a recent trigger status change - one that's happened within the last 30 minutes. These triangles are shown if the *Mark elements on trigger status change* check-box is marked in map **configuration**.

Links

Clicking on a map element opens a menu with some available links.

Controls

Two control buttons are available in the title bar:

-  - add map to the favourites widget in the **Dashboard**
-  - use the full browser window to display the map

Referencing a network map

Network maps can be referenced by both `sysmapid` and `mapname` GET parameters. For example,

`http://zabbix/zabbix/maps.php?mapname=Local%20network`

will open the map with that name (Local network).

If both `sysmapid` (map ID) and `mapname` (map name) are specified, `mapname` has higher priority.

10 Discovery

Overview

In the *Monitoring* → *Discovery* section results of **network discovery** are shown. Discovered devices are sorted by the discovery rule.

STATUS OF DISCOVERY			
Discovery rules			Discovery rule: all
Discovered device	Monitored host	Uptime/Downtime	SNMPv2 agent: .1.3.6.1.2.1.1.1.0
Local network_Z (3 devices)			
192.168.3.14	Baseline Switch 2250-SFP	21 days, 16:01:56	
192.168.3.9	192.168.3.9	5 days, 14:42:42	

If a device is already monitored, the host name will be listed in the *Monitored host* column, and the duration of the device being discovered or lost after previous discovery is shown in the *Uptime/Downtime* column.

After that follow the columns showing the state of individual services for each discovered device. A tooltip for each cell will show individual service uptime or downtime.

Attention:

Only those services that have been found on at least one device will have a column showing their state.

11 IT services

Overview

In the *Monitoring* → *IT services* section the status of **IT services** is displayed.

A list of the existing IT services is displayed along with data of their status and SLA. From the dropdown in the upper right corner you can select a desired period for display.

Displayed data:

Parameter	Description
<i>Service</i>	Service name.
<i>Status</i>	Status of service: OK - no problems (trigger colour and severity) - indicates a problem and its severity
<i>Reason</i>	Indicates the reason of problem (if any).
<i>Problem time</i>	Displays SLA bar. Green/red ratio indicates the proportion of availability/problems. The bar displays the last 20% of SLA (from 80% to 100%). The bar contains a link to a graph of availability data.
<i>SLA/Acceptable SLA</i>	Displays current SLA/expected SLA value. If current value is below the acceptable level, it is displayed in red.

You can also click on the service name to access the *IT Services Availability Report*.

IT SERVICES AVAILABILITY REPORT "Zabbix server"					
			Period: Monthly Year: 2011		
Month	Ok	Problems	Downtime	Percentage	SLA
Jan 2011	31d 0h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
Feb 2011	28d 0h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
Mar 2011	30d 23h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
Apr 2011	30d 0h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
May 2011	31d 0h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
Jun 2011	30d 0h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
Jul 2011	31d 0h 0m	0d 0h 0m	0d 0h 0m	100.00%	99.0500
Aug 2011	29d 16h 49m	0d 0h 29m	0d 0h 0m	99.93%	99.0500

Here you can assess IT service availability data over a longer period of time on daily/weekly/monthly/yearly basis.

2 Inventory

Overview

The Inventory menu features sections providing an overview of host inventory data by a chosen parameter as well as the ability to view host inventory details.

1 Overview

Overview

The *Inventory* → *Overview* section provides ways of having an overview of **host inventory** data.

For an overview to be displayed, choose a host group (or all groups) and the inventory field by which to display data. The number of hosts corresponding to each entry of the chosen field will be displayed.

HOST INVENTORY OVERVIEW

Hosts

Group

Discovered hosts

Grouping by

Type

Type	Host count ↓
Switch	2
Unknown	1
Workstation	1

Zabbix 2.0.0 Copyright 2001-2012 by Zabbix SIA

Type

Site address C

Site city

Site country

Site notes

Site rack location

Site state / province

Software

Software (Full details)

Software application A

Software application B

Software application C

Software application D

Software application E

Tag

Type

Type (Full details)

URL A

URL B

URL C

Vendor

The completeness of an overview depends on how much inventory information is maintained with the hosts.

Numbers in the *Host count* column are links; they lead to these hosts being filtered out in the *Host Inventories* table.

HOST INVENTORIES

Hosts

Group

Discovered hosts

Displaying 1 to 2 of 2 found

Filter

Field:

Type

 exactly

Switch

Filter

Reset

Host ↑	Group	Name	Type	OS	Serial number A	Tag	MAC address A
Baseline Switch 2250-SFP	Discovered hosts, Switches	Baseline Switch 2250-SFP Plus	Switch				
ProCurve J4900B Switch 2626	Discovered hosts, Switches	ProCurve Switch 2626	Switch				

2 Hosts

Overview

In the *Inventory* → *Hosts* section **inventory data** of hosts are displayed.

Select a group from the dropdown in the upper right corner to display the inventory data of hosts in that group. You can also filter the hosts by any inventory field to display only the hosts you are interested in.

HOST INVENTORIES							
Hosts							Group Discovered hosts
Displaying 1 to 1 of 1 found							
Filter Field: Notes like occasion Filter Reset							
Host	Group	Name	Type	OS	Serial number A	Tag	MAC address A
New host	Discovered hosts	mnd	Workstation	Linux mnd 2.6.38-15-generic #61-Ubuntu SMP Tue Jun 12 19:1	CZC9XX3XCP		[eth0] 00:24:21:05:66:fd

To display all host inventories, select "all" in the group dropdown, clear the comparison field in the filter and press "Filter".

While only some key inventory fields are displayed in the table, you can also view all available inventory information for that host. To do that, click on the host name in the first column.

Host Inventory	
Type	Workstation
Name	mnd
OS	Linux mnd 2.6.38-15-generic #61-Ubuntu SMP Tue Jun 12 19:1
OS (Short)	Ubuntu 11.04
Serial number A	CZC9XX3XCP
MAC address A	[eth0] 00:24:21:05:66:fd
Hardware (Full details)	processor 0: AuthenticAMD AMD Athlon(tm) Dual Core Processor 5400B working at 1000MHz (maximum 2800MHz) processor 1: AuthenticAMD AMD Athlon(tm) Dual Core Processor 5400B working at 1000MHz (maximum 2800MHz)
Location	Head Office
Notes	Occasionally annoying
Vendor	http://welcome.hp.com/country/us/en/welcome.html
Site city	Riga
Cancel	

The completeness of inventory data depends on how much inventory information is maintained with the host.

3 Reports

Overview

The Reports menu features several sections that contain a variety of predefined and user-customizable reports focused on displaying an overview of such parameters as the status of Zabbix, triggers and gathered data.

1 Status of Zabbix

Overview

In *Reports* → *Status of Zabbix* a summary of key system data is displayed.

STATUS OF ZABBIX		
Report		
Parameter	Value	Details
Zabbix server is running	Yes	localhost:10051
Number of hosts (monitored/not monitored/templates)	43	17 / 0 / 26
Number of items (monitored/disabled/not supported)	867	860 / 1 / 6
Number of triggers (enabled/disabled)[problem/unknown/ok]	146	144 / 2 [1 / 0 / 143]
Number of users (online)	3	2
Required server performance, new values per second	10.97	-

This report is also displayed as a widget in the **Dashboard**.

Displayed data

Parameter	Value	Details
<i>Zabbix server is running</i>	Status of Zabbix server: Yes - server is running No - server is not running <i>Note:</i> To make sure the web frontend knows that the server is running there must be at least one trapper process started on the server (StartTrappers parameter in <code>zabbix_server.conf</code> file>0).	Location and port of Zabbix server.
<i>Number of hosts</i>	Total number of hosts configured is displayed. Templates are counted as a type of host too.	Number of monitored hosts/not monitored hosts/templates.
<i>Number of items</i>	Total number of items is displayed. Only items assigned to enabled hosts are counted.	Number of monitored/disabled/unsupported items.
<i>Number of triggers</i>	Total number of triggers is displayed. Only triggers assigned to enabled hosts and depending on enabled items are counted.	Number of enabled/disabled triggers. [Triggers in problem/unknown/ok state.]
<i>Number of users</i>	Total number of users configured is displayed.	Number of users online.
<i>Required server performance, new values per second</i>	The expected number of new values processed by Zabbix server per second is displayed.	<i>Required server performance</i> is an estimate and can be useful as a guideline. For precise numbers of values processed, use the <code>zabbix[wcache,values,all]</code> internal item. Enabled items from monitored hosts are included in the calculation. Log items are counted as one value per item update interval. Regular interval values are counted; flexible interval values are not. The calculation is not adjusted during a "nodata" maintenance period. Trapper items are not counted.

2 Availability report

Overview

In *Reports* → *Availability report* you can see what proportion of time each trigger has been in problem/ok/unknown state. The percentage of time for each state is displayed.

Thus it is easy to determine the availability situation of various elements on your system.

AVAILABILITY REPORT

Report

Mode By host

Filter

Template groupZabbix servers

TemplateZabbix server

From01 / 08 / 2012 00 : 00

Till01 / 09 / 2012 00 : 00

FilterReset

Name	Problems	Ok	Unknown	Graph
/etc/passwd has been changed on Zabbix server	0.0000%	81.9361%	18.0639%	Show
Configured max number of opened files is too low on Zabbix server	0.0000%	81.9381%	18.0619%	Show
Configured max number of processes is too low on Zabbix server	0.0000%	81.9380%	18.0620%	Show
Disk I/O is overloaded on Zabbix server	0.0269%	82.0425%	17.9307%	Show
Free disk space is less than 20% on volume /	0.0000%	82.0737%	17.9263%	Show
Free inodes is less than 20% on volume /	0.0000%	82.0739%	17.9261%	Show
Host information was changed on Zabbix server	0.0000%	81.9363%	18.0637%	Show
Host name of zabbix_agentd was changed on Zabbix server	0.0000%	81.9352%	18.0648%	Show
Hostname was changed on Zabbix server	0.0000%	81.9367%	18.0633%	Show
Lack of available memory on server Zabbix server	0.0000%	82.0705%	17.9295%	Show
Less than 25% free in the configuration cache	0.0000%	82.1217%	17.8783%	Show
Less than 25% free in the history cache	0.0000%	82.1216%	17.8784%	Show
Less than 25% free in the text history cache	0.0000%	82.1238%	17.8762%	Show
Less than 25% free in the trends cache	0.0000%	82.1238%	17.8762%	Show
More than 100 items having missing data for more than 10 minutes	0.0224%	82.1061%	17.8716%	Show
Processor load is too high on Zabbix server	0.0291%	82.0406%	17.9303%	Show
Too many processes on Zabbix server	0.0000%	82.0722%	17.9278%	Show
Too many processes running on Zabbix server	0.0000%	82.0723%	17.9277%	Show
Version of zabbix_agentd was changed on Zabbix server	0.0000%	81.9381%	18.0619%	Show
Web scenario failed	99.5180%	0.0082%	0.4738%	Show
Zabbix agent on Zabbix server unreachable for 5min	0.0288%	82.0817%	17.8895%	Show

From the dropdown in the upper right corner you can choose the selection mode - whether to display triggers by hosts or by triggers belonging to a template. Then in the filter you can narrow down the selection to the desired options and the time period.

AVAILABILITY REPORT

Report

ModeBy trigger template

Filter

Template groupTemplates

TemplateTemplate OS Linux

Template triggerZabbix agent on Template OS Linux un

Filter by host groupall

Period

From01 / 08 / 201200 : 00


Till01 / 09 / 201200 : 00

FilterReset

Host	Name	Problems	Ok	Unknown	Graph
New host	Zabbix agent on New host unreachable for 5min	1.4490%	80.6537%	17.8972%	Show
Zabbix server	Zabbix agent on Zabbix server unreachable for 5min	0.0288%	82.0817%	17.8895%	Show

The name of the trigger is a link to the latest events of that trigger.

Clicking on *Show* in the Graph column displays a bar graph where availability information is displayed in bar format each bar representing a past week of the current year.



The green part of a bar stands for OK time, red for problem time, and light grey for the unknown time when no information was gathered.

3 Triggers top 100

Overview

In *Reports* → *Triggers top 100* you can see the triggers that have changed their state most often within the period of evaluation, sorted by the number of status changes.

From the dropdown in the upper right corner you can choose the time period for evaluation - day, week, month, year.

MOST BUSY TRIGGERS TOP 100			
Report			
Month			
Host	Trigger	Severity	Number of status changes
New host	Disk I/O is overloaded on New host	Warning	32
ProCurve J4900B Switch 2626	Operational status was changed on ProCurve J4900B Switch 2626 interface 16	Information	28
New host	Zabbix agent on New host unreachable for 5min	Average	23
Baseline Switch 2250-SFP	Operational status was changed on Baseline Switch 2250-SFP interface Ethernet0/3	Information	20
ProCurve J4900B Switch 2626	Operational status was changed on ProCurve J4900B Switch 2626 interface 18	Information	19
ProCurve J4900B Switch 2626	Operational status was changed on ProCurve J4900B Switch 2626 interface 15	Information	10
Zabbix server	Disk I/O is overloaded on Zabbix server	Warning	10
Baseline Switch 2250-SFP	Operational status was changed on Baseline Switch 2250-SFP interface Ethernet0/5	Information	8
New host	Processor load is too high on New host	Warning	8
Zabbix server	Processor load is too high on Zabbix server	Warning	8
ProCurve J4900B Switch 2626	Operational status was changed on ProCurve J4900B Switch 2626 interface 12	Information	6
Zabbix server	Zabbix server has just been restarted	Information	6
Zabbix server	Zabbix agent on Zabbix server unreachable for 5min	Average	5
Baseline Switch 2250-SFP	Operational status was changed on Baseline Switch 2250-SFP interface Ethernet0/3	Information	4
New host	New host has just been restarted	Information	4
ProCurve J4900B Switch 2626	Operational status was changed on ProCurve J4900B Switch 2626 interface 16	Information	4
Zabbix server	Zabbix discoverer processes more than 75% busy	Warning	4
Zabbix server	More than 100 items having missing data for more than 10 minutes	Warning	2
Zabbix server	Web scenario failed	Information	2
Zabbix server	Zabbix housekeeper processes more than 75% busy	Average	2

Both host and trigger column entries are links that offer some useful options:

- for host - links to user-defined scripts and latest data
- for trigger - links to latest events, the trigger configuration form and a simple graph

4 Bar reports

Overview

In the *Reports* → *Bar reports* section you can create some customized bar reports on run-time. Reports can be viewed, but are not saved.

From the dropdown in the upper right corner you can choose one of the three types of available bar reports. Then use the *Filter* options to create the report.

Parameter	Description
<i>Title</i>	Name of the report.
<i>X label</i>	Label displayed below the X axis.
<i>Y label</i>	Label displayed alongside the Y axis.
<i>Legend</i>	With this checkbox marked, a legend will be displayed alongside the report.
<i>Scale</i>	Picking a scale will separate out the value bars for either every hour/day/week/month/year. So, for example, picking a <i>daily</i> scale will display one bar for the values of one day.
<i>Period</i>	This parameter is available for the first and third report type. Enter the start and end of the evaluation period. With the second report type, several custom periods, each displayed in different colour, can be entered.
<i>Items</i>	Click on <i>Add</i> to select the items whose data you wish to display.

Specifically for the third report type:

<i>Groups</i>	From the <i>Other groups</i> pane select host groups. Item values for any host in the group having that item will be displayed.
<i>Hosts</i>	From the <i>Other hosts</i> pane select hosts. Item values for any selected host having that item will be displayed.
<i>Average</i>	Select whether to display averaged data for an hour/day/week/month/year.
<i>Item</i>	Select the item whose data you wish to display.
<i>Palette</i>	Pick a palette of colours for displaying side-by-side bars and colour intensity (middle/darken/brighter).

Item data comparison

The first bar report offers a possibility to simply compare item values side by side.

Bar reports

Report Reports Distribution of values for multiple periods

Filter

Title: Report 1: Incoming vs outgoing traffic

X label: Avg per day

Y label: Traffic amount

Legend: ☒


Scale: Daily

Period: From 31 / 08 / 2012 00 : 00 To 06 / 09 / 2012 00 : 00

Items:

- ☐ Incoming network traffic on eth0 New host: Incoming network traffic on eth0 avg Left
- ☐ Outgoing network traffic on eth0 New host: Outgoing network traffic on eth0 avg Left

Add Delete selected Show Reset



Period data comparison

The second bar report offers a possibility to compare the values of one or several items in custom periods.

Report Reports Distribution of values for multiple items

Filter

Title: Report 2: Working hours vs non-working hours

X label: Parameters

Y label: %

Legend: ☐

Period:


- ☐ 03 Sep 2012 09:00:00 - 03 Sep 2012 18:00:00 03 Sep 2012 09:00:00 03 Sep 2012 18:00:00
- ☐ 03 Sep 2012 18:00:00 - 04 Sep 2012 09:00:00 03 Sep 2012 18:00:00 04 Sep 2012 09:00:00
- ☐ 04 Sep 2012 09:00:00 - 04 Sep 2012 18:00:00 04 Sep 2012 09:00:00 04 Sep 2012 18:00:00
- ☐ 04 Sep 2012 18:00:00 - 05 Sep 2012 09:00:00 04 Sep 2012 18:00:00 05 Sep 2012 09:00:00

Add Delete selected

Items:


- ☐ CPU iowait time New host: CPU iowait time avg
- ☐ CPU idle time New host: CPU idle time avg
- ☐ CPU system time New host: CPU system time avg
- ☐ CPU user time New host: CPU user time avg
- ☐ CPU nice time New host: CPU nice time avg

Add Delete selected Show Reset



One item data comparison

The third bar report offers a possibility to compare the values of one item for different hosts/predefined intervals (hourly/daily/weekly/monthly/yearly).



4 Configuration

Overview

The Configuration menu contains sections for setting up major Zabbix functions, such as hosts and host groups, data gathering, data thresholds, sending problem notifications, creating data visualisation and others.

1 Host groups

Overview

In the *Configuration* → *Host groups* section users can configure and maintain host groups. A host group can contain both templates and hosts.

A listing of existing host groups with their details is displayed.

CONFIGURATION OF HOST GROUPS

Create host group

Host groups

Displaying 1 to 6 of 6 found

☐	Name	#	Members
☒	Discovered hosts	Templates (0) Hosts (1)	192.168.1.17
☐	Linux servers	Templates (0) Hosts (12)	host 1, host 2, host 3, host 4, host 5, host 6, host 7, host 8, host 9, host 10, host 11, host 12
☐	Switches	Templates (0) Hosts (2)	Baseline_Switch 2250-SFP, ProCurve J4900B Switch 2626
☐	Templates	Templates (26) Hosts (0)	Template OS Linux , Template App Zabbix Server , Template App Zabbix Agent , Template App Agentless , Template SNMP Interfaces , Template SNMP Generic , Template SNMP Device , Template SNMP OS Windows , Template SNMP Disks , Template SNMP OS Linux , Template SNMP Processors , Template IPMI Intel SR1530 , Template IPMI Intel SR1630 , Template App MySQL , Template OS OpenBSD , Template OS FreeBSD , Template OS AIX , Template OS HP-UX , Template OS Solaris , Template OS Mac OS X , Template OS Windows , Template JMX Generic , Template JMX Tomcat , Template HP Procurve2 , Template HP Procurve , C Template
☐	Workstations	Templates (1) Hosts (1)	Template OS Linux New host
☐	Zabbix servers	Templates (0) Hosts (1)	Zabbix server

Enable selected

Enable selected

Disable selected

Delete selected

Go (1)

Zabbix 2.0.0 Copyright 2001-2012 by Zabbix SIA

Connected as 'Admin'

Displayed data:

Column	Description
<i>Name</i>	Name of the host group. Clicking on the group name opens the host group configuration form.
<i>#</i>	Number of templates and hosts in the group (displayed in parentheses). Clicking on "Templates" or "Hosts" will, in the whole listing of templates or hosts, filter out those that belong to the group.
<i>Members</i>	Names of group members. Template names are displayed in grey, monitored host names in blue and non-monitored host names in red. Clicking on a name will open the template/host configuration form.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the status of all hosts in the group to "Monitored"
- *Disable selected* - change the status of all hosts in the group to "Not monitored"
- *Delete selected* - delete the host groups

To use these options, mark the check-boxes before the respective host groups, then select the required option and click on "Go".

2 Templates

Overview

In the *Configuration* → *Templates* section users can configure and maintain templates.

A listing of existing templates with their details is displayed.

CONFIGURATION OF TEMPLATES								Create template	Import
Templates								Group SNMP Templates	
Displaying 1 to 7 of 7 found									
☐ Templates	Applications	Items	Triggers	Graphs	Screens	Discovery	Linked templates	Linked to	
☐ Template SNMP Device	Applications (2)	Items (6)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (1)	Template SNMP Generic , Template SNMP Interfaces	Baseline Switch 2250-SFP , ProCurve J4900B Switch 2626	
☐ Template SNMP Disks	Applications (1)	Items (0)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (1)	-	Template SNMP OS Linux , Template SNMP OS Windows	
☐ Template SNMP Generic	Applications (1)	Items (5)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (0)	-	Template SNMP Device , Template SNMP OS Linux , Template SNMP OS Windows	
☐ Template SNMP Interfaces	Applications (1)	Items (1)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (1)	-	Template SNMP Device , Template SNMP OS Linux , Template SNMP OS Windows	
☐ Template SNMP OS Linux	Applications (4)	Items (6)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (3)	Template SNMP Disks , Template SNMP Generic , Template SNMP Interfaces , Template SNMP Processors	-	
☐ Template SNMP OS Windows	Applications (4)	Items (6)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (3)	Template SNMP Disks , Template SNMP Generic , Template SNMP Interfaces , Template SNMP Processors	-	
☐ Template SNMP Processors	Applications (1)	Items (0)	Triggers (0)	Graphs (0)	Screens (0)	Discovery (1)	-	Template SNMP OS Linux , Template SNMP OS Windows	
Export selected Go (0)									

From the dropdown to the right in the title bar you can choose whether to display all templates or only those belonging to a group.

Displayed data:

Column	Description
<i>Templates</i>	Name of the template. Clicking on the template name opens the template configuration form .
<i>Elements (Applications, Items, Triggers, Graphs, Screens, Discovery)</i>	Number of the respective elements in the template (displayed in parentheses). Clicking on the element name will, in the whole listing of that element, filter out those that belong to the template.
<i>Linked templates</i>	Templates that are linked to the template, in a nested setup where the template will inherit all elements of the linked templates.
<i>Linked to</i>	The hosts and templates that the template is linked to.

To configure a new template, click on the *Create template* button in the top right-hand corner. To import a template from an XML file, click on the *Import* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Export selected* - export the template to an XML file
- *Delete selected* - delete the template while leaving its linked elements (items, triggers etc.) with the hosts
- *Delete selected with linked elements* - delete the template and its linked elements from the hosts

To use these options, mark the check-boxes before the respective templates, then select the required option and click on "Go".

3 Hosts

Overview

In the *Configuration* → *Hosts* section users can configure and maintain hosts.

A listing of existing hosts with their details is displayed.

From the dropdown to the right in the *Hosts* bar you can choose whether to display all hosts or only those belonging to one particular group.

CONFIGURATION OF HOSTS

Create hostImport

Hosts

GroupDiscovered hosts

Displaying 1 to 3 of 3 found

Filter

☐	Name	Applications	Items	Triggers	Graphs	Discovery	Interface	Templates	Status	Availability
☐	Baseline Switch 2250-SFP	Applications (2)	Items (447)	Triggers (55)	Graphs (55)	Discovery (1)	192.168.3.14: 161	Template SNMP Device (Template SNMP Generic, Template SNMP Interfaces)	Monitored	   
☐	New host	Applications (10)	Items (45)	Triggers (17)	Graphs (7)	Discovery (2)	192.168.3.39: 32050	Template OS Linux (Template App Zabbix Agent)	Monitored	   
☐	ProCurve J4900B Switch 2626	Applications (2)	Items (246)	Triggers (30)	Graphs (30)	Discovery (1)	192.168.3.7: 161	Template SNMP Device (Template SNMP Generic, Template SNMP Interfaces)	Monitored	   

Export selected

Go (0)

Export selected

Mass update

Enable selected

Disable selected

Delete selected

Zabbix 2.0.0 Copyright 2001-2012 by Zabbix SIA

Connected as 'Admin'

Displayed data:

Column	Description
<i>Name</i>	Name of the host. Clicking on the host name opens the host configuration form .
<i>Elements (Applications, Items, Triggers, Graphs, Discovery)</i>	Clicking on the element name will display items, triggers etc. of the host. The number of the respective elements is displayed in parentheses.
<i>Interface</i>	The main interface of the host is displayed.
<i>Templates</i>	The templates linked to the host are displayed. If other templates are contained in the linked template, those are displayed in parentheses, separated by a comma. Clicking on a template name will open its configuration form.
<i>Status</i>	Host status is displayed - <i>Monitored</i> or <i>Not monitored</i> . By clicking on the status you can change it.
<i>Availability</i>	Availability of the host is displayed. Four icons each represent a supported interface (Zabbix agent, SNMP, IPMI, JMX). If the interface is configured and available, it is displayed in green. If it is configured and unavailable, it is displayed in red, and, upon mouseover, will display details of why the interface cannot be reached.

To configure a new host, click on the *Create host* button in the top right-hand corner. To import a host from an XML file, click on the *Import* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Export selected* - export the hosts to an XML file
- *Mass update* - **update several properties** for a number of hosts at once
- *Enable selected* - change host status to *Monitored*
- *Disable selected* - change host status to *Not monitored*
- *Delete selected* - delete the hosts

To use these options, mark the check-boxes before the respective hosts, then select the required option and click on "Go".

Filter

As the list may contain very many hosts, it may be needed to filter out the ones you really need.

The narrow blue bar just below the *Hosts* bar is actually a link to the filter. If you click on it, a filter becomes available where you can filter hosts by name, DNS, IP or port number.

Hosts

GroupDiscovered hosts

Displaying 1 to 2 of 2 found

Filter

Name like:

DNS like:

IP like:

Port:161

Filter

Reset

1 Applications

Overview

The application list for a template can be accessed from *Configuration* → *Templates* and then clicking on Applications for the respective template.

The application list for a host can be accessed from *Configuration* → *Hosts* and then clicking on Applications for the respective host.

A list of existing applications is displayed.

CONFIGURATION OF APPLICATIONS

Create application

Applications

Group allHost New host

Displaying 1 to 11 of 11 found

« Host listHost: New hostMonitored

Applications (11)Items (43)Triggers (18)Graphs (7)Discovery rules (2)

☐	Application 	Show
☐	Template OS Linux: CPU	Items (13)
☐	Template OS Linux: Filesystems	Items (5)
☐	Template OS Linux: General	Items (5)
☐	Template OS Linux: Memory	Items (5)
☐	Template OS Linux: Network interfaces	Items (4)
☐	Template OS Linux: OS	Items (8)
☐	Template OS Linux: Performance	Items (13)
☐	Template OS Linux: Processes	Items (2)
☐	Template OS Linux: Security	Items (2)
☐	SNMP data	Items (1)
☐	Template App Zabbix Agent: Zabbix agent	Items (4)

Enable selected

Go (0)

Displayed data:

Column	Description
Name	Name of the application, displayed as a blue link for directly created applications. Clicking on the application name link opens the application configuration form. If the host application belongs to a template, the template name is displayed before the application name, as a grey link. Clicking on the template link will open the application list on the template level.
Show	Click on Items to view the items contained in the application. The number of items is displayed in parentheses.

To configure a new application, click on the *Create application* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change application status to *Enabled*
- *Disable selected* - change application status to *Disabled*
- *Delete selected* - delete the applications

To use these options, mark the check-boxes before the respective applications, then select the required option and click on "Go".

2 Items

Overview

The item list for a template can be accessed from *Configuration* → *Templates* and then clicking on Items for the respective template.

The item list for a host can be accessed from *Configuration* → *Hosts* and then clicking on Items for the respective host.

A list of existing items is displayed.

CONFIGURATION OF ITEMS Create item										
Items										
Displaying 1 to 15 of 15 found										
Filter										
Host list Host: New host Monitored  Applications (11) Items (42) Triggers (18) Graphs (7) Discovery rules (2)										
Wizard	Name	Triggers	Key	Interval	History	Trends	Type	Applications	Status	Error
☐	Template App Zabbix Agent: Agent ping	Triggers (1)	agent.ping	60	7	365	Zabbix agent	Zabbix agent	Enabled	
☐	A test item - SNMP uptime		Uptime	30	7		SNMPv2 agent	SNMP data	Enabled	
☐	Template OS Linux: Available memory	Triggers (1)	vm.memory.size[available]	60	7	365	Zabbix agent	Memory	Enabled	
☐	Template OS Linux: Checksum of /etc/passwd	Triggers (1)	vfs.file.cksum[/etc/passwd]	3600	7	365	Zabbix agent	Security	Enabled	
☐	Template OS Linux: Context switches per second		system.cpu.switches	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU steal time		system.cpu.util[,steal]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU softirq time		system.cpu.util[,softirq]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU system time		system.cpu.util[,system]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU user time		system.cpu.util[,user]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU iowait time	Triggers (1)	system.cpu.util[,iowait]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU nice time		system.cpu.util[,nice]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU idle time		system.cpu.util[,idle]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Template OS Linux: CPU interrupt time		system.cpu.util[,interrupt]	60	7	365	Zabbix agent	CPU, Performance	Enabled	
☐	Mounted filesystem discovery : Free disk space on /		vfs.fs.size[/,free]	60	7	365	Zabbix agent	Filesystems	Enabled	
☐	Mounted filesystem discovery : Free disk space on / (percentage)	Triggers (1)	vfs.fs.size[/,pfree]	60	7	365	Zabbix agent	Filesystems	Enabled	
Enable selected Go (0)										

Displayed data:

Column	Description
Wizard	The wizard icon is a link to a wizard for creating a trigger based on the item.
Name	Name of the item, mostly displayed as a blue link except for items created from item prototypes. Clicking on the item name link opens the item configuration form . If the host item belongs to a template, the template name is displayed before the item name, as a grey link. Clicking on the template link will open the item list on the template level. If the item has been created from an item prototype, its name is preceded by the low level discovery rule name, in khaki. Clicking on the discovery rule name will open the item prototype list.
Triggers	Moving the mouse over Triggers will display an info box displaying the triggers associated with the item.
Key	The number of the triggers is displayed in parentheses. Item key is displayed.
Interval	Frequency of the check is displayed.
History	How many days item data history will be kept is displayed.
Trends	How many days item trends history will be kept is displayed.
Type	Item type is displayed (Zabbix agent, SNMP agent, simple check, etc).
Applications	Item applications are displayed.
Status	Item status is displayed - <i>Enabled</i> , <i>Disabled</i> or <i>Not supported</i> . By clicking on the status you can change it - from Enabled to Disabled or from Disabled and Not supported to Enabled.
Error	A green checkbox icon is displayed if there are no errors. A red square icon with a cross is displayed if there are errors. Move the mouse over the icon and you will see a tooltip with the error description. Only green icon is displayed if the item is disabled.

To configure a new item, click on the *Create item* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change item status to *Enabled*
- *Disable selected* - change item status to *Disabled*
- *Mass update* - **update several properties** for a number of items at once
- *Copy selected to...* - copy the items to other hosts or templates
- *Clear history for selected* - delete history and trend data for items

- *Delete selected* - delete the items

To use these options, mark the check-boxes before the respective items, then select the required option and click on "Go".

Filter

As the list may contain very many items, it may be needed to filter out the ones you really need.

The narrow blue bar just below the *Items* bar is actually a link to the filter. If you click on it, a filter becomes available where you can filter items by several properties.

CONFIGURATION OF ITEMS

Create item

Items

Displaying 1 to 2 of 2 found

Host group:

Select

Type:

all

Type of information:

all

Status:

all

Host:

New host

Select

Update interval (in sec):

Triggers:

all

Application:

Select

Keep history (in days):

Template:

all

Name like:

Keep trends (in days):

Key like:

Filter

Reset

Subfilter [affects only filtered data!]

Applications

CPU (0) , [Filesystems \(+2\)](#) , [General \(+3\)](#) , [Memory \(+2\)](#) , Network interfaces (0) , [OS \(+5\)](#) , [Performance \(2\)](#) , [Processes \(+2\)](#) , [Security \(+1\)](#) , SNMP data (0) , [Zabbix agent \(+3\)](#)

Types

SNMPv2 agent (0) , [Zabbix agent \(2\)](#)

Type of information

Character (0) , [Numeric \(float\) \(2\)](#) , Numeric (unsigned) (0) , Text (0)

Status

Disabled (0) , [Enabled \(2\)](#)

Template

Not Templated items (0) , [Templated items \(2\)](#)

With triggers

[Without triggers \(+11\)](#) , [With triggers \(2\)](#)

Interval

30 (0) , [60 \(2\)](#) , 600 (0) , 3600 (0)

« Host list

Host: [New host](#)

Monitored

☒

☐

☐

☐

[Applications \(11\)](#)

[Items \(42\)](#)

[Triggers \(18\)](#)

[Graphs \(7\)](#)

[Discovery rules \(2\)](#)

☐	Wizard	Name	Triggers	Key	Interval	History	Trends	Type	Applications	Status	Error
☐		Template OS Linux: CPU iowait time	Triggers (1)	system.cpu.util[iowait]	60	7	365	Zabbix agent	CPU, Performance	Enabled	☒
☐		Template OS Linux: Processor load (1 min average per core)	Triggers (2)	system.cpu.load[percpu,avg1]	60	7	365	Zabbix agent	CPU, Performance	Enabled	☒

The **Subfilter** below the filter offers further filtering options (for the data already filtered). The links in red are groups of items with a common parameter value. If you click on the link it turns green and only the items with this parameter value remain in the list.

3 Triggers

Overview

The trigger list for a template can be accessed from *Configuration* → *Templates* and then clicking on Triggers for the respective template.

The trigger list for a host can be accessed from *Configuration* → *Hosts* and then clicking on Triggers for the respective host.

A list of existing triggers is displayed. By default, only the enabled triggers are displayed. To display disabled triggers as well, use the *Show disabled triggers* link to the right in the Triggers bar.

CONFIGURATION OF TRIGGERS

Create trigger

Triggers

Group all Host New host

Displaying 1 to 18 of 18 found

[Hide disabled triggers]

« Host list

Host: New host

Monitored

Applications (11)

Items (43)

Triggers (18)

Graphs (7)

Discovery rules (2)

☐	Severity	Name	Expression	Status	Error
☐	Warning	Template OS Linux: /etc/passwd has been changed on {HOST.NAME}	{mnd desk:vfs.file.cksum[/etc/passwd].diff(0)}>0	Enabled	☒
☐	Information	Template OS Linux: Configured max number of opened files is too low on {HOST.NAME}	{mnd desk:kernel.maxfiles.last(0)}<1024	Enabled	☒
☐	Information	Template OS Linux: Configured max number of processes is too low on {HOST.NAME}	{mnd desk:kernel.maxproc.last(0)}<256	Enabled	☒
☐	Warning	Template OS Linux: Disk I/O is overloaded on {HOST.NAME}	{mnd desk:system.cpu.util[iowait].last(0)}>20	Enabled	☒
☐	Warning	Mounted filesystem discovery: Free disk space is less than 20% on volume /	{mnd desk:vfs.fs.size[/,pfree].last(0)}<20	Enabled	☒
☐	Warning	Mounted filesystem discovery: Free inodes is less than 20% on volume /	{mnd desk:vfs.fs.inode[/,pfree].last(0)}<20	Enabled	☒
☐	Information	Template OS Linux: Host information was changed on {HOST.NAME}	{mnd desk:system.uname.diff(0)}>0	Enabled	☒
☐	Information	Template App Zabbix Agent: Host name of zabbix_agentd was changed on {HOST.NAME}	{mnd desk:agent.hostname.diff(0)}>0	Enabled	☒
☐	Information	Template OS Linux: Hostname was changed on {HOST.NAME}	{mnd desk:system.hostname.diff(0)}>0	Enabled	☒
☐	Average	Template OS Linux: Lack of available memory on server {HOST.NAME}	{mnd desk:vm.memory.size[available].last(0)}<20M	Enabled	☒
☐	Warning	Template OS Linux: Lack of free swap space on {HOST.NAME}	{mnd desk:system.swap.size[pfree].last(0)}<50	Enabled	☒
☐	Information	Processor load has gone over 1 on {HOST.NAME}	{mnd desk:system.cpu.load[percpu,avg1].last(0)}>1	Enabled	☒
☐	Warning	Template OS Linux: Processor load is too high on {HOST.NAME}	{mnd desk:system.cpu.load[percpu,avg1].last(0)}>5	Enabled	☒
☐	Warning	Template OS Linux: Too many processes on {HOST.NAME}	{mnd desk:proc.num[,].last(0)}>300	Enabled	☒
☐	Warning	Template OS Linux: Too many processes running on {HOST.NAME}	{mnd desk:proc.num[,run].last(0)}>30	Enabled	☒
☐	Information	Template App Zabbix Agent: Version of zabbix_agentd was changed on {HOST.NAME}	{mnd desk:agent.version.diff(0)}>0	Enabled	☒
☐	Average	Template App Zabbix Agent: Zabbix agent on {HOST.NAME} is unreachable for 5 minutes	{mnd desk:agent.ping.nodata(5m)}=1	Enabled	☒
☐	Information	Template OS Linux: {HOST.NAME} has just been restarted	{mnd desk:system.uptime.change(0)}<0	Enabled	☒

Enable selected

Go (0)

Displayed data:

Column	Description
Severity	Severity of the trigger is displayed by both name and cell background colour.
Name	Name of the trigger, mostly displayed as a blue link except for triggers created from trigger prototypes. Clicking on the trigger name link opens the trigger configuration form . If the host trigger belongs to a template, the template name is displayed before the trigger name, as a grey link. Clicking on the template link will open the trigger list on the template level. If the trigger has been created from a trigger prototype, its name is preceded by the low level discovery rule name, in khaki. Clicking on the discovery rule name will open the trigger prototype list.
Expression	Trigger expression is displayed. The host-item part of the expression is displayed as a link, leading to the item configuration form.
Status	Trigger status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it.
Error	A green checkbox icon displayed if there are no errors. A red icon with a cross is displayed if there are errors.

To configure a new trigger, click on the *Create trigger* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change trigger status to *Enabled*
- *Disable selected* - change trigger status to *Disabled*
- *Mass update* - update several properties for a number of triggers at once
- *Copy selected to...* - copy the triggers to other hosts or templates
- *Delete selected* - delete the triggers

To use these options, mark the check-boxes before the respective triggers, then select the required option and click on "Go".

4 Graphs

Overview

The custom graph list for a template can be accessed from *Configuration* → *Templates* and then clicking on *Graphs* for the respective template.

The custom graph list for a host can be accessed from *Configuration* → *Hosts* and then clicking on *Graphs* for the respective host.

A list of existing graphs is displayed.

CONFIGURATION OF GRAPHS

Create graph

Graphs

Group allHost New host

Displaying 1 to 7 of 7 found

« Host list

Host: New host

Monitored

Applications (11)

Items (43)

Triggers (18)

Graphs (7)

Discovery rules (2)

☐	Name	Width	Height	Graph type
☐	Template OS Linux: CPU jumps	900	200	Normal
☐	Template OS Linux: CPU load	900	200	Normal
☐	Template OS Linux: CPU utilization	900	200	Stacked
☐	Mounted filesystem discovery: Disk space usage /	600	340	Pie
☐	Network interface discovery: Network traffic on eth0	900	200	Normal
☐	Network interface discovery: Network traffic on vboxnet0	900	200	Normal
☐	Template OS Linux: Swap usage	600	340	Pie

Copy selected to ...

Go (0)

Displayed data:

Column	Description
Name	Name of the custom graph, mostly displayed as a blue link except for graphs created from graph prototypes. Clicking on the graph name link opens the graph configuration form. If the host graph belongs to a template, the template name is displayed before the graph name, as a grey link. Clicking on the template link will open the graph list on the template level. If the graph has been created from a graph prototype, its name is preceded by the low level discovery rule name, in khaki. Clicking on the discovery rule name will open the graph prototype list.
Width	Graph width is displayed.
Height	Graph height is displayed.
Graph type	Graph type is displayed - Normal, Stacked, Pie or Exploded.

To configure a new graph, click on the *Create graph* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Copy selected to...* - copy the graphs to other hosts or templates
- *Delete selected* - delete the graphs

To use these options, mark the check-boxes before the respective graphs, then select the required option and click on "Go".

5 Discovery rules

Overview

The list of low-level discovery rules for a template can be accessed from *Configuration* → *Templates* and then clicking on *Discovery* for the respective template.

The list of low-level discovery rules for a host can be accessed from *Configuration* → *Hosts* and then clicking on *Discovery* for the respective host.

A list of existing low-level discovery rules is displayed.

CONFIGURATION OF DISCOVERY RULES										Create discovery rule
Discovery rules										
Displaying 1 to 2 of 2 found										
Host list Host: New host Monitored    Applications (11) Items (43) Triggers (18) Graphs (7) Discovery rules (2)										
☐	Name 	Items	Triggers	Graphs	Key	Interval	Type	Status	Error	
☐	Template OS Linux: Mounted filesystem discovery	Item prototypes (5)	Trigger prototypes (2)	Graph prototypes (1)	vfs.fs.discovery	3600	Zabbix agent	Enabled		
☐	Template OS Linux: Network interface discovery	Item prototypes (2)	Trigger prototypes (0)	Graph prototypes (1)	net.if.discovery	3600	Zabbix agent	Enabled		
Enable selected  Go (0)										

Displayed data:

Column	Description
<i>Name</i>	Name of the rule, displayed as a blue link. Clicking on the rule name opens the low-level discovery rule configuration form. If the discovery rule belongs to a template, the template name is displayed before the rule name, as a grey link. Clicking on the template link will open the rule list on the template level.
<i>Items</i>	A link to the list of item prototypes is displayed. The number of existing item prototypes is displayed in parentheses.
<i>Triggers</i>	A link to the list of trigger prototypes is displayed. The number of existing trigger prototypes is displayed in parentheses.
<i>Graphs</i>	A link to the list of graph prototypes displayed. The number of existing graph prototypes is displayed in parentheses.
<i>Key</i>	The item key used for discovery is displayed.
<i>Interval</i>	The frequency of performing discovery is displayed.
<i>Type</i>	The item type used for discovery is displayed (Zabbix agent, SNMP agent, etc).
<i>Status</i>	Discovery rule status is displayed - <i>Enabled</i> , <i>Disabled</i> or <i>Not supported</i> . By clicking on the status you can change it - from Enabled to Disabled or from Disabled and Not supported to Enabled.
<i>Error</i>	A green checkbox icon is displayed if there are no errors. A red square icon with a cross is displayed if there are errors.

To configure a new low-level discovery rule, click on the *Create discovery rule* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the low-level discovery rule status to *Enabled*
- *Disable selected* - change the low-level discovery rule status to *Disabled*
- *Delete selected* - delete the low-level discovery rules

To use these options, mark the check-boxes before the respective discovery rules, then select the required option and click on "Go".

4 Maintenance

Overview

In the *Configuration* → *Maintenance* section users can configure and maintain maintenance periods for hosts.

A listing of existing maintenance periods with their details is displayed.

From the dropdown to the right in the *Maintenance periods* bar you can choose whether to display all maintenance periods or only those belonging to one particular group.

CONFIGURATION OF MAINTENANCE PERIODS Create maintenance period

Maintenance periods Group

Displaying 1 to 1 of 1 found

☐	Name	Type	State	Description
☐	Weekly maintenance	No data collection	Active	Friday 16:00 to 17:00

Delete selected Go (0)

Displayed data:

Column	Description
<i>Name</i>	Name of the maintenance period. Clicking on the maintenance period name opens the maintenance period configuration form .
<i>Type</i>	The type of maintenance is displayed: <i>With data collection</i> or <i>No data collection</i>
<i>State</i>	The state of the maintenance period: Approaching - will become active soon Active - is active Expired - is not active any more
<i>Description</i>	Description of the maintenance period is displayed.

To configure a new maintenance period, click on the *Create maintenance period* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers one mass-editing option:

- *Delete selected* - delete the maintenance periods

To use this option, mark the check-boxes before the respective maintenance periods and click on "Go".

5 Web

Overview

In the *Configuration* → *Web* section users can configure and maintain web scenarios.

A listing of existing web scenarios with their details is displayed. Web scenarios are grouped by the applications they are part of.

From the dropdown to the right in the *Scenarios* bar you can choose whether to display all web scenarios or only those belonging to one particular group and host. Additionally you can choose to hide disabled scenarios (or show them again) by clicking on the respective link.

CONFIGURATION OF WEB MONITORING Create scenario

Scenarios Group Host

Displaying 1 to 1 of 1 found [\[Hide disabled scenarios \]](#)

☐	Name	Number of steps	Update interval	Status
☒ Web scenarios (1 scenario)				
☐	Availability of zabbix	4	60	Enabled

Enable selected Go (0)

Displayed data:

Column	Description
<i>Name</i>	Name of the web scenario. Clicking on the web scenario name opens the web scenario configuration form .
<i>Number of steps</i>	The number of steps contained in the scenario.
<i>Update interval</i>	How often the scenario is performed.
<i>Status</i>	Web scenario status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it.

To configure a new web scenario, click on the *Create scenario* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the scenario status to *Enabled*
- *Disable selected* - change the scenario status to *Disabled*
- *Clear history for selected* - clear history and trend data for the scenarios
- *Delete selected* - delete the web scenarios

To use these options, mark the check-boxes before the respective web scenarios, then select the required option and click on "Go".

6 Actions

Overview

In the *Configuration* → *Actions* section users can configure and maintain actions.

A listing of existing actions with their details is displayed. The actions displayed are actions assigned to the selected event source (triggers, discovery, auto-registration).

To view actions assigned to a different event source, change the source from the dropdown to the right in the *Actions* bar.

CONFIGURATION OF ACTIONS

Create action

Actions

Event source: Triggers

Displaying 1 to 1 of 1 found

☐	Name	Conditions	Operations	Status
☐	Problem notification	Maintenance status not in "maintenance" Trigger value = "PROBLEM"	Send message to user groups: Zabbix administrators Send message to user groups: Zabbix administrators Send message to user groups: Database manager Send message to user groups: Zabbix administrators Send message to user groups: Zabbix administrators	Enabled

Enable selected Go (0)

Displayed data:

Column	Description
<i>Name</i>	Name of the action. Clicking on the action name opens the action configuration form .
<i>Conditions</i>	Action conditions are displayed.
<i>Operations</i>	Action operations are displayed.
<i>Status</i>	Action status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it. If an action is disabled during an escalation in progress (like a message being sent), the message in progress will be sent and then one more message on the escalation will be sent. The follow-up message will have the following text at the beginning of the message body: <i>NOTE: Escalation cancelled: action '<Action name>' disabled</i> . This way the recipient is informed that the escalation is cancelled and no more steps will be executed.

To configure a new action, click on the *Create action* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the action status to *Enabled*
- *Disable selected* - change the action status to *Disabled*
- *Delete selected* - delete the actions

To use these options, mark the check-boxes before the respective actions, then select the required option and click on "Go".

7 Screens

Overview

In the *Configuration* → *Screens* section users can configure and maintain screens.

A listing of existing screens with their details is displayed.

☐	Name	Dimension (cols x rows)	Screen
☐	Another screen	3 x 5	Edit
☐	Zabbix server	2 x 2	Edit

Export selected Go (0)

Displayed data:

Column	Description
Name	Name of the screen. By clicking on the screen name you can open the grid of screen elements for editing.
Dimensions	The number of columns and rows of the screen.
Screen	Click on the <i>Edit</i> link to edit general screen properties (name and dimensions).

To create a **new screen**, click on the *Create screen* button in the top right-hand corner. To import a screen from an XML file, click on the *Import* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Export selected* - export the screens to an XML file
- *Delete selected* - delete the screens

To use these options, mark the check-boxes before the respective screens, then select the required option and click on "Go".

8 Slide shows

Overview

In the *Configuration* → *Slide shows* section users can configure and maintain slide shows.

A listing of existing slide shows with their details is displayed.

☐	Name	Delay	Count of slides
☐	First slide show	15	2

Delete selected Go (0)

Displayed data:

Column	Description
Name	Name of the slide show. Clicking on the slide show name opens the slide show configuration form .
Delay	The default duration of showing one slide is displayed.
Count of slides	The number of slides in the slide show is displayed.

To configure a new slide show, click on the *Create slide show* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers one mass-editing option:

- *Delete selected* - delete the slide shows

To use this option, mark the check-boxes before the respective slide shows and click on "Go".

9 Maps

Overview

In the *Configuration* → *Maps* section users can configure and maintain network maps.

A listing of existing maps with their details is displayed.

CONFIGURATION OF NETWORK MAPS

Create map

Import

Maps

Displaying 1 to 1 of 1 found

☐	Name	Width	Height	Edit
☐	Local network	680	300	Edit

Export selected

Go (0)

Displayed data:

Column	Description
<i>Name</i>	Name of the map. By clicking on the map name you can access the grid for adding map elements .
<i>Width</i>	Map width is displayed.
<i>Height</i>	Map height is displayed.
<i>Edit</i>	Click on the <i>Edit</i> link to edit general map properties.

To create a **new map**, click on the *Create map* button in the top right-hand corner. To import a map from an XML file, click on the *Import* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Export selected* - export the maps to an XML file
- *Delete selected* - delete the maps

To use these options, mark the check-boxes before the respective maps, then select the required option and click on "Go".

10 Discovery

Overview

In the *Configuration* → *Discovery* section users can configure and maintain discovery rules.

A listing of existing discovery rules with their details is displayed.

CONFIGURATION OF DISCOVERY RULE					Create discovery rule
Discovery rule					
Displaying 1 to 1 of 1 found					
☐	<u>Name</u> ↑	<u>IP range</u>	<u>Delay</u>	<u>Checks</u>	<u>Status</u>
☐	Local network_Z	192.168.0.1-127	3600	SNMPv2 agent, Zabbix agent	Enabled
Enable selected Go (0)					

Displayed data:

Column	Description
<i>Name</i>	Name of the discovery rule. Clicking on the discovery rule name opens the discovery rule configuration form .
<i>IP range</i>	The range of IP addresses to use for network scanning is displayed.
<i>Delay</i>	The frequency of performing discovery displayed.
<i>Checks</i>	The types of checks used for discovery are displayed.
<i>Status</i>	Action status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it.

To configure a new discovery rule, click on the *Create discovery rule* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the discovery rule status to *Enabled*
- *Disable selected* - change the discovery rule status to *Disabled*
- *Delete selected* - delete the discovery rules

To use these options, mark the check-boxes before the respective discovery rules, then select the required option and click on "Go".

11 IT services

Overview

In the *Configuration* → *IT services* section users can configure and maintain an IT services hierarchy.

When you first open this section it only contains a *root* entry.

You can use it as a starting point of building the hierarchy of monitored infrastructure. Click on it and add services and then other services below the ones you have added.

CONFIGURATION OF IT SERVICES		
IT services		
Service	Status calculation	Trigger
root		
[-] Network equipment	Problem, if at least one child has a problem	-
[-] Switch	Switch 1 has a problem	-
[-] Switch	Add service	-
[-] Applications	Edit service	-
[-] VMs	Delete service	-
[-] Workstations	Do not calculate	-
[-] Host1	Problem, if at least one child has a problem	Zabbix agent on New host unreachable for 5min
[-] Host2	Problem, if at least one child has a problem	Zabbix agent on Zabbix server unreachable for 5min

For details on adding services, see the **IT services** section.

5 Administration

Overview

The Administration menu is for administrative functions of Zabbix. This menu is available to users of **Super Administrators** type only.

1 General

Overview

The *Administration* → *General* section contains a number of screens for setting frontend-related defaults and customizing Zabbix.

The dropdown to the right allows you to switch between different configuration screens.

ZABBIX
Help | Get support | Print | Profile | Logout

Monitoring | Inventory | Reports | Configuration | Administration

General | DM | Authentication | Users | Media types | Scripts | Audit | Queue | Notifications | Installation

History: Dashboard » Configuration of GUI

CONFIGURATION OF GUI

GUI

GUI
Housekeeper
Images
Icon mapping
Regular expressions
Macros
Value mapping
Working time
Trigger severities
Trigger displaying options
Other

Default theme
Original blue

Dropdown first entry
All
remember selected

Search/Filter elements limit
1000

Max count of elements to show inside table cell
50

Enable event acknowledges
☒

Show events not older than (in days)
7

Max count of events per trigger to show
100

Show warning if Zabbix server is down
☒

Save

1 GUI

This screen provides customization of several frontend-related defaults.

GUI

Default theme
Original blue

Dropdown first entry
All
remember selected

Search/Filter elements limit
1000

Max count of elements to show inside table cell
50

Enable event acknowledges
☒

Show events not older than (in days)
7

Max count of events per trigger to show
100

Show warning if Zabbix server is down
☒

Configuration parameters:

Parameter	Description
<i>Default theme</i>	Default theme for users who have not set a specific one in their profiles.
<i>Dropdown first entry</i>	Whether first entry in element selection dropdowns should be <i>All</i> or <i>None</i> . With <i>remember selected</i> checked, the last selected element in the dropdown will be remembered (instead of the default) when navigating to another page.
<i>Search/Filter elements limit</i>	Maximum amount of elements (rows) that will be displayed in a web-interface list, like, for example, in <i>Monitoring</i> → <i>Events</i> or <i>Configuration</i> → <i>Hosts</i> . <i>Note:</i> If set to, for example, '50', only the first 50 elements will be displayed in all affected frontend lists. If some list contains more than fifty elements, the indication of that will be the '+' sign in "Displaying 1 to 50 of 50+ found". Also, if filtering is used and still there are more than 50 matches, only the first 50 will be displayed.

Parameter	Description
<i>Max count of elements
to show inside table cell</i>	For entries that are displayed in a single table cell, no more than configured here will be shown.
<i>Enable event acknowledges</i>	This parameter defines if event acknowledgments are activated in Zabbix interface.
<i>Show events not older
than (in days)</i>	This parameter defines for how many days events are displayed in Status of Triggers screen. Default is 7 days.
<i>Max count of events per trigger to show</i>	Maximum number of event to show for each trigger in Status of Triggers screen. Default is 100.
<i>Show warning if Zabbix server is down</i>	This parameter enables a warning message to be displayed atop the browser window if Zabbix server cannot be reached (may be down). The message remains visible even if the user has scrolled down the page. If the mouse is moved over it, the message is temporarily hidden to reveal the contents below. This parameter is supported since Zabbix 2.0.1 version.

2 Housekeeper

The housekeeper is a periodical process, executed by Zabbix server. The process removes outdated information and information deleted by user.


This screen allows to set for how long information will be kept before it is deleted by the housekeeper.

Configuration parameters:

Parameter	Description
<i>Do not keep actions older than (in days)</i>	This parameter defines how many days of executed action (e-mail, jabber, SMS, etc) history Zabbix will keep in the database. Older actions will be removed.
<i>Do not keep events older than (in days)</i>	This parameter defines how many days of event history Zabbix will keep in the database. Older events will be removed.

3 Images

The Images section displays all the images available in Zabbix. Images are stored in the database.



The *Type* dropdown allows you to switch between icon and background images:

- Icons are used to display **network map** elements
- Backgrounds are used as background images of network maps

Adding image

You can add your own image by clicking on the *Create image* button in the top right corner.

Image

Name

Type

Icon

Upload

Browse...

Image attributes:

Parameter	Description
<i>Name</i>	Unique name of an image.
<i>Type</i>	Set Icon or Background type.
<i>Upload</i>	Select the file (PNG, JPEG) from a local system to be uploaded to Zabbix.

Note:

Maximum size of the upload file is limited by value of `ZBX_MAX_IMAGE_SIZE` that is 1024x1024 bytes or 1 MB.

The upload of an image may fail if the image size is close to 1 MB and the `max_allowed_packet` MySQL configuration parameter is at a default of 1MB. In this case, increase the `max_allowed_packet` parameter.

4 Icon mapping

This section allows to create the mapping of certain hosts with certain icons. Host inventory field information is used to create the mapping.

The mappings can then be used in **network map configuration** to assign appropriate icons to matching hosts automatically. To create a new icon map, click on *Create icon map* in the top right corner.

Icon map

Name

Host type

Mappings

Inventory field

Expression

Icon

1: Type

server

Rackmountable_5U_server_3D_(96)

Remove

2: Type

router

Router_symbol_(96)

Remove

3: Type

workstation

Workstation_(96)

Remove

Add

Default

Server_(96)

Configuration parameters:

Parameter	Description
Name	Unique name of icon map.
Mappings	A list of mappings. The order of mappings determines which one will have priority. You can move mappings up and down the list with drag-and-drop.
Inventory field	Host inventory field that will be looked into to seek a match.
Expression	Regular expression describing the match.
Icon	Icon to use if a match for the expression is found.
Default	Default icon to use.

5 Regular expressions

This section allows to create custom regular expressions that can be used in several places in the frontend. See **Regular expressions** section for details.

6 Macros

This section allows to define system-wide macros.

Macros

Macro

Value

{SNMP_COMMUNITY}

→ public

Remove

{SMACRO}

→ value

Remove

Add

See **User macros** section for more details.

7 Value mapping

This section allows to create value maps that allow for human-readable representation of incoming data in Zabbix frontend. See **Value mapping** section for more details.

8 Working time

Working time is system-wide parameter, which defines working time. Working time is displayed as a white background in graphs, while non-working time is displayed in grey.

Working time

Working time

1-5,09:00-18:00;

See **Time period specification** page for description of the time format.

9 Trigger severities

This section allows to customize **trigger severity** names and colors.

Trigger severities

	Custom severity	Colour
Not classified	>Custom name here<	DBDBDB
Information	Information	D6F6FF
Warning	Warning	FFF6A5
Average	Average	FFB689
High	High	FF9999
Disaster	Disaster	FF3838

Info

Custom severity names affect all locales and require r

Cancel

Save

Reset defaults

You can enter new names and color codes or click on the color to select another from the provided palette.

See **Customising trigger severities** page for more information.

10 Trigger displaying options

This section allows to customize how trigger status is displayed in the frontend.

Trigger displaying options

	Colour	Blinking
Unacknowledged PROBLEM events	DC0000	☒
Acknowledged PROBLEM events	DC0000	☒
Unacknowledged OK events	00AA00	☒
Acknowledged OK events	00AA00	☒

Display OK triggers for

1800

seconds

On status change triggers blink for

1800

seconds

Save

Reset defaults

The colors for acknowledged/unacknowledged events can be customized and blinking enabled or disabled. Also the time period for displaying OK triggers and for blinking upon trigger status change can be customized.

11 Other parameters

This section allows to configure several other frontend parameters.

Other parameters

Refresh unsupported items (in sec)

600

Group for discovered hosts

Discovered hosts

User group for database down message

Zabbix administrators

Log unmatched SNMP traps

☒

Parameter	Description
<i>Refresh unsupported items (in sec)</i>	Some items may become unsupported due to errors in user parameters or because of an item not being supported by agent. Zabbix can be configured to periodically make unsupported items active. Zabbix will activate unsupported item every N seconds set here. If set to 0, the automatic activation will be disabled. The configured value also applies to how often Zabbix proxies reactivate unsupported items.
<i>Group for discovered hosts</i>	Hosts discovered by network discovery and agent auto-registration will be automatically placed in the host group, selected here.
<i>User group for database down message</i>	User group for sending alarm message or 'None'. Availability of Zabbix server depends on availability of backend database. It cannot work without a database. Database watchdog , a special Zabbix server process, will alarm selected users in case of disaster. If the database is down, the watchdog will send notifications to the user group set here, using all configured user media entries. Zabbix server will not stop; it will wait until the database is back again to continue processing. <i>Note:</i> Until Zabbix version 1.8.2 database watchdog was supported for MySQL only. Since 1.8.2, it is supported for all databases.
<i>Log unmatched SNMP traps</i>	Log SNMP trap if no corresponding SNMP interfaces have been found.

2 DM

Overview

In the *Administration* → *DM* section **distributed monitoring** options (proxies or nodes) can be configured in the Zabbix front-end. The dropdown in the top right-hand corner allows to switch between proxy or node screens.

Proxies

A listing of existing proxies with their details is displayed.

CONFIGURATION OF PROXIES

Proxies

Create proxy

Proxies

Displaying 1 to 1 of 1 found

☐	Name 	Mode	Last seen (age)	Host count	Item count	Required performance (vps)	Hosts
☐	W1 proxy	Active	-	6	201	2.87	New host , srv_01 , srv_02 , srv_03 , srv_04 , srv_05

Enable selected 

Go (0)

Displayed data:

Column	Description
<i>Name</i>	Name of the proxy. Clicking on the proxy name opens the proxy configuration form .
<i>Mode</i>	Proxy mode is displayed - <i>Active</i> or <i>Passive</i> .
<i>Last seen (age)</i>	The time when the proxy was last seen by the server is displayed.
<i>Host count</i>	The number of hosts monitored by the proxy is displayed.
<i>Item count</i>	The number of items monitored by the proxy is displayed.
<i>Required performance (vps)</i>	Required proxy performance is displayed (the number of values that need to be collected per second).
<i>Hosts</i>	All hosts monitored by the proxy are listed. Clicking on the host name opens the host configuration form.

To configure a new proxy, click on the *Create proxy* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the status of hosts monitored by the proxy to *Monitored*
- *Disable selected* - change the status of hosts monitored by the proxy to *Not monitored*
- *Delete selected* - delete the proxies

To use these options, mark the check-boxes before the respective proxies, then select the required option and click on "Go".

Nodes

If node setup is not configured, this screen is empty and displays a "Your setup is not configured for distributed monitoring" message.

See how to configure a **node setup**.

3 Authentication

Overview

In *Administration* → *Authentication* the user authentication method to Zabbix can be changed. The available methods are internal, LDAP and HTTP authentication.

By default, internal Zabbix authentication is used. To change, click on the button with the method name and press *Save*.

Internal

Internal Zabbix authentication is used.

LDAP

External LDAP authentication can be used to check user names and passwords. Note that a user must exist in Zabbix as well, however its Zabbix password will not be used.

Zabbix LDAP authentication works at least with Microsoft Active Directory and OpenLDAP.

LDAP authentication

Default authentication

Internal

LDAP

HTTP

LDAP host

Port

Base DN

Search attribute

Bind DN

Bind password

Test authentication [must be a valid LDAP user]

Login

User password

Save

Test

Configuration parameters:

Parameter	Description
<i>LDAP host</i>	Name of LDAP server. For example: ldap://ldap.zabbix.com For secure LDAP server use ldaps protocol. ldaps://ldap.zabbix.com With OpenLDAP 2.x.x and later, a full LDAP URI of the form ldap://hostname:port or ldaps://hostname:port may be used.
<i>Port</i>	Port of LDAP server. Default is 389. For secure LDAP connection port number is normally 636.
<i>Base DN</i>	Not used when using full LDAP URIs. Base path to search accounts: ou=Users,ou=system (for OpenLDAP), DC=company,DC=com (for Microsoft Active Directory)
<i>Search attribute</i>	LDAP account attribute used for search: uid (for OpenLDAP), sAMAccountName (for Microsoft Active Directory)
<i>Bind DN</i>	LDAP account for binding and searching over the LDAP server, examples: uid=ldap_search,ou=system (for OpenLDAP), CN=ldap_search,OU=user_group,DC=company,DC=com (for Microsoft Active Directory)
<i>Bind password</i>	Required, anonymous binding is not supported. LDAP password of the account for binding and searching over the LDAP server.
<i>Test authentication</i>	Header of a section for testing
<i>Login</i>	Name of a test user (which is currently logged in the Zabbix frontend). This user name must exist in the LDAP server. Zabbix will not activate LDAP authentication if it is unable to authenticate the test user.
<i>User password</i>	LDAP password of the test user.

Note:

It is recommended to create a separate LDAP account (*Bind DN*) to perform binding and searching over the LDAP server with minimal privileges in the LDAP instead of using real user accounts (used for logging in the Zabbix frontend). Such an approach provides more security and does not require changing the *Bind password* when the user changes his own password in the LDAP server.

In the table above it's *ldap_search* account name.

Note:

Some user groups can still be authenticated by Zabbix. These groups must have **frontend access** set to Internal.

HTTP

Apache-based (HTTP) authentication can be used to check user names and passwords. Note that a user must exist in Zabbix as well, however its Zabbix password will not be used.

Attention:

Be careful! Make sure that Apache authentication is configured and works properly before switching it on.

Note:

In case of Apache authentication all users (even with **frontend access** set to Internal) will be authenticated by Apache, not by Zabbix!

4 Users**Overview**

In the *Administration* → *Users* section both user groups and users of the system are maintained.

By default the user group screen is displayed. To switch to the user screen and back, use the dropdown in the top right-hand corner.

User groups

A listing of existing user groups with their details is displayed.

CONFIGURATION OF USERS AND USER GROUPS					
					User groups ▾ Create user group
User groups					
Displaying 1 to 6 of 6 found					
☐	Name ↕	#	Members	Status	Frontend access
☐	Disabled	Users (0)		Disabled	System default
☐	Enabled debug mode	Users (0)		Enabled	System default
☐	Guests	Users (1)	guest	Enabled	System default
☐	IT Management	Users (0)		Enabled	System default
☐	No access to the frontend	Users (0)		Enabled	Disabled
☐	Zabbix administrators	Users (1)	Admin	Enabled	System default
Enable selected ▾ Go (0)					

Displayed data:

Column	Description
<i>Name</i>	Name of the user group. Clicking on the user group name opens the user group configuration form .
<i>#</i>	The number of users in the group (displayed in parentheses). Clicking on <i>Users</i> will display the respective users filtered out in the user list.
<i>Members</i>	Aliases of individual users in the user group. Clicking on the alias will open the user configuration form.
<i>Status</i>	User group status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it.

Column	Description
<i>Frontend access</i>	Frontend access level is displayed: System default - Zabbix, LDAP or HTTP authentication; depending on the chosen authentication method Internal - the user is authenticated by Zabbix regardless of system settings Disabled - frontend access for this user is disabled. By clicking on the current level you can change it.
<i>Debug mode</i>	Debug mode status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it.

To configure a new user group, click on the *Create user group* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the user group status to *Enabled*
- *Disable selected* - change the user group status to *Disabled*
- *Enable DEBUG* - enable debug mode for the user groups
- *Disable DEBUG* - disable debug mode for the user groups
- *Delete selected* - delete the user groups

To use these options, mark the check-boxes before the respective user groups, then select the required option and click on "Go".

Users

A listing of existing users with their details is displayed.

CONFIGURATION OF USERS AND USER GROUPS

Users

Users

Create user

Users

User group All

Displaying 1 to 2 of 2 found

☐	Alias ↑	Name	Surname	User type	Groups	Is online?	Login	Frontend access	Debug mode	Status
☐	Admin	Zabbix	Administrator	Zabbix Super Admin	Zabbix administrators	Yes (Mon, 24 Sep 2012 10:59:21 +0300)	Ok	System default	Disabled	Enabled
☐	guest	Default	User	Zabbix User	Guests	No (Wed, 19 Sep 2012 10:21:30 +0300)	Ok	System default	Disabled	Enabled

Unblock selected

Go (0)

From the dropdown to the right in the *Users* bar you can choose whether to display all users or those belonging to one particular group.

Displayed data:

Column	Description
<i>Alias</i>	Alias of the user, used for logging into Zabbix. Clicking on the alias opens the user configuration form .
<i>Name</i>	First name of the user.
<i>Surname</i>	Second name of the user.
<i>User type</i>	User type is displayed - <i>Zabbix Super Admin</i> , <i>Zabbix Admin</i> or <i>Zabbix User</i> .
<i>Groups</i>	Groups that the user is member of are listed. Clicking on the user group name opens the user group configuration form.
<i>Is online?</i>	The on-line status of the user is displayed - <i>Yes</i> or <i>No</i> . The time of last user activity is displayed in parentheses.
<i>Login</i>	The login status of the user is displayed - <i>Ok</i> or <i>Blocked</i> . A user can become temporarily blocked upon more than five unsuccessful login attempts. By clicking on <i>Blocked</i> you can unblock the user.
<i>Frontend access</i>	Frontend access level is displayed - <i>System default</i> , <i>Internal</i> or <i>Disabled</i> , depending on the one set for the whole user group.
<i>Debug mode</i>	Debug mode status is displayed - <i>Enabled</i> or <i>Disabled</i> , depending on the one set for the whole user group.
<i>Status</i>	User status is displayed - <i>Enabled</i> or <i>Disabled</i> , depending on the one set for the whole user group.

To configure a new user, click on the *Create user* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Unblock selected* - re-enable system access to blocked users
- *Delete selected* - delete the users

To use these options, mark the check-boxes before the respective users, then select the required option and click on "Go".

5 Media types

Overview

In the *Administration* → *Media types* section users can configure and maintain media type information.

Media type information contains general instructions for using a medium as delivery channel for notifications. Specific details, such as the individual e-mail addresses to send a notification to are kept with individual users.

A listing of existing media types with their details is displayed.

CONFIGURATION OF MEDIA TYPES					Create media type
Media types					
Displaying 1 to 3 of 3 found					
☐	Description	Type	Status	Used in actions	Details
☐	Email	Email	Enabled	Report problems to Zabbix administrators	SMTP server: "mail.company.com", SMTP helo: "company.com", SMTP email: "zabbix@company.com"
☐	Jabber	Jabber	Enabled	-	Jabber identifier: "jabber@company.com"
☐	SMS	SMS	Enabled	-	GSM modem: "/dev/ttyS0"
Enable selected Go (0)					

Displayed data:

Column	Description
Description	Description of the media type. Clicking on the description opens the media type configuration form .
Type	Type of the media (e-mail, SMS, etc) is displayed.
Status	Media type status is displayed - <i>Enabled</i> or <i>Disabled</i> . By clicking on the status you can change it.
Used in actions	All actions where the media type is used directly (selected in the <i>Send only to</i> dropdown) are displayed. Clicking on the action name opens the action configuration form.
Details	Detailed information of the media type is displayed.

To configure a new media type, click on the *Create media type* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers some mass-editing options:

- *Enable selected* - change the media type status to *Enabled*
- *Disable selected* - change the media type status to *Disabled*
- *Delete selected* - delete the media types

To use these options, mark the check-boxes before the respective media types, then select the required option and click on "Go".

6 Scripts

Overview

In the *Administration* → *Scripts* section user-defined scripts can be configured and maintained.

These scripts, depending on the set user permissions, then become available for execution by clicking on the host in various frontend locations (*Dashboard*, *Latest data*, *Status of triggers*, *Events*, *Maps*) and can also be run as an action operation. The scripts are executed on the Zabbix server or agent.

A listing of existing scripts with their details is displayed.

CONFIGURATION OF SCRIPTS Create script

Scripts

Displaying 1 to 3 of 3 found

☐ Name	Type	Execute on	Commands	User group	Host group	Host access
☐ Detect operating system	Script	Server	sudo /usr/bin/nmap -O {HOST.CONN} 2>&1	Zabbix administrators	All	Read
☐ Ping	Script	Server	/bin/ping -c 3 {HOST.CONN} 2>&1	All	All	Read
☐ Traceroute	Script	Server	/usr/bin/traceroute {HOST.CONN} 2>&1	All	All	Read

Delete selected
Go (0)

Displayed data:

Column	Description
Name	Name of the script. Clicking on the script name opens the script configuration form .
Type	Script type is displayed - <i>Script</i> or <i>IPMI</i> command.
Execute on	It is displayed whether the script will be executed on Zabbix server or agent.
Commands	All commands to be executed within the script are displayed.
User group	The user group that the script is available to is displayed (or <i>All</i> for all user groups).
Host group	The host group that the script is available for is displayed (or <i>All</i> for all host groups).
Host access	The permission level for the host group is displayed - <i>Read</i> or <i>Write</i> . Only users with the required permission level will have access to executing the script.

To configure a new script, click on the *Create script* button in the top right-hand corner.

Mass editing options

A dropdown below the list offers one mass-editing option:

- *Delete selected* - delete the scripts

To use this option, mark the check-boxes before the respective scripts and click on "Go".

Configuring a script

Script

Name

Type

Script

Execute on

☐ Zabbix agent
☒ Zabbix server

Commands

Description

User groups

Zabbix administrators

Host groups

All

Required host permissions

Read

Enable confirmation

☐

Confirmation text

Test confirmation

Save

Clone

Delete

Cancel

Script attributes:

Parameter	Description
<i>Name</i>	Unique name of the script.
<i>Type</i>	Select script type - Script or IPMI command . A special dropdown selection for scripts containing IPMI commands is available since Zabbix 2.0 version (previously a special syntax of <i>IPMI <command></i> had to be used in the command field).
<i>Execute on</i>	Select the radio button whether to execute the script on Zabbix server or agent. The option to execute scripts on Zabbix agent is available since Zabbix 2.0 version (providing remote commands are enabled in the EnableRemoteCommands parameter in Zabbix agent configuration file).
<i>Commands</i>	Enter full path to the commands to be executed within the script. The following macros are supported in the commands: {HOST.CONN}, {HOST.IP}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}. If a macro may resolve to a value with spaces (for example, host name), don't forget to quote as needed.
<i>Description</i>	Enter a description for the script.
<i>User groups</i>	Select the user group that the script will be available to (or <i>All</i> for all user groups).
<i>Host groups</i>	Select the host group that the script will be available for (or <i>All</i> for all host groups).
<i>Required host permissions</i>	Select the permission level for the host group - <i>Read</i> or <i>Write</i> . Only users with the required permission level will have access to executing the script.
<i>Enable confirmation</i>	Mark the checkbox to display a confirmation message before executing the script. This feature might be especially useful with potentially dangerous operations (like a reboot script) or ones that might take a long time.
<i>Confirmation text</i>	Enter a custom confirmation text for the confirmation popup enabled with the checkbox above (for example, <i>Remote system will be rebooted. Are you sure?</i>). To see how the text will look like, click on <i>Test confirmation</i> next to the field.

7 Audit

Overview

In the *Administration* → *Audit* section users can view records of changes made in the frontend and details of executed actions.

By default frontend audit records are displayed. To switch to action details and back, use the dropdown in the top right-hand corner.

Logs

In this screen audit logs of various changes made in the frontend can be seen. You can use the filter, located below the *Logs* bar, to narrow down the records by user, activity type, affected resource and the time period.

AUDIT LOGS							
Logs							
Displaying 1 to 12 of 12 found							
Filter User Admin Select Action Update ▼ Resource All ▼ Filter Reset							
Zoom: 1h 2h 3h 6h 12h 1d All 24.09.2012 11:03 - 25.09.2012 11:03 01d 00h 00m 00:00							
Time	User	IP	Resource	Action	ID	Description	Details
25 Sep 2012 09:55:32	Admin	192.168.3.39	Action	Updated	0		Name: Report problems to Zabbix administrators 2
25 Sep 2012 09:54:53	Admin	192.168.3.39	Action	Updated	0		Name: Report problems to Zabbix administrators 2
25 Sep 2012 09:53:30	Admin	192.168.3.39	Action	Updated	0		Name: Report problems to Zabbix administrators 2
25 Sep 2012 09:52:09	Admin	192.168.3.39	Action	Updated	0		Actions [3] disabled
25 Sep 2012 09:42:20	Admin	192.168.3.39	Media type	Updated	0		Media type [Email]
25 Sep 2012 09:40:21	Admin	192.168.3.39	Action	Updated	0		Name: Report problems to Zabbix administrators
25 Sep 2012 09:39:57	Admin	192.168.3.39	Action	Updated	0		Name: Report problems to Zabbix administrators
25 Sep 2012 09:38:42	Admin	192.168.3.39	User	Updated	0		User alias [Admin] name [Zabbix] surname [Administrator]
25 Sep 2012 09:37:57	Admin	192.168.3.39	Action	Updated	0		Name: Report problems to Zabbix administrators
24 Sep 2012 14:58:38	Admin	192.168.3.39	User group	Updated	0		Group name [Guests]
24 Sep 2012 14:57:45	Admin	192.168.3.39	User group	Updated	0		Group name [Guests]
24 Sep 2012 14:51:53	Admin	192.168.3.39	User	Updated	0		Unblocked user alias [Newbie] name [Newbie] surname [First]

Displayed data:

Column	Description
<i>Time</i>	Timestamp of the audit record.
<i>User</i>	User of the activity.
<i>IP</i>	IP that was used in the activity.
<i>Resource</i>	Affected resource is displayed.
<i>Action</i>	Activity type is displayed - <i>Login, Logout, Added, Updated, Deleted, Enabled or Disabled.</i>
<i>ID</i>	ID of the affected resource is displayed.
<i>Description</i>	Description of the resource is displayed.
<i>Details</i>	Detailed information on the performed activity is displayed.

Actions

In this screen details of executed actions (notifications or remote commands) are displayed.

You can use the filter, located below the *Actions* bar, to narrow down the records by recipient of e-mail and time period.

AUDIT ACTIONS							Actions
Actions							
Displaying 1 to 12 of 12 found							
Filter							
Recipient Select							
Filter Reset							
Zoom: 1h 2h 3h 6h 12h 1d All 26.09.2012 01:02 - 26.09.2012 13:02 (now)							
1d 12h 1h 1h 12h 1d 30m 12h 00m (fixed)							
Time	Type	Status	Retries left	Recipient(s)	Message	Error	
26 Sep 2012 13:02:19	Email	In progress	3	daredevil_x@inbox.lv	Subject: OK: Processor load has gone over 1 on New host Message: Trigger: Processor load has gone over 1 on New host Trigger status: OK Trigger severity: Information Trigger URL: Item values: 1. Processor load (1 min average per core) (New host:system.cpu.load[percpu,avg1]): 0.855		
26 Sep 2012 12:59:18		executed			Command: Zabbix server::capache2 restart		
26 Sep 2012 12:55:17	Email	sent		daredevil_x@inbox.lv	Subject: PROBLEM: Processor load has gone over 1 on New host Message: Trigger: Processor load has gone over 1 on New host Trigger status: PROBLEM Trigger severity: Information Trigger URL: Item values: 1. Processor load (1 min average per core) (New host:system.cpu.load[percpu,avg1]): 1.095		

Displayed data:

Column	Description
<i>Time</i>	Timestamp of the action.
<i>Type</i>	Action type is displayed - <i>Email</i> or <i>Command</i> .
<i>Status</i>	Action status is displayed: <i>in progress</i> - action is in progress <i>sent</i> - notification has been sent <i>executed</i> - command has been executed <i>not sent</i> - action has not been completed
<i>Retries left</i>	The remaining number of times the server will try to send the notification is displayed.
<i>Recipient(s)</i>	Notification recipient(s) e-mail addresses are displayed.
<i>Message</i>	The content of the message/remote command is displayed.
<i>Error</i>	Error information regarding the action execution is displayed.

8 Queue

Overview

In the *Administration* → *Queue* section items that are waiting to be updated are displayed.

Ideally, when you open this section it should all be "green" meaning no items in the queue. If all items are updated without delay, there are none waiting. However, due to lacking server performance, connection problems or problems with agents, some items may get delayed and the information is displayed in this section. For more details, see the **Queue** section.

From the dropdown in the upper right corner you can select:

- queue overview by item type
- queue overview by proxy
- list of delayed items

Overview by item type

In this screen it is easy to locate if the problem is related to one or several item types.

QUEUE OF ITEMS TO BE UPDATED							Overview
Items	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 m	Overview Overview by proxy Details
Zabbix agent	0	0	0	24	1	0	
Zabbix agent (active)	0	0	0	0	0	0	
Simple check	0	0	0	0	0	0	
SNMPv1 agent	0	0	0	0	0	0	
SNMPv2 agent	0	0	0	0	0	0	
SNMPv3 agent	0	0	0	0	0	0	
Zabbix internal	0	0	0	0	0	0	
Zabbix aggregate	0	0	0	0	0	0	
External check	0	0	0	0	0	0	
Database monitor	0	0	0	0	0	0	
IPMI agent	0	0	0	0	0	0	
SSH agent	0	0	0	0	0	0	
TELNET agent	0	0	0	0	0	0	
JMX agent	0	0	0	0	0	0	
Calculated	0	0	0	0	0	0	

Each line contains an item type. Each column shows the number of waiting items - waiting for 5-10 seconds/10-30 seconds/30-60 seconds/1-5 minutes/5-10 minutes or over 10 minutes respectively.

Overview by proxy

In this screen it is easy to locate if the problem is related to one of the proxies or the server.

QUEUE OF ITEMS TO BE UPDATED							Overview by proxy
Proxy	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes	
W1 proxy	0	0	0	0	0	0	
Server	0	0	0	26	0	0	
Total: 2							

Each line contains a proxy, with the server last in the list. Each column shows the number of waiting items - waiting for 5-10 seconds/10-30 seconds/30-60 seconds/1-5 minutes/5-10 minutes or over 10 minutes respectively.

List of waiting items

In this screen, each waiting item is listed.

QUEUE OF ITEMS TO BE UPDATED				Details
Next check	Delayed by	Host	Name	
27 Sep 2012 10:50:15	5m 45s	New host	Processor load (5 min average per core)	
27 Sep 2012 10:50:16	5m 44s	New host	Context switches per second	
27 Sep 2012 10:50:17	5m 43s	New host	CPU idle time	
27 Sep 2012 10:50:32	5m 28s	New host	Number of logged in users	
27 Sep 2012 10:51:09	4m 51s	New host	Number of running processes	
27 Sep 2012 10:51:10	4m 50s	New host	Number of processes	
27 Sep 2012 10:51:12	4m 48s	New host	Interrupts per second	
27 Sep 2012 10:51:14	4m 46s	New host	Processor load (1 min average per core)	
27 Sep 2012 10:53:13	2m 47s	New host	Processor load (15 min average per core)	
27 Sep 2012 10:53:18	2m 42s	New host	CPU interrupt time	
27 Sep 2012 10:53:19	2m 41s	New host	CPU iowait time	
27 Sep 2012 10:53:20	2m 40s	New host	CPU nice time	
27 Sep 2012 10:53:21	2m 39s	New host	CPU softirq time	
27 Sep 2012 10:53:22	2m 38s	New host	CPU steal time	
27 Sep 2012 10:53:23	2m 37s	New host	CPU system time	
27 Sep 2012 10:53:24	2m 36s	New host	CPU user time	
27 Sep 2012 10:53:26	2m 34s	New host	Host local time	
27 Sep 2012 10:53:27	2m 33s	New host	Free swap space	
27 Sep 2012 10:53:28	2m 32s	New host	Free swap space in %	
Total: 19				

Displayed data:

Column	Description
<i>Next check</i>	The time when the check was due is displayed.
<i>Delayed by</i>	The length of the delay is displayed.
<i>Host</i>	Host of the item is displayed.
<i>Name</i>	Name of the waiting item is displayed.

9 Notifications

Overview

In the *Administration* → *Notifications* section a report on the number of notifications sent to each user is displayed.

From the dropdowns in the top right-hand corner you can choose the media type (or all), period (data for each day/week/month/year) and year for the notifications sent.

Notifications			
Month	Admin	guest	
Jan 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Feb 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Mar 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Apr 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
May 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Jun 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Jul 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Aug 2012	0 (0/0/0)	0 (0/0/0)	0 (0/0/0)
Sep 2012	29 (29/0/0)	0 (0/0/0)	0 (0/0/0)

all ([Email](#) / [Jabber](#) / [SMS](#))

Each column displays totals per one system user.

10 Installation

Overview

In the *Administration* → *Installation* section Zabbix frontend can be **reinstalled**.



To continue with the installation, click on *Next*. To exit the installation, click on *Cancel*.

2 User profile

Overview

In the user profile you can customize some Zabbix frontend features, such as the interface language, color theme, number of rows displayed in the lists etc. The changes made here will apply for the user only.

To access the user profile configuration form, click on **Profile** in the upper right corner of Zabbix window.

Configuration

The **User** tab allows you to set various user preferences.

UserMediaMessaging

Password

Change password

Language

English (en_GB)

Theme

System default

Auto-login

☒

Auto-logout (min 90 seconds)

☐

90

Refresh (in seconds)

30

Rows per page

50

URL (after login)

Save

Cancel

Parameter	Description
Password	Click on the link to display two fields for entering a new password.
Language	Select the interface language of your choice.
Theme	Select a color theme specifically for your profile.
Auto-login	With this checkbox marked you will be logged in automatically, without having to enter your username and password.
Auto-logout (min 90 seconds)	With this checkbox marked you will be logged out automatically, after the set amount of seconds (minimum 90 seconds). However, this option will not work if the "Show warning if Zabbix server is down" global configuration option (since Zabbix 2.0.1) is enabled and Zabbix frontend is kept opened or when, for example, Monitoring menu pages perform background information refreshes. Additionally, Auto-logout will be deactivated in the user profile when logging in with Remember me for 30 days checked.
Refresh (in seconds)	You can set how often the information in the pages will be refreshed on the Monitoring menu, except for Dashboard, which uses its own refresh parameters for every widget.
Rows per page	You can set how many rows will be displayed per page in the lists. Fewer rows (and fewer records to display) mean faster loading times.

Parameter	Description
URL (after login)	You can set a specific URL to be displayed after the login. Instead of the default <i>Monitoring → Dashboard</i> it can be, for example, the URL of <i>Monitoring → Triggers</i> .

Note:

If some language is not available for selection in the user profile it means that a locale for it is not installed on the web server. See the [link](#) at the bottom of this page to find out how to install them.

The **Media** tab allows you to specify the **media** details for the user, such as the types, the addresses to use and when to use them to deliver notifications.

User	Media	Messaging
	Media	
☐	Email	user@company.com
		1-5,09:00-18:00;
		NIWAHD
		Enabled
		Edit
☐	Jabber	jabber@company.com
		1-7,00:00-24:00;
		NIWAHD
		Enabled
		Edit
Add Delete selected		

Note:

Only **admin level** users (Admin and Super Admin) can change their own media details.

The **Messaging** tab allows you to set **global notifications**.

See also

1. [How to install additional locales to be able to select unavailable languages in the user profile](#)

1 Global notifications

Overview

Global notifications are a way of displaying issues that are currently happening right on the screen you're at in Zabbix frontend.

Without global notifications, working in some other location than *Status of triggers* or *Dashboard* pages would not show any information regarding issues that are currently happening. Global notifications will display this information regardless of where you are.

Global notifications involve both showing a message and **playing a sound**.

Configuration

Global notifications can be enabled per user in the *Messaging* tab of **profile configuration**.

UserMediaMessaging

Frontend messaging ☒

Message timeout (seconds)

60

Play sound

Once

Trigger severity

☒ Recovery

no_sound

Play

Stop

☒ Not classified

no_sound

Play

Stop

☒ Information

no_sound

Play

Stop

☒ Warning

no_sound

Play

Stop

☒ Average

no_sound

Play

Stop

☒ High

no_sound

Play

Stop

☒ Disaster

no_sound


Play

Stop

Parameter	Description
<i>Frontend messaging</i>	Mark the checkbox to enable global notifications.
<i>Message timeout</i>	You can set for how long the message will be displayed. By default, messages will stay on screen for 60 seconds.
<i>Play sound</i>	You can set how long the sound will be played. Once - sound is played once and fully. 10 seconds - sound is repeated for 10 seconds. Message timeout - sound is repeated while the message is visible.
<i>Trigger severity</i>	You can set the trigger severities that global notifications and sounds will be activated for. You can also select the sounds appropriate for various severities. If no severity is marked then no messages will be displayed at all. Also, recovery messages will only be displayed for those severities that are marked. So if you mark <i>Recovery</i> and <i>Disaster</i> , global notifications will be displayed for the problems and the recoveries of disaster severity triggers.

Global messages displayed

As the messages arrive, they are displayed in a floating section on the right hand side. This section can be repositioned vertically.



For this section, several controls are available:

-  **Move** button allows to reposition the section vertically. This can also be done by dragging section header;
-  **Snooze** button silences currently active alarm sound;
-  **Mute/Unmute** button switches between playing and not playing the alarm sounds;
-  **Clear** button removes all currently visible messages.

2 Sound in browsers

Overview

For the sounds to be played in Zabbix frontend, *Frontend messaging* must be enabled in the user profile *Messaging* tab, with all trigger severities checked, and sounds should also be enabled in the global notification pop-up window.

The sounds of Zabbix frontend have been successfully tested in the following web browser versions and no additional configuration was required:

- Firefox 3.5.16 on Linux
- Opera 11.01 on Linux
- Google Chrome 9.0 on Windows
- Firefox 3.5.16 on Windows
- IE7 browser on Windows
- Opera v11.01 on Windows
- Chrome v9.0 on Windows
- Safari v5.0 on Windows, but this browser requires *Quick Time Player* to be installed

Additional requirements

Firefox v 3.5.16

For playing wav files in the Firefox browser you can use one of the following applications:

- Windows Media Player
- Quick Time plug-in.

Then, in *Tools* → *Options* → *Applications*, in "Wave sound (audio/wav)" set Windows Media Player to play these files.

Safari 5.0

Quick Time Player is required.

Microsoft Internet Explorer

To play sounds in MSIE7 and MSIE8:

- In *Tools* → *Internet Options* → *Advanced* enable *Play sounds in webpages*
- In *Tools* → *Manage Add-ons...* enable **Windows Media Player**
- In the Windows Media Player, in *Tools*→*Options*→*File Types* enable *Windows audio file (wav)*

In the Windows Media Player, in *Tools*→*Options* tab, "File Types" is only available if the user is a member of "Power Users" or "Administrators" group, i.e. a regular user does not have access to this tab and does not see it.

An additional thing - if IE does not have some *.wav file in the local cache directory (%userprofile%\Local Settings\Temporary Internet Files) the sound will not play the first time.

Known not to work


Browsers where the sound did not work:

- Opera 10.11 on Linux.

3 Global search

It is possible to search for various entities in the Zabbix frontend. Search input box is located in the upper right corner. Search can be started by pressing enter or clicking on the *Search* button.

If there is a host that starts with the entered string, a dropdown will appear, listing all such hosts:



In the search results, it is possible to collapse each individual block. Enabled hosts will be displayed in blue, disabled hosts in red.

Entities searched It is possible to search for these entities and their properties:

- Hosts
 - Visible name (or host name if visible name is not defined)
 - IP address
 - DNS name
- Templates
 - Name
- Host groups
 - Name

Links available For entities found the following links are available:

- Hosts
 - Monitoring
 - * Latest data
 - * Triggers
 - * Events
 - * Templated screens
 - Configuration
 - * Host properties
 - * Applications
 - * Items
 - * Triggers
 - * Graphs
- Host groups
 - Monitoring

- * Latest data
- * Triggers
- * Events
- Configuration
 - * Host group properties
 - * Host group members (hosts and templates; separate links since Zabbix 2.0.2)
- Templates
 - Configuration
 - * Template properties
 - * Applications
 - * Items
 - * Triggers
 - * Graphs
 - * Templated screens

Below each block the amount of entities found and displayed is shown, for example, *Displaying 13 of 13 found*. The amount of displayed entries in each block is limited to 100.

For all configuration entities amount of entities found is displayed in parenthesis. If no entities of that type are found, the entry is greyed out and link is not provided.

4 Frontend maintenance mode

Overview

Zabbix web frontend can be temporarily disabled in order to prohibit access to it. This can be useful for protecting the Zabbix database from any changes initiated by users, thus protecting the integrity of database.

Zabbix database can be stopped and maintenance tasks can be performed while Zabbix frontend is in maintenance mode.

Users from a defined range of IP addresses will be able to work with the frontend normally during maintenance mode.

Configuration

In order to enable maintenance mode, the `maintenance.inc.php` file (located in `/conf` of the Zabbix HTML document directory on the webserver) must be modified to uncomment the following lines:

```
// Maintenance mode
define('ZBX_DENY_GUI_ACCESS',1);

// IP range, who allowed to connect to FrontEnd
$ZBX_GUI_ACCESS_IP_RANGE = array('127.0.0.1');

// MSG showed on Warning screen!
$_REQUEST['warning_msg'] = 'We are upgrading MySQL database till 15:00. Stay tuned...';
```

Parameter	Details
ZBX_DENY_GUI_ACCESS	Enable maintenance mode: 1 - maintenance mode is enabled, disabled otherwise
ZBX_GUI_ACCESS_IP_RANGE	Connections from these IP addresses will be allowed during the maintenance mode. For example: 192.168.1.1-255
warning_msg	A message you can enter to inform users about the maintenance.

Display

The following screen will be displayed when trying to access the Zabbix frontend while in maintenance mode. The screen is refreshed every 30 seconds in order to return to a normal state without user intervention when the maintenance is over.



IP addresses defined in `ZBX_GUI_ACCESS_IP_RANGE` will be able to access the frontend as always.

5 Definitions

Overview

While many things in the frontend can be configured using the frontend itself, some customisations are currently only possible by editing a definitions file.

This file is `defines.inc.php` located in `/include` of the Zabbix HTML document directory.

Parameters

Parameters in this file that could be of interest to users:

- `ZBX_LOGIN_ATTEMPTS`

Number of unsuccessful login attempts that is allowed to an existing system user before a login block is applied (see `ZBX_LOGIN_BLOCK`). By default 5 attempts. Once the set number of login attempts is tried unsuccessfully, each additional unsuccessful attempt results in a login block. Used with **internal** authentication only.

- `ZBX_LOGIN_BLOCK`

Number of seconds for blocking a user from accessing Zabbix frontend after a number of unsuccessful login attempts (see `ZBX_LOGIN_ATTEMPTS`). By default 30 seconds. Used with **internal** authentication only.

- `ZBX_PERIOD_DEFAULT`

Default graph period, in seconds. One hour by default.

- `ZBX_MIN_PERIOD`

Minimum graph period, in seconds. One hour by default.

- `ZBX_MAX_PERIOD`

Maximum graph period, in seconds. Two years by default since 1.6.7, one year before that.

- `GRAPH_Y_AXIS_SIDE_DEFAULT`

Default location of Y axis in simple graphs and default value for drop down box when adding items to custom graphs. Possible values: 0 - left, 1 - right.

Default: 0

- `ZBX_UNITS_ROUND_OFF_THRESHOLD`

Threshold value for roundoff constants. Values less than it will be rounded to ZBX_UNITS_ROUNDOFF_LOWER_LIMIT number of digits after comma, greater to ZBX_UNITS_ROUNDOFF_UPPER_LIMIT.

Default: 0.01

- ZBX_UNITS_ROUNDOFF_UPPER_LIMIT

Number of digits after comma, when value is greater than roundoff threshold

Default: 2

- ZBX_UNITS_ROUNDOFF_LOWER_LIMIT

Number of digits after comma, when value is less than roundoff threshold

Default: 6

- ZBX_HISTORY_DATA_UPKEEP (available since 1.8.4)

Number of days, which will reflect on frontend choice when deciding which history or trends table to process for selected period on data graphing. When this define is:

- * less than zero - Zabbix takes item values for selected graph period configured in item "keep in history";
- * equal to zero - Zabbix takes item values only from trends;
- * greater than zero - Zabbix overwrites item "keep in history" configured value with this define;

This define could be useful for partitioned history data storage.

Default: -1

- DEFAULT_LATEST_ISSUES_CNT

Controls how many issues are shown in the dashboard's *Last n issues* widget. By default 20 issues are shown.

- SCREEN_REFRESH_TIMEOUT (available since 2.0.4)

Used in screens and defines the timeout seconds for a screen element update. When the defined number of seconds after launching an update pass and the screen element has still not been updated, the screen element will be darkened.

Default: 30

- SCREEN_REFRESH_RESPONSIVENESS (available since 2.0.4)

Used in screens and defines the number of seconds after which query skipping will be switched off. Otherwise, if a screen element is in update status all queries on update are skipped until a response is received. With this parameter in use, another update query might be sent after N seconds without having to wait for the response to the first one.

Default: 10

6 Creating your own theme

Overview

By default, Zabbix provides a number of predefined themes. You may follow the step-by-step procedure provided here in order to create your own. Feel free to share result of your work with Zabbix community if you created something nice.

Step 1

To define your own theme you'll need to create a CSS file and save it as *styles/themes/mytheme/main.css*. You can either copy the files from a different theme and create your theme based on it or start from scratch. The rules in the *main.css* file will extend the ones that are defined in the base Zabbix CSS files located in the *styles* folder. Any theme-specific images must be placed in the *styles/themes/mytheme/images* folder.

Step 2

Add your theme to the list of themes returned by the `Z::getThemes()` method. You can do this by overriding the `ZBase::getThemes()` method in the `Z` class. This can be done by adding the following code before the closing brace in *include/classes/core/Z.php*:

```
public static function getThemes() {
    return array_merge(parent::getThemes(), array(
        'mytheme' => _('My theme')
    ));
}
```

Attention:

Note that the name you specify within the first pair of quotes must match the name of the directory under which the theme files have been saved.

To add multiple themes, just list them under the first theme, for example:

```
public static function getThemes() {
    return array_merge(parent::getThemes(), array(
        'mytheme' => _('My theme'),
        'anothertheme' => _('Another theme'),
        'onemoretheme' => _('One more theme')
    ));
}
```

Note that every theme except the last one must have a trailing comma.

Note:

To change graph colours, entry must be added in the database table *graph_theme*.

Step 3

Activate the new theme.

In Zabbix GUI, you may either set this theme to be the default one or change your theme in the user profile.

Enjoy the new look and feel!

7 Debug mode**Overview**

Debug mode may be used to diagnose performance problems with frontend pages.

Configuration

Debug mode can be activated for individual users who belong to a user group:

- when configuring a **user group**;
- when viewing configured **user groups**.

When *Debug mode* is enabled for a user group, its users will see a *Debug* button in the upper right corner of the browser window.

Clicking on the *Debug* button opens a new window below the page contents which contains the SQL statistics of the page, along with a list of API calls and individual SQL statements:

In case of performance problems with the page, this window may be used to search for the root cause of the problem.

Warning:

Enabled *Debug mode* negatively affects frontend performance.

16. Appendixes

Please use the sidebar to access content in the Appendixes section.

Zabbix API

The Zabbix API is a powerful tool for extending Zabbix, developing apps for your favorite platform or integrating it with third party software

This section provides an overview of the functions provided by the Zabbix API and will help you find your way around the available classes and methods.

Monitoring The Zabbix API allows you to access history and other data gathered during monitoring.

History

Retrieve historical values gathered by Zabbix monitoring processes for presentation or further processing.

[History API](#)

Events

Retrieve events generated by triggers, network discovery and other Zabbix systems for more flexible situation management or third-party tool integration.

[Event API](#)

Service monitoring

Retrieve detailed service layer availability information about any IT service.

[IT service SLA calculation](#)

Configuration The Zabbix API allows you to manage the configuration of your monitoring system.

Hosts and host groups

Manage host groups, hosts and everything related to them, including host interfaces, host macros and maintenance periods.

[Host API](#) | [Host group API](#) | [Host interface API](#) | [User macro API](#) | [Maintenance API](#)

Items and applications

Define items to monitor. Create or remove applications and assign items to them.

[Item API](#) | [Application API](#)

Triggers

Configure triggers to notify you about problems in your system. Manage trigger dependencies.

[Trigger API](#)

Graphs

Edit graphs or separate graph items for better presentation of the gathered data.

[Graph API](#) | [Graph item API](#)

Templates

Manage templates and link them to hosts or other templates.

[Template API](#)

Export and import

Export and import Zabbix configuration data for configuration backups, migration or large-scale configuration updates.

[Configuration API](#)

Low-level discovery

Configure low-level discovery rules as well as item, trigger and graph prototypes to monitor dynamic entities.

[LLD rule API](#) | [Item prototype API](#) | [Trigger prototype API](#) | [Graph prototype API](#)

Screens

Edit global and template-level screens or each screen item individually.

[Screen API](#) | [Screen item API](#) | [Template screen API](#) | [Template screen item API](#)

Actions and alerts

Define actions and operations to notify users about certain events or automatically execute remote commands. Gain access to information about generated alerts and their receivers.

[Action API](#) | [Alert API](#)

IT services

Manage IT services for service-level monitoring and retrieve detailed SLA information about any service.

IT service API

Maps

Configure maps to create detailed dynamic representations of your IT infrastructure.

Map API

Web monitoring

Configure web scenarios to monitor your web applications and services.

Web scenario API

Network discovery

Manage network-level discovery rules to automatically find and monitor new hosts. Gain full access to information about discovered services and hosts.

[Discovery rule API](#) | [Discovery check API](#) | [Discovery host API](#) | [Discovery service API](#)

Administration With the Zabbix API you can change administration settings of your monitoring system.

Users

Add users that will have access to Zabbix, assign them to user groups and grant permissions. Configure media types and the ways users will receive alerts.

[User API](#) | [User group API](#) | [Media type API](#) | [Media API](#)

General

Change certain global configuration options.

[Icon map API](#) | [Image API](#) | [User macro API](#)

Proxies

Manage the proxies used in your distributed monitoring setup.

Proxy API

Scripts

Configure and execute scripts to help you with your daily tasks.

Script API

API information Retrieve the version of the Zabbix API so that your application could use version-specific features.

API info API

Action

This class is designed to work with actions.

Object references:

- [Action](#)
- [Action condition](#)
- [Action operation](#)

Available methods:

- [action.create](#) - create new actions
- [action.delete](#) - delete actions
- [action.exists](#) - check if an action exists
- [action.get](#) - retrieve actions
- [action.update](#) - update actions

> Action object

The following objects are directly related to the `action` API.

Action

The action object has the following properties.

Property	Type	Description
<code>actionid</code>	string	(<i>readonly</i>) ID of the action.
<code>esc_period</code> (required)	integer	Default operation step duration. Must be greater than 60 seconds.
<code>evaltype</code> (required)	integer	Action condition evaluation method. Possible values: 0 - AND / OR; 1 - AND; 2 - OR.
<code>eventsource</code> (required)	integer	Type of events that the action will handle. Refer to the event "source" property for a list of supported event types.
<code>name</code> (required)	string	Name of the action.
<code>def_longdata</code>	string	Problem message text.
<code>def_shortdata</code>	string	Problem message subject.
<code>r_longdata</code>	string	Recovery message text.
<code>r_shortdata</code>	string	Recovery message subject.
<code>recovery_msg</code>	integer	Whether recovery messages are enabled. Possible values: 0 - (<i>default</i>) disabled; 1 - enabled.
<code>status</code>	integer	Whether the action is enabled or disabled. Possible values: 0 - (<i>default</i>) enabled; 1 - disabled.

Action condition

The action condition object defines a condition that must be met to perform the configured action operations. It has the following properties.

Property	Type	Description
<code>conditionid</code>	string	(<i>readonly</i>) ID of the action condition.

Property	Type	Description
conditiontype (required)	integer	Type of condition. Possible values for trigger actions: 0 - host group; 1 - host; 2 - trigger; 3 - trigger name; 4 - trigger severity; 5 - trigger value; 6 - time period; 13 - host template; 15 - application; 16 - maintenance status; 17 - node. Possible values for discovery actions: 7 - host IP; 8 - discovered service type; 9 - discovered service port; 10 - discovery status; 11 - uptime or downtime duration; 12 - received value; 18 - discovery rule; 19 - discovery check; 20 - proxy; 21 - discovery object. Possible values for auto-registration actions: 20 - proxy; 22 - host name.
value (required)	string	Value to compare with.
actionid	string	(<i>readonly</i>) ID of the action that the condition belongs to.
operator	integer	Condition operator. Possible values: 0 - (<i>default</i>) =; 1 - <>; 2 - like; 3 - not like; 4 - in; 5 - >=; 6 - <=; 7 - not in.

The following operators and values are supported for each condition type.

Condition	Condition name	Supported operators	Expected value
0	Host group	=, <>	Host group ID.
1	Host	=, <>	Host ID.
2	Trigger	=, <>	Trigger ID.
3	Trigger name	like, not like	Trigger name.
4	Trigger severity	=, <>, >=, <=	Trigger severity. Refer to the trigger "severity" property for a list of supported trigger severities.
5	Trigger value	=	Trigger value. Refer to the trigger "value" property for a list of supported trigger values.

Condition	Condition name	Supported operators	Expected value
6	Time period	in, not in	Time when the event was triggered as a time period .
7	Host IP	=, <>	One or several IP ranges to check separated by commas. Refer to the network discovery configuration section for more information on supported formats of IP ranges.
8	Discovered service type	=, <>	Type of discovered service. The type of service matches the type of the discovery check used to detect the service. Refer to the discovery check "type" property for a list of supported types.
9	Discovered service port	=, <>	One or several port ranges separated by commas.
10	Discovery status	=	Status of a discovered object. Possible values: 0 - host or service up; 1 - host or service down; 2 - host or service discovered; 3 - host or service lost.
11	Uptime or downtime duration	>=, <=	Time indicating how long has the discovered object been in the current status in seconds.
12	Received values	=, <>, >=, <=, like, not like	Value returned when performing a Zabbix agent, SNMPv1, SNMPv2 or SNMPv3 discovery check.
13	Host template	=, <>	Linked template ID.
15	Application	=, like, not like	Name of the application.
16	Maintenance status	in, not in	No value required: using the "in" operator means that the host must be in maintenance, "not in" - not in maintenance.
17	Node	=, <>	ID of the distributed monitoring node.
18	Discovery rule	=, <>	ID of the discovery rule.
19	Discovery check	=, <>	ID of the discovery check.
20	Proxy	=, <>	ID of the proxy.
21	Discovery object	=	Type of object that triggered the discovery event. Possible values: 1 - discovered host; 2 - discovered service.
22	Host name	like, not like	Host name.

Action operation

The action operation object defines an operation that will be performed when an action is executed. It has the following properties.

Property	Type	Description
operationid	string	(<i>readonly</i>) ID of the action operation.
operationtype (required)	integer	Type of operation. Possible values: 0 - send message; 1 - remote command; 2 - add host; 3 - remove host; 4 - add to host group; 5 - remove from host group; 6 - link to template; 7 - unlink from template; 8 - enable host; 9 - disable host.
actionid	string	ID of the action that the operation belongs to.
esc_period	integer	Duration of an escalation step in seconds. Must be greater than 60 seconds. If set to 0, the default action escalation period will be used.
esc_step_from	integer	Default: 0. Step to start escalation from.
esc_step_to	integer	Default: 1. Step to end escalation at.
evaltype	integer	Default: 1. Operation condition evaluation method.
opcommand	object	Possible values: 0 - (<i>default</i>) AND / OR; 1 - AND; 2 - OR. Object containing the data about the command run by the operation. The operation command object is described in detail below .
opcommand_grp	array	Required for remote command operations. Host groups to run remote commands on. Each object has the following properties: opcommand_grpid - (<i>string, readonly</i>) ID of the object; operationid - (<i>string</i>) ID of the operation; groupid - (<i>string</i>) ID of the host group.
opcommand_hst	array	Required for remote command operations if opcommand_hst is not set. Host to run remote commands on. Each object has the following properties: opcommand_hstid - (<i>string, readonly</i>) ID of the object; operationid - (<i>string</i>) ID of the operation; hostid - (<i>string</i>) ID of the host; if set to 0 the command will be run on the current host.
opconditions	array	Required for remote command operations if opcommand_grp is not set. Operation conditions used for trigger actions. The operation condition object is described in detail below .

Property	Type	Description
opgroup	array	Host groups to add hosts to. Each object has the following properties: operationid - (<i>string</i>) ID of the operation; groupid - (<i>string</i>) ID of the host group.
opmessage	object	Required for "add to host group" and "remove from host group" operations. Object containing the data about the message sent by the operation. The operation message object is described in detail below .
opmessage_grp	array	Required for message operations. User groups to send messages to. Each object has the following properties: operationid - (<i>string</i>) ID of the operation; usrgrp - (<i>string</i>) ID of the user group.
opmessage_usr	array	Required for message operations if opmessage_usr is not set. Users to send messages to. Each object has the following properties: operationid - (<i>string</i>) ID of the operation; userid - (<i>string</i>) ID of the user.
optemplate	array	Required for message operations if opmessage_grp is not set. Templates to link the hosts to to. Each object has the following properties: operationid - (<i>string</i>) ID of the operation; templateid - (<i>string</i>) ID of the template. Required for "link to template" and "unlink from template" operations.

Action operation command

The operation command object contains data about the command that will be run by the operation.

Property	Type	Description
operationid	string	(<i>readonly</i>) ID of the operation.
command (required)	string	Command to run.
type (required)	integer	Type of operation command. Possible values: 0 - custom script; 1 - IPMI; 2 - SSH; 3 - Telnet; 4 - global script.

Property	Type	Description
authtype	integer	Authentication method used for SSH commands. Possible values: 0 - password; 1 - public key.
execute_on	integer	Required for SSH commands. Target on which the custom script operation command will be executed. Possible values: 0 - Zabbix agent; 1 - Zabbix server.
password	string	Required for custom script commands. Password used for SSH commands with password authentication and Telnet commands.
port	string	Port number used for SSH and Telnet commands.
privatekey	string	Name of the private key file used for SSH commands with public key authentication.
publickey	string	Required for SSH commands with public key authentication. Name of the public key file used for SSH commands with public key authentication.
scriptid	string	Required for SSH commands with public key authentication. ID of the script used for global script commands.
username	string	Required for global script commands. User name used for authentication.
		Required for SSH and Telnet commands.

Action operation message

The operation message object contains data about the message that will be sent by the operation.

Property	Type	Description
operationid	string	(<i>readonly</i>) ID of the action operation.
default_msg	integer	Whether to use the default action message text and subject. Possible values: 0 - (<i>default</i>) use the data from the operation; 1 - use the data from the action.
mediatypeid	string	ID of the media type that will be used to send the message.
message	string	Operation message text.
subject	string	Operation message subject.

Action operation condition

The action operation condition object defines a condition that must be met to perform the current operation. It has the following properties.

Property	Type	Description
opconditionid	string	(<i>readonly</i>) ID of the action operation condition

Property	Type	Description
conditiontype (required)	integer	Type of condition. Possible values: 14 - event acknowledged.
value (required)	string	Value to compare with.
operationid	string	(<i>readonly</i>) ID of the operation.
operator	integer	Condition operator. Possible values: 0 - (<i>default</i>) =.

The following operators and values are supported for each operation condition type.

Condition	Condition name	Supported operators	Expected value
14	Event acknowledged	=	Whether the event is acknowledged. Possible values: 0 - not acknowledged; 1 - acknowledged.

action.create

Description

`object action.create(object/array actions)`

This method allows to create new actions.

Parameters

(object/array) Actions to create.

Additionally to the **standard action properties**, the method accepts the following parameters.

Parameter	Type	Description
operations (required)	array	Action operations to create for the action.
conditions	array	Action conditions to create for the action.

Return values

(object) Returns an object containing the IDs of the created actions under the `actionids` property. The order of the returned IDs matches the order of the passed actions.

Examples

Create a trigger action

Create an action that will be run when a trigger from host "30045" that has the word "memory" in its name goes into problem state. The action must first send a message to all users in user group "7". If the event is not resolved in 4 minutes, it will run script "3" on all hosts in group "2".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "action.create",
  "params": {
    "name": "Trigger action",
    "eventsources": 0,
    "evaltype": 0,
```



```

"status": 0,
"esc_period": 120,
"def_shortcode": "{TRIGGER.NAME}: {TRIGGER.STATUS}",
"def_longdata": "{TRIGGER.NAME}: {TRIGGER.STATUS}\r\nLast value: {ITEM.LASTVALUE}\r\n\r\n{TRIGGER.",
"conditions": [
    {
        "conditiontype": 1,
        "operator": 0,
        "value": "30045"
    },
    {
        "conditiontype": 3,
        "operator": 2,
        "value": "memory"
    }
],
"operations": [
    {
        "operationtype": 0,
        "esc_period": 0,
        "esc_step_from": 1,
        "esc_step_to": 2,
        "evaltype": 0,
        "opmessage_grp": [
            {
                "usrgrp": "7"
            }
        ],
        "opmessage": {
            "default_msg": 1,
            "mediatypeid": "1"
        }
    },
    {
        "operationtype": 1,
        "esc_step_from": 3,
        "esc_step_to": 4,
        "evaltype": 0,
        "opconditions": [
            {
                "conditiontype": 14,
                "operator": 0,
                "value": "0"
            }
        ],
        "opcommand_grp": [
            {
                "groupid": "2"
            }
        ],
        "opcommand": {
            "type": 4,
            "scriptid": "3"
        }
    }
],
"auth": "038e1d7b1735c6a5436ee9eae095879e",
"id": 1
}

```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "actionids": [
      "17"
    ]
  },
  "id": 1
}
```

Create a discovery action

Create an action that will link discovered hosts to template "30085".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "action.create",
  "params": {
    "name": "Discovery action",
    "eventsources": 1,
    "status": 0,
    "esc_period": 0,
    "evaltype": 0,
    "conditions": [
      {
        "conditiontype": 21,
        "value": "1"
      },
      {
        "conditiontype": 10,
        "value": "2"
      }
    ],
    "operations": [
      {
        "esc_step_from": 1,
        "esc_period": 0,
        "optemplate": [
          {
            "templateid": "30085"
          }
        ],
        "operationtype": 6,
        "esc_step_to": 1
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "actionids": [
      "18"
    ]
  },
  "id": 1
}
```

See also

- Action condition
- Action operation

Source

CAction::create() in *frontends/php/api/classes/CAction.php*.

action.delete

Description

object action.delete(array actionIds)

This method allows to delete actions.

Parameters

(array) IDs of the actions to delete.

Return values

(object) Returns an object containing the IDs of the deleted actions under the `actionids` property.

Examples

Delete multiple actions

Delete two actions.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "action.delete",
  "params": [
    "17",
    "18"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "actionids": [
      "17",
      "18"
    ]
  },
  "id": 1
}
```

Source

CAction::delete() in *frontends/php/api/classes/CAction.php*.

action.exists

Description

boolean action.exists(object filter)

This method checks if at least one action that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
actionid	string/array	IDs of actions.
name	string/array	Names of actions.
node	string	Name of the node the actions must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the actions must belong to.

Return values

(boolean) Returns true if at least one action that matches the given filter criteria exists.

Examples

Check action by name

Check if an action named "Auto discovery. Linux servers." exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "action.exists",
  "params": {
    "name": "Auto discovery. Linux servers."
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

CAction::exists() in *frontends/php/api/classes/CAction.php*.

action.get

Description

integer/array action.get(object parameters)

The method allows to retrieve actions according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
actionids	string/array	Return only actions with the given IDs.
groupids	string/array	Return only actions that use the given host groups in action conditions.
hostids	string/array	Return only actions that use the given hosts in action conditions.
triggerids	string/array	Return only actions that use the given triggers in action conditions.
mediatypeids	string/array	Return only actions that use the given media types to send messages.
usrgrpids	string/array	Return only actions that are configured to send messages to the given user groups.

Parameter	Type	Description
userids	string/array	Return only actions that are configured to send messages to the given users.
scriptids	string/array	Return only actions that are configured to run the given scripts.
selectConditions	query	Return an object containing the action conditions in the <code>conditions</code> property.
selectOperations	query	Does not support <code>shorten</code> or arrays of properties. Return an object containing the action operations in the <code>operations</code> property.
sortfield	string/array	Does not support <code>shorten</code> or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>actionid</code> , <code>name</code> and <code>status</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve discovery actions

Retrieve all configured discovery actions together with action conditions and operations.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "action.get",
  "params": {
    "output": "extend",
    "selectOperations": "extend",
    "selectConditions": "extend",
    "filter": {
      "eventsources": 1
    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
```

```

"result": [
  {
    "actionid": "2",
    "name": "Auto discovery. Linux servers.",
    "eventsources": "1",
    "evaltype": "0",
    "status": "1",
    "esc_period": "0",
    "def_shortdata": "",
    "def_longdata": "",
    "recovery_msg": "0",
    "r_shortdata": "",
    "r_longdata": "",
    "conditions": {
      "2": {
        "conditionid": "2",
        "actionid": "2",
        "conditiontype": "10",
        "operator": "0",
        "value": "0"
      },
      "3": {
        "conditionid": "3",
        "actionid": "2",
        "conditiontype": "8",
        "operator": "0",
        "value": "9"
      },
      "4": {
        "conditionid": "4",
        "actionid": "2",
        "conditiontype": "12",
        "operator": "2",
        "value": "Linux"
      }
    },
    "operations": {
      "1": {
        "operationid": "1",
        "actionid": "2",
        "operationtype": "6",
        "esc_period": "0",
        "esc_step_from": "1",
        "esc_step_to": "1",
        "evaltype": "0",
        "opconditions": [],
        "optemplate": [
          {
            "operationid": "1",
            "templateid": "10001"
          }
        ]
      },
      "2": {
        "operationid": "2",
        "actionid": "2",
        "operationtype": "4",
        "esc_period": "0",
        "esc_step_from": "1",
        "esc_step_to": "1",
        "evaltype": "0",
        "opconditions": [],

```

```

        "opgroup": [
            {
                "operationid": "2",
                "groupid": "2"
            }
        ]
    }
},
    "id": 1
}

```

See also

- [Action condition](#)
- [Action operation](#)

Source

CAction::get() in *frontends/php/api/classes/CAction.php*.

action.update

Description

`object action.update(object/array actions)`

This method allows to update existing actions.

Parameters

(object/array) Action properties to be updated.

The `actionid` property must be defined for each action, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard action properties](#), the method accepts the following parameters.

Parameter	Type	Description
conditions	array	Action conditions to replace existing conditions.
operations	array	Action operations to replace existing operations.

Return values

(object) Returns an object containing the IDs of the updated actions under the `actionids` property.

Examples

Disable action

Disable action, that is, set its status to "1".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "action.update",
    "params": {
        "actionid": "2",
        "status": 1
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "actionids": [
      "2"
    ]
  },
  "id": 1
}
```

See also

- [Action condition](#)
- [Action operation](#)

Source

CAction::update() in *frontends/php/api/classes/CAction.php*.

Alert

This class is designed to work with alerts.

Object references:

- [Alert](#)

Available methods:

- [alert.get](#) - retrieve alerts

> Alert object

The following objects are directly related to the alert API.

Alert

Note:

Alerts are created by the Zabbix server and cannot be modified via the API.

The alert object contains information about whether certain action operations have been executed successfully. It has the following properties.

Property	Type	Description
alertid	string	ID of the alert.
actionid	string	ID of the action that generated the alert.
alerttype	integer	Alert type. Possible values: 0 - message; 1 - remote command.
clock	timestamp	Time when the alert was generated.
error	string	Error text if there are problems sending a message or running a command.
esc_step	integer	Action escalation step during which the alert was generated.
eventid	string	ID of the event that triggered the action.
mediatypeid	string	ID of the media type that was used to send the message.
message	string	Message text. Used for message alerts.
retries	integer	Number of times Zabbix tried to send the message.
sendto	string	Address, user name or other identifier of the recipient. Used for message alerts.

Property	Type	Description
status	integer	Status indicating whether the action operation has been executed successfully. Possible values for message alerts: 0 - message not sent; 1 - message sent; 2 - failed after a number of retries. Possible values for command alerts: 1 - command run; 2 - tried to run the command on the Zabbix agent but it was unavailable.
subject	string	Message subject. Used for message alerts.
userid	string	ID of the user that the message was sent to.

alert.get

Description

`integer/array alert.get(object parameters)`

The method allows to retrieve alerts according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
alertids	string/array	Return only alerts with the given IDs.
actionids	string/array	Return only alerts generated by the given actions.
eventids	string/array	Return only alerts generated by the given events.
groupids	string/array	Return only alerts generated by objects from the given host groups.
hostids	string/array	Return only alerts generated by objects from the given hosts.
mediatypeids	string/array	Return only message alerts that used the given media types.
triggerids	string/array	Return only alerts generated by the given triggers.
userids	string/array	Return only message alerts that were sent to the given users.
time_from	timestamp	Return only alerts that have been generated after the given time.
time_till	timestamp	Return only alerts that have been generated before the given time.
selectHosts	query	Return the hosts that triggered the action operation in the <code>hosts</code> property.
selectMediatypes	query	Does not support <code>shorten</code> or arrays of properties. Return the media type that was used for the message alert as an array in the <code>mediatypes</code> property.
selectUsers	query	Does not support <code>shorten</code> or arrays of properties. Return the user that the message was addressed to as an array in the <code>users</code> property.
sortfield	string/array	Does not support <code>shorten</code> or arrays of properties. Sort the result by the given properties. Possible values are: <code>alertid</code> , <code>clock</code> , <code>eventid</code> and <code>status</code> .

Parameter	Type	Description
countOutput	flag	These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve alerts by action ID

Retrieve all alerts generated by action "3".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "alert.get",
  "params": {
    "output": "extend",
    "actionids": "3"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "alertid": "1",
      "actionid": "3",
      "eventid": "21243",
      "userid": "1",
      "clock": "1362128008",
      "mediatypeid": "1",
      "sendto": "support@company.com",
      "subject": "PROBLEM: Zabbix agent on Linux server is unreachable for 5 minutes: ",
      "message": "Trigger: Zabbix agent on Linux server is unreachable for 5 minutes: \nTrigger stat",
      "status": "0",
      "retries": "3",
      "error": "",
      "esc_step": "1",
      "alerttype": "0"
    }
  ],
}
```

```
    "id": 1
}
```

See also

- [Host](#)
- [Media type](#)
- [User](#)

Source

CAAlert::get() in *frontends/php/api/classes/CAAlert.php*.

API info

This class is designed to retrieve meta information about the API.

Available methods:

- [apiinfo.version](#) - retrieving the version of the Zabbix API

apiinfo.version

Description

`string apiinfo.version(array)`

This method allows to retrieve the version of the Zabbix API.

Parameters

(array) The method accepts an empty array.

Return values

Attention:

This method is available to unauthenticated users and should be called without the `auth` parameter in the JSON-RPC request. Starting from Zabbix 2.4 the method will return an error if the `auth` parameter is given.

(string) Returns the version of the Zabbix API.

Note:

Starting from Zabbix 2.0.4 the version of the API matches the version of Zabbix. In earlier 2.0 releases the method will always return 1.4.

Examples

Retrieving the version of the API

Retrieve the version of the Zabbix API.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apiinfo.version",
    "params": [],
    "id": 1,
    "auth": "16a46baf181ef9602e1687f3110abf8a"
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "2.0.5",
    "id": 1
}
```

Source

CAPInfo::version() in *frontends/php/api/classes/CAPInfo.php*.

Application

This class is designed to work with applications.

Object references:

- **Application**

Available methods:

- **application.create** - creating new applications
- **application.delete** - deleting applications
- **application.exists** - checking if applications exist
- **application.get** - retrieving application
- **application.massadd** - updating application
- **application.update** - adding items to applications

> Application object

The following objects are directly related to the application API.

Application

The application object has the following properties.

Property	Type	Description
applicationid	string	(<i>readonly</i>) ID of the application.
hostid (required)	string	ID of the host that the application belongs to.
name (required)	string	Cannot be updated. Name of the application
templateid	string	(<i>readonly</i>) ID of the parent template application.

application.create

Description

`object application.create(object/array applications)`

This method allows to create new applications.

Parameters

(object/array) Applications to create.

The method accepts applications with the **standard application properties**.

Return values

(object) Returns an object containing the IDs of the created applications under the `applicationids` property. The order of the returned IDs matches the order of the passed applications.

Examples

Creating an application

Create an application to store SNMP items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "application.create",
  "params": {
    "name": "SNMP Items",
    "hostid": "10050"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "applicationids": [
      "356"
    ]
  },
  "id": 1
}
```

Source

CApplication::create() in *frontends/php/api/classes/CApplication.php*.

application.delete

Description

object application.delete(array applicationIds)

This method allows to delete applications.

Parameters

(array) IDs of the applications to delete.

Return values

(object) Returns an object containing the IDs of the deleted applications under the applicationids property.

Examples

Deleting multiple applications

Delete two applications.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "application.delete",
  "params": [
    "356",
    "358"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "applicationids": [
      "356",
      "358"
    ]
  }
}
```

```

    ]
  },
  "id": 1
}

```

Source

CApplication::delete() in *frontends/php/api/classes/CApplication.php*.

application.exists

Description

boolean application.exists(object filter)

This method checks if at least one application that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
hostid	string/array	IDs of the hosts the applications must belong to.
name	string/array	Names of the applications
node	string	Name of the node the applications must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the applications must belong to.

Return values

(boolean) Returns true if at least one application that matches the given filter criteria exists.

Examples

Check application on host

Check if application "Memory" exists on host "10084."

Request:

```

{
  "jsonrpc": "2.0",
  "method": "application.exists",
  "params": {
    "hostid": "10084",
    "name": "Memory"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}

```

Source

CApplication::exists() in *frontends/php/api/classes/CApplication.php*.

application.get

Description

integer/array application.get(object parameters)

The method allows to retrieve applications according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
applicationids	string/array	Return only applications with the given IDs.
groupids	string/array	Return only applications that belong to hosts from the given host groups.
hostids	string/array	Return only applications that belong to the given hosts.
inherited	boolean	If set to true return only applications inherited from a template.
itemids	string/array	Return only applications that contain the given items.
templated	boolean	If set to true return only applications that belong to templates.
templateids	string/array	Return only applications that belong to the given templates.
expandData	flag	Return the name of the host that the application belongs to as a property of the application object.
selectHosts	query	Return the hosts that the application belongs to in the hosts property.
selectItems	query	Does not support shorten. Return the items contained in the application in the items property.
sortfield	string/array	Does not support shorten or arrays of values. Sort the result by the given properties.
countOutput	flag	Possible values are: applicationid and name. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving applications from a host

Retrieve all applications from a host sorted by name.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "application.get",
  "params": {
    "output": "extend",
    "hostids": "10001",
    "sortfield": "name"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "applicationid": "13",
      "hostid": "10001",
      "name": "CPU",
      "templateid": "0"
    },
    {
      "applicationid": "5",
      "hostid": "10001",
      "name": "Filesystems",
      "templateid": "0"
    },
    {
      "applicationid": "21",
      "hostid": "10001",
      "name": "General",
      "templateid": "0"
    },
    {
      "applicationid": "15",
      "hostid": "10001",
      "name": "Memory",
      "templateid": "0"
    }
  ],
  "id": 1
}
```

See also

- [Host](#)
- [Item](#)

Source

CApplication::get() in *frontends/php/api/classes/CApplication.php*.

application.massadd

Description

object application.massadd(object parameters)

This method allows to simultaneously add multiple items to the given applications.

Parameters

(object) Parameters containing the IDs of the applications to update and the items to add to the applications.

The method accepts the following parameters.

Parameter	Type	Description
applications (required)	array/object	Applications to be updated. The applications must have the <code>applicationid</code> property defined.
items	array/object	Items to add to the given applications. The items must have the <code>itemid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated applications under the `applicationids` property.

Examples

Adding items to multiple applications

Add the given items to two applications.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "application.massadd",
  "params": {
    "applications": [
      {
        "applicationid": "247"
      },
      {
        "applicationid": "246"
      }
    ],
    "items": [
      {
        "itemid": "22800"
      },
      {
        "itemid": "22801"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "applicationids": [
      "247",
      "246"
    ]
  },
  "id": 1
}
```

See also

- [Item](#)

Source

CApplication::massAdd() in *frontends/php/api/classes/CApplication.php*.

application.update

Description

object application.update(object/array applications)

This method allows to update existing applications.

Parameters

(object/array) **Application properties** to be updated.

The applicationid property must be defined for each application, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated applications under the applicationids property.

Examples

Changing the name of an application

Change the name of the application to "Processes and performance".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "application.update",
  "params": {
    "applicationid": "13",
    "name": "Processes and performance"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "applicationids": [
      "13"
    ]
  },
  "id": 1
}
```

Source

CApplication::update() in *frontends/php/api/classes/CApplication.php*.

Configuration

This class is designed to export and import Zabbix configuration data.

Available methods:

- **configuration.export** - exporting the configuration
- **configuration.import** - importing the configuration

configuration.export

Description

string configuration.export(object parameters)

This method allows to export configuration data as a serialized string.

Parameters

(object) Parameters defining the objects to be exported and the format to use.

Parameter	Type	Description
format (required)	string	Format in which the data must be exported. Possible values: json - JSON; xml - XML.
options (required)	object	Objects to be exported. The options object has the following parameters: groups - (array) IDs of host groups to export; hosts - (array) IDs of hosts to export; images - (array) IDs of images to export; maps - (array) IDs of maps to export. screens - (array) IDs of screens to export; templates - (array) IDs of templates to export;

Return values

(string) Returns a serialized string containing the requested configuration data.

Examples

Exporting a host

Export the configuration of a host as an XML string.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "configuration.export",
  "params": {
    "options": {
      "hosts": [
        "10161"
      ]
    },
    "format": "xml"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": "<!--?xml version=\"1.0\" encoding=\"UTF-8\"?-->\n<zabbix_export><version>2.0</version><date>2016-05-27 12:00:00</date><groups>\n</groups><hosts>\n</hosts></zabbix_export>",
  "id": 1
}
```

Source

CConfiguration::export() in frontends/php/api/classes/CConfiguration.php.

configuration.import

Description

boolean configuration.import(object parameters)

This method allows to import configuration data from a serialized string.

Parameters

(object) Parameters containing the data to import and rules how the data should be handled.

Parameter	Type	Description
format (required)	string	Format of the serialized string. Possible values: json - JSON; xml - XML.
source (required)	string	Serialized string containing the configuration data.
rules (required)	object	Rules on how new and existing objects should be imported. The rules parameter is described in detail in the table below.

Note:

If no rules are given, the configuration will not be updated.

The rules object supports the following parameters.

Parameter	Type	Description
applications	object	Rules on how to import applications. Supported parameters: createMissing - (boolean) if set to true, new applications will be created; default: false; updateExisting - (boolean) if set to true, existing applications will be updated; default: false.
discoveryRules	object	Rules on how to import LLD rules. Supported parameters: createMissing - (boolean) if set to true, new LLD rules will be created; default: false; updateExisting - (boolean) if set to true, existing LLD rules will be updated; default: false.
graphs	object	Rules on how to import graphs. Supported parameters: createMissing - (boolean) if set to true, new graphs will be created; default: false; updateExisting - (boolean) if set to true, existing graphs will be updated; default: false.
groups	object	Rules on how to import host groups. Supported parameters: createMissing - (boolean) if set to true, new host groups will be created; default: false.

Parameter	Type	Description
hosts	object	Rules on how to import hosts. Supported parameters: createMissing - (boolean) if set to true, new hosts will be created; default: false; updateExisting - (boolean) if set to true, existing hosts will be updated; default: false.
images	object	Rules on how to import images. Supported parameters: createMissing - (boolean) if set to true, new images will be created; default: false; updateExisting - (boolean) if set to true, existing images will be updated; default: false.
items	object	Rules on how to import items. Supported parameters: createMissing - (boolean) if set to true, new items will be created; default: false; updateExisting - (boolean) if set to true, existing items will be updated; default: false.
maps	object	Rules on how to import maps. Supported parameters: createMissing - (boolean) if set to true, new maps will be created; default: false; updateExisting - (boolean) if set to true, existing maps will be updated; default: false.
screens	object	Rules on how to import screens. Supported parameters: createMissing - (boolean) if set to true, new screens will be created; default: false; updateExisting - (boolean) if set to true, existing screens will be updated; default: false.
templateLinkage	object	Rules on how to import template links. Supported parameters: createMissing - (boolean) if set to true, new links between templates and host will be created; default: false.
templates	object	Rules on how to import templates. Supported parameters: createMissing - (boolean) if set to true, new templates will be created; default: false; updateExisting - (boolean) if set to true, existing templates will be updated; default: false.
templateScreens	object	Rules on how to import template screens. Supported parameters: createMissing - (boolean) if set to true, new template screens will be created; default: false; updateExisting - (boolean) if set to true, existing template screens will be updated; default: false.

Parameter	Type	Description
triggers	object	Rules on how to import triggers. Supported parameters: createMissing - (boolean) if set to true, new triggers will be created; default: false; updateExisting - (boolean) if set to true, existing triggers will be updated; default: false.

Return values

(boolean) Returns true if importing has been successful.

Examples

Importing hosts and items

Import the host and items contained in the XML string. Leave everything else unchanged.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "configuration.import",
  "params": {
    "format": "xml",
    "rules": {
      "hosts": {
        "createMissing": true,
        "updateExisting": true
      },
      "items": {
        "createMissing": true,
        "updateExisting": true
      }
    },
    "source": "<!--?xml version='1.0' encoding='UTF-8'?--><zabbix_export><version>2.0</version><data>"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}

```

Source

CConfiguration::import() in *frontends/php/api/classes/CConfiguration.php*.

Discovered host

This class is designed to work with discovered hosts.

Object references:

- Discovered host

Available methods:

- dhost.exists - check if a discovered host exists
- dhost.get - retrieve discovered hosts

> Discovered host object

The following objects are directly related to the `dhost` API.

Discovered host

Note:

Discovered host are created by the Zabbix server and cannot be modified via the API.

The discovered host object contains information about a host discovered by a network discovery rule. It has the following properties.

Property	Type	Description
<code>dhostid</code>	string	ID of the discovered host.
<code>druleid</code>	string	ID of the discovery rule that detected the host.
<code>lastdown</code>	timestamp	Time when the discovered host last went down.
<code>lastup</code>	timestamp	Time when the discovered host last went up.
<code>status</code>	integer	Whether the discovered host is up or down. A host is up if it has at least one active discovered service.
		Possible values: 0 - host up; 1 - host down.

`dhost.exists`

Description

`boolean dhost.exists(object filter)`

This method checks if at least one discovered host that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
<code>dhostid</code>	string/array	IDs of the discovered hosts.
<code>node</code>	string	Name of the node the discovered hosts must belong to.
<code>nodeids</code>	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the discovered hosts must belong to.

Return values

(boolean) Returns `true` if at least one discovered host that matches the given filter criteria exists.

Examples

Check multiple discovered hosts

Check if discovered hosts with IDs "1" and "2" exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "dhost.exists",
  "params": {
    "dhostid": [
      "1",
      "2"
    ]
  }
}
```

```

    },
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}

```

Source

CDHost::exists() in *frontends/php/api/classes/CDHost.php*.

dhost.get

Description

integer/array dhost.get(object parameters)

The method allows to retrieve discovered hosts according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dhostids	string/array	Return only discovered hosts with the given IDs.
druleids	string/array	Return only discovered hosts that have been created by the given discovery rules.
dserviceids	string/array	Return only discovered hosts that are running the given services.
selectDRules	query	Return the discovery rule that detected the host as an array in the drules property.
selectDServices	query	Does not support shorten. Return the discovered services running on the host in the dservices property.
limitSelects	integer	Supports count; does not support shorten. Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: selectDServices - results will be sorted by dserviceid. Sort the result by the given properties.
countOutput	flag	Possible values are: dhostid and druleid. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	

Parameter	Type	Description
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve discovered hosts by discovery rule

Retrieve all hosts and the discovered services they are running that have been detected by discovery rule "4".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "dhost.get",
  "params": {
    "output": "extend",
    "selectDServices": "extend",
    "druleids": "4"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "dservices": [
        {
          "dserviceid": "1",
          "dhostid": "1",
          "type": "4",
          "key_": "",
          "value": "",
          "port": "80",
          "status": "0",
          "lastup": "1337697227",
          "lastdown": "0",
          "dcheckid": "5",
          "ip": "192.168.1.1",
          "dns": "station.company.lan"
        }
      ],
      "dhostid": "1",
      "druleid": "4",
      "status": "0",
      "lastup": "1337697227",
      "lastdown": "0"
    },
    {
      "dservices": [
        {
          "dserviceid": "2",
          "dhostid": "2",
          "type": "4",

```

```

        "key_": "",
        "value": "",
        "port": "80",
        "status": "0",
        "lastup": "1337697234",
        "lastdown": "0",
        "dcheckid": "5",
        "ip": "192.168.1.4",
        "dns": "john.company.lan"
    }
],
"dhostid": "2",
"druleid": "4",
"status": "0",
"lastup": "1337697234",
"lastdown": "0"
},
{
    "dservices": [
        {
            "dserviceid": "3",
            "dhostid": "3",
            "type": "4",
            "key_": "",
            "value": "",
            "port": "80",
            "status": "0",
            "lastup": "1337697234",
            "lastdown": "0",
            "dcheckid": "5",
            "ip": "192.168.1.26",
            "dns": "printer.company.lan"
        }
    ],
    "dhostid": "3",
    "druleid": "4",
    "status": "0",
    "lastup": "1337697234",
    "lastdown": "0"
},
{
    "dservices": [
        {
            "dserviceid": "4",
            "dhostid": "4",
            "type": "4",
            "key_": "",
            "value": "",
            "port": "80",
            "status": "0",
            "lastup": "1337697234",
            "lastdown": "0",
            "dcheckid": "5",
            "ip": "192.168.1.7",
            "dns": "mail.company.lan"
        }
    ],
    "dhostid": "4",
    "druleid": "4",
    "status": "0",
    "lastup": "1337697234",
    "lastdown": "0"
}

```

```

    }
  ],
  "id": 1
}

```

See also

- [Discovered service](#)
- [Discovery rule](#)

Source

CDHost::get() in *frontends/php/api/classes/CDHost.php*.

Discovered service

This class is designed to work with discovered services.

Object references:

- [Discovered service](#)

Available methods:

- [dservice.exists](#) - check if a discovered service exists
- [dservice.get](#) - retrieve discovered services

> Discovered service object

The following objects are directly related to the `dservice` API.

Discovered service

Note:

Discovered services are created by the Zabbix server and cannot be modified via the API.

The discovered service object contains information about a service discovered by a network discovery rule on a host. It has the following properties.

Property	Type	Description
<code>dserviceid</code>	string	ID of the discovered service.
<code>dcheckid</code>	string	ID of the discovery check used to detect the service.
<code>dhostid</code>	string	ID of the discovered host running the service.
<code>dns</code>	string	DNS of the host running the service.
<code>ip</code>	string	IP address of the host running the service.
<code>key_</code>	string	Key used by a Zabbix agent discovery check to locate the service.
<code>lastdown</code>	timestamp	Time when the discovered service last went down.
<code>lastup</code>	timestamp	Time when the discovered service last went up.
<code>port</code>	integer	Service port number.
<code>status</code>	integer	Status of the service.
		Possible values: 0 - service up; 1 - service down.
<code>type</code>	integer	Type of discovered service. The type of service matches the type of the discovery check used to detect the service.

Refer to the [discovery check "type" property](#) for a list of supported types.

Property	Type	Description
value	string	Value returned by the service when performing a Zabbix agent, SNMPv1, SNMPv2 or SNMPv3 discovery check.

dservice.exists

Description

`boolean dservice.exists(object filter)`

This method checks if at least one discovered service that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
dserviceid	string/array	IDs of discovered services.
node	string	Name of the node the discovered services must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the discovered services must belong to.

Return values

(boolean) Returns `true` if at least one discovered service that matches the given filter criteria exists.

Examples

Check multiple discovered services

Check if discovered services with IDs "121" and "73" exist.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "dservice.exists",
  "params": {
    "dserviceid": [
      "121",
      "73"
    ]
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

`CDService::exists()` in *frontends/php/api/classes/CDService.php*.

dservice.get

Description

integer/array dservice.get(object parameters)

The method allows to retrieve discovered services according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dserviceids	string/array	Return only discovered services with the given IDs.
dhostids	string/array	Return only discovered services that belong to the given discovered hosts.
dcheckids	string/array	Return only discovered services that have been detected by the given discovery checks.
druleids	string/array	Return only discovered services that have been detected by the given discovery rules.
selectDRules	query	Return the discovery rule that detected the service as an array in the drules property.
selectDHosts	query	Does not support shorten. Return the discovered host that service belongs to as an array in the dhosts property.
selectHosts	query	Does not support shorten. Return the hosts with the same IP address as the service in the hosts property.
limitSelects	integer	Supports count; does not support shorten. Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: selectHosts - result will be sorted by hostid. Sort the result by the given properties.
countOutput	flag	Possible values are: dserviceid, dhostid and ip. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve services discovered on a host

Retrieve all discovered services detected on discovered host "11".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "dservice.get",
  "params": {
    "output": "extend",
    "dhostids": "11"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "dserviceid": "12",
      "dhostid": "11",
      "type": "4",
      "key_": "",
      "value": "",
      "port": "80",
      "status": "1",
      "lastup": "0",
      "lastdown": "1348650607",
      "dcheckid": "5",
      "ip": "192.168.1.134",
      "dns": "john.local"
    },
    {
      "dserviceid": "13",
      "dhostid": "11",
      "type": "3",
      "key_": "",
      "value": "",
      "port": "21",
      "status": "1",
      "lastup": "0",
      "lastdown": "1348650610",
      "dcheckid": "6",
      "ip": "192.168.1.134",
      "dns": "john.local"
    }
  ],
  "id": 1
}
```

See also

- [Discovered host](#)
- [Discovery check](#)
- [Host](#)

Source

CDService::get() in *frontends/php/api/classes/CDService.php*.

Discovery check

This class is designed to work with discovery checks.

Object references:

- [Discovery check](#)

Available methods:

- [dcheck.get](#) - retrieve discovery checks

> Discovery check object

The following objects are directly related to the `dcheck` API.

Discovery check

The discovery check object defines a specific check performed by a network discovery rule. It has the following properties.

Property	Type	Description
dcheckid	string	(<i>readonly</i>) ID of the discovery check.
druleid	string	ID of the discovery rule that the check belongs to.
key_	string	The value of this property differs depending on the type type of the check: - key to query for Zabbix agent checks, required; - SNMP OID for SNMPv1, SNMPv2 and SNMPv3 checks, required.
ports	string	One or several port ranges to check separated by commas. Used for all checks except for ICMP.
snmp_community	string	Default: 0. SNMP community.
snmpv3_authpassphrase	string	Required for SNMPv1, SNMPv2 and SNMPv3 agent checks. Auth passphrase used for SNMPv3 agent checks with security level set to <i>authNoPriv</i> or <i>authPriv</i> .
snmpv3_privpassphrase	string	Priv passphrase used for SNMPv3 agent checks with security level set to <i>authPriv</i> .
snmpv3_securitylevel	string	Security level used for SNMPv3 agent checks. Possible values: 0 - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv.
snmpv3_securityname	string	Security name used for SNMPv3 agent checks.

type	integer	Type of check. Possible values: 0 - <i>(default)</i> SSH; 1 - LDAP; 2 - SMTP; 3 - FTP; 4 - HTTP; 5 - POP; 6 - NNTP; 7 - IMAP; 8 - TCP; 9 - Zabbix agent; 10 - SNMPv1 agent; 11 - SNMPv2 agent; 12 - ICMP ping; 13 - SNMPv3 agent; 14 - HTTPS; 15 - Telnet.
uniq	integer	Whether to use this check as a device uniqueness criteria. Only a single unique check can be configured for a discovery rule. Used for Zabbix agent, SNMPv1, SNMPv2 and SNMPv3 agent checks. Possible values: 0 - <i>(default)</i> do not use this check as a uniqueness criteria; 1 - use this check as a uniqueness criteria.

dcheck.get

integer/array dcheck.get(object parameters)

The method allows to retrieve discovery checks according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dcheckids	string/array	Return only discovery checks with the given IDs.
druleids	string/array	Return only discovery checks that belong to the given discovery rules.
dserviceids	string/array	Return only discovery checks that have detected the given discovered services.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: dcheckid and druleid. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	

Parameter	Type	Description
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve discovery checks for a discovery rule

Retrieve all discovery checks used by discovery rule "6".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "dcheck.get",
  "params": {
    "output": "extend",
    "dcheckids": "6"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "dcheckid": "6",
      "druleid": "4",
      "type": "3",
      "key_": "",
      "snmp_community": "",
      "ports": "21",
      "snmpv3_securityname": "",
      "snmpv3_securitylevel": "0",
      "snmpv3_authpassphrase": "",
      "snmpv3_privpassphrase": "",
      "uniq": "0"
    }
  ],
  "id": 1
}
```

Source

`CDCheck::get()` in *frontends/php/api/classes/CDCheck.php*.

Discovery rule

This class is designed to work with network discovery rules.

Note:

This API is meant to work with network discovery rules. For the low-level discovery rules see the [LLD rule API](#).

Object references:

- [Discovery rule](#)

Available methods:

- [drule.create](#) - create new discovery rules
- [drule.delete](#) - delete discovery rules
- [drule.exists](#) - check if a discovery rule exists
- [drule.get](#) - retrieve discovery rules
- [drule.isreadable](#) - check if discovery rules are readable
- [drule.iswritable](#) - check if discovery rules are writable
- [drule.update](#) - update discovery rules

> Discovery rule object

The following objects are directly related to the `drule` API.

Discovery rule

The discovery rule object defines a network discovery rule. It has the following properties.

Property	Type	Description
<code>druleid</code>	string	(<i>readonly</i>) ID of the discovery rule.
<code>iprange</code> (required)	string	One or several IP ranges to check separated by commas. Refer to the network discovery configuration section for more information on supported formats of IP ranges.
<code>name</code> (required)	string	Name of the discovery rule.
<code>delay</code>	integer	Execution interval of the discovery rule in seconds. Default: 3600.
<code>nextcheck</code>	timestamp	(<i>readonly</i>) Time when the discovery rule will be executed next.
<code>proxy_hostid</code>	string	ID of the proxy used for discovery.
<code>status</code>	integer	Whether the discovery rule is enabled. Possible values: 0 - (<i>default</i>) enabled; 1 - disabled.

`drule.create`

Description

object `drule.create(object/array $discoveryRules)`

This method allows to create new discovery rules.

Parameters

(object/array) Discovery rules to create.

Additionally to the [standard discovery rule properties](#), the method accepts the following parameters.

Parameter	Type	Description
<code>dchecks</code> (required)	array	Discovery checks to create for the discovery rule.

Return values

(object) Returns an object containing the IDs of the created discovery rules under the `druleids` property. The order of the returned IDs matches the order of the passed discovery rules.

Examples

Create a discovery rule

Create a discovery rule to find machines running the Zabbix agent in the local network. The rule must use a single Zabbix agent check on port 10050.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.create",
  "params": {
    "name": "Zabbix agent discovery",
    "iprange": "192.168.1.1-255",
    "dchecks": [
      {
        "type": "9",
        "key_": "system.uname",
        "ports": "10050",
        "uniq": "0"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "druleids": [
      "6"
    ]
  },
  "id": 1
}
```

See also

- [Discovery check](#)

Source

`CDRule::create()` in *frontends/php/api/classes/CDRule.php*.

drule.delete

Description

`object drule.delete(array discoveryRuleIds)`

This method allows to delete discovery rules.

Parameters

(array) IDs of the discovery rules to delete.

Return values

(object) Returns an object containing the IDs of the deleted discovery rules under the `druleids` property.

Examples

Delete multiple discovery rules

Delete two discovery rules.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.delete",
  "params": [
    "4",
    "6"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "druleids": [
      "4",
      "6"
    ]
  },
  "id": 1
}
```

Source

CDRule::delete() in *frontends/php/api/classes/CDRule.php*.

drule.exists

Description

`boolean drule.exists(object filter)`

This method checks if at least one discovery rule that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
druleids	string/array	IDs of discovery rules.
name	string/array	Names of discovery rules.
node	string	Name of the node the discovery rules must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the discovery rules must belong to.

Return values

(boolean) Returns true if at least one discovery rule that matches the given filter criteria exists.

Examples

Check a discovery rule by name

Check if a discovery rule called "Local network" exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.exists",
  "params": {
    "name": "Local network"
  }
}
```

```

    },
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}

```

See also

- [drule.isreadable](#)
- [drule.iswritable](#)

Source

CDRule::exists() in *frontends/php/api/classes/CDRule.php*.

drule.get

Description

integer/array drule.get(object parameters)

The method allows to retrieve discovery rules according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dhostids	string/array	Return only discovery rules that created the given discovered hosts.
druleids	string/array	Return only discovery rules with the given IDs.
dserviceids	string/array	Return only discovery rules that created the given discovered services.
selectDChecks	query	Return discovery checks used by the discovery rule in the dchecks property.
selectDHosts	query	Supports count; does not support shorten. Return the discovered hosts that the discovery rule created in the dhosts property.
limitSelects	integer	Supports count; does not support shorten. Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: selectDChecks - results will be sorted by dcheckid; selectDHosts - results will be sorted by dhostsId. Sort the result by the given properties.
countOutput	flag	Possible values are: druleid and name. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	

Parameter	Type	Description
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve all discovery rules

Retrieve all configured discovery rules and the discovery checks they use.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.get",
  "params": {
    "output": "extend",
    "selectDChecks": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "druleid": "2",
      "proxy_hostid": "0",
      "name": "Local network",
      "iprange": "192.168.3.1-255",
      "delay": "5",
      "nextcheck": "1348754327",
      "status": "0",
      "dchecks": [
        {
          "dcheckid": "7",
          "druleid": "2",
          "type": "3",
          "key_": "",
          "snmp_community": "",
          "ports": "21",
          "snmpv3_securityname": "",
          "snmpv3_securitylevel": "0",
          "snmpv3_authpassphrase": "",
          "snmpv3_privpassphrase": "",
          "uniq": "0"
        },
        {
          "dcheckid": "8",
          "druleid": "2",

```

```

        "type": "4",
        "key_": "",
        "snmp_community": "",
        "ports": "80",
        "snmpv3_securityname": "",
        "snmpv3_securitylevel": "0",
        "snmpv3_authpassphrase": "",
        "snmpv3_privpassphrase": "",
        "uniq": "0"
    }
],
},
{
    "druleid": "6",
    "proxy_hostid": "0",
    "name": "Zabbix agent discovery",
    "iprange": "192.168.1.1-255",
    "delay": "3600",
    "nextcheck": "0",
    "status": "0",
    "dchecks": [
        {
            "dcheckid": "10",
            "druleid": "6",
            "type": "9",
            "key_": "system.uptime",
            "snmp_community": "",
            "ports": "10050",
            "snmpv3_securityname": "",
            "snmpv3_securitylevel": "0",
            "snmpv3_authpassphrase": "",
            "snmpv3_privpassphrase": "",
            "uniq": "0"
        }
    ]
}
],
    "id": 1
}

```

See also

- [Discovered host](#)
- [Discovery check](#)

Source

CDRule::get() in *frontends/php/api/classes/CDRule.php*.

drule.isreadable

Description

boolean drule.isreadable(array discoveryRuleIds)

This method checks if the given discovery rules are available for reading.

Parameters

(array) IDs of the discovery rules to check.

Return values

(boolean) Returns true if the given discovery rules are available for reading.

Examples

Check multiple discovery rules

Check if the two discovery rules are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.isreadable",
  "params": [
    "5",
    "8"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [drule.exists](#)
- [drule.iswritable](#)

Source

CDRule::isReadable() in *frontends/php/api/classes/CDRule.php*.

drule.iswritable

Description

boolean drule.iswritable(array discoveryRuleIds)

This method checks if the given discovery rules are available for writing.

Parameters

(array) IDs of the discovery rules to check.

Return values

(boolean) Returns true if the given discovery rules are available for writing.

Examples

Check multiple discovery rules

Check if the two discovery rules are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.iswritable",
  "params": [
    "5",
    "8"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
}
```



```
    "id": 1
}
```

See also

- [drule.isreadable](#)
- [drule.exists](#)

Source

CDRule::isWritable() in *frontends/php/api/classes/CDRule.php*.

drule.update

Description

`object drule.update(object/array discoveryRules)`

This method allows to update existing discovery rules.

Parameters

(object/array) Discovery rule properties to be updated.

The `druleid` property must be defined for each discovery rule, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard discovery rule properties](#), the method accepts the following parameters.

Parameter	Type	Description
dchecks	array	Discovery checks to replace existing checks.

Return values

(object) Returns an object containing the IDs of the updated discovery rules under the `druleids` property.

Examples

Change the IP range of a discovery rule

Change the IP range of a discovery rule to "192.168.2.1-255".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "drule.update",
  "params": {
    "druleid": "6",
    "iprange": "192.168.2.1-255"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "druleids": [
      "6"
    ]
  },
  "id": 1
}
```

See also

- [Discovery check](#)

Source

CDRule::update() in *frontends/php/api/classes/CDRule.php*.

Event

This class is designed to work with events.

Object references:

- **Event**

Available methods:

- **event.get** - retrieving events
- **event.acknowledge** - acknowledging events

> **Event object**

The following objects are directly related to the event API.

Event

Note:
Events are created by the Zabbix server and cannot be modified via the API.

The event object has the following properties.

Property	Type	Description
eventid	string	ID of the event.
acknowledged	integer	Whether the event has been acknowledged.
clock	timestamp	Time when the event was created.
ns	integer	Nanoseconds when the event was created.
object	integer	Type of object that is related to the event. Possible values for trigger events: 0 - trigger. Possible values for discovery events: 1 - discovered host; 2 - discovered service. Possible values for auto-registration events: 3 - auto-registered host.
objectid	string	ID of the related object.
source	integer	Type of the event. Possible values: 0 - event created by a trigger; 1 - event created by a discovery rule; 2 - event created by active agent auto-registration.

Property	Type	Description
value	integer	State of the related object. Possible values for trigger events: 0 - OK; 1 - problem; 2 - unknown. Possible values for discovery events: 0 - host or service up; 1 - host or service down; 2 - host or service discovered; 3 - host or service lost.
value_changed	integer	This parameter is not used for active agent auto-registration events. Whether the state of the related object has changed since the previous event.

event.acknowledge

Description

object event.acknowledge(string/array parameters)

This method allows to acknowledge events and add an acknowledgement message. If an event is already acknowledged, a new message will still be added.

Parameters

(string/array) Parameters containing the IDs of the events acknowledge and a message.

Parameter	Type	Description
eventids (required)	string/object	IDs of the events to acknowledge.
message	string	Text of the acknowledgement message.

Note:

Only trigger events with value_changed set to 1 can be acknowledged.

Return values

(object) Returns an object containing the IDs of the acknowledged events under the eventids property.

Examples

Acknowledging an event

Acknowledge a single event and leave a message.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "event.acknowledge",
  "params": {
    "eventids": "20427",
    "message": "Problem resolved."
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "eventids": [
      "20427"
    ]
  },
  "id": 1
}
```

Source

CEvent::acknowledge() in *frontends/php/api/classes/CEvent.php*.

event.get

Description

integer/array event.get(object parameters)

The method allows to retrieve events according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
eventids	string/array	Return only events with the given IDs.
groupids	string/array	Return only events created by triggers that belong to the given host groups.
hostids	string/array	Return only events created by triggers that belong to the given hosts.
triggerids	string/array	Return only events that have been created by the given triggers.
object	integer	Return only events created by objects of the given type.
source	integer	Refer to the event object page for a list of supported object types. Return only events with the given type.
acknowledged	boolean	Refer to the event object page for a list of supported event types. If set to <code>true</code> return only acknowledged events.
value	integer/array	Return only events with the given values.
time_from	timestamp	Return only events that have been created after or at the given time.
time_till	timestamp	Return only events that have been created before or at the given time.
eventid_from	string	Return only events with IDs greater or equal to the given ID.
eventid_till	string	Return only events with IDs less or equal to the given ID.
selectHosts	query	Return hosts containing the trigger that created the event in the <code>hosts</code> property.
selectItems	query	Does not support <code>shorten</code> or arrays of properties. Return items contained in the trigger that created the event in the <code>items</code> property.

Does not support `shorten` or arrays of properties.

Parameter	Type	Description
selectTriggers	query	Return the trigger that created the event as an array in the <code>triggers</code> property.
select_alerts	query	Does not support <code>shorten</code> or arrays of properties. Return alerts generated by the event in the <code>alerts</code> property. Alerts are sorted in reverse chronological order.
select_acknowledges	query	Does not support <code>shorten</code> or arrays of properties. Return event's acknowledges in the <code>acknowledges</code> property. Acknowledges are sorted in reverse chronological order. The event acknowledgement object has the following properties: <code>acknowledgeid</code> - (string) acknowledgement's ID; <code>userid</code> - (string) ID of the user that acknowledged the event; <code>eventid</code> - (string) ID of the acknowledged event; <code>clock</code> - (timestamp) time when the event was acknowledged; <code>message</code> - (string) text of the acknowledgement message; <code>alias</code> - (string) alias of the user that acknowledged the event.
sortfield	string/array	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>eventid</code> , <code>object</code> and <code>objectid</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving trigger events

Retrieve the latest events from trigger "13926."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "event.get",
  "params": {
    "output": "extend",
    "select_acknowledges": "extend",
    "triggerids": "13926",
    "sortfield": "eventid",
    "sortorder": "DESC"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "acknowledges": [
        {
          "acknowledgeid": "1",
          "userid": "1",
          "eventid": "9695",
          "clock": "1350640590",
          "message": "Problem resolved.\n\r----[BULK ACKNOWLEDGE]----",
          "alias": "Admin"
        }
      ],
      "eventid": "9695",
      "source": "0",
      "object": "0",
      "objectid": "13926",
      "clock": "1347970410",
      "value": "1",
      "acknowledged": "1",
      "ns": "413316245",
      "value_changed": "1"
    },
    {
      "acknowledges": [],
      "eventid": "9671",
      "source": "0",
      "object": "0",
      "objectid": "13926",
      "clock": "1347970347",
      "value": "2",
      "acknowledged": "0",
      "ns": "0",
      "value_changed": "0"
    }
  ],
  "id": 1
}
```

Retrieving events by time period

Retrieve all events that have been created between October 9 and 10, 2012, in reverse chronological order.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "event.get",
```

```

    "params": {
        "output": "extend",
        "time_from": "1349797228",
        "time_till": "1350661228",
        "sortfield": "eventid",
        "sortorder": "desc"
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": [
        {
            "eventid": "20616",
            "source": "0",
            "object": "0",
            "objectid": "14282",
            "clock": "1350477814",
            "value": "2",
            "acknowledged": "0",
            "ns": "0",
            "value_changed": "0"
        },
        {
            "eventid": "20617",
            "source": "0",
            "object": "0",
            "objectid": "14283",
            "clock": "1350477814",
            "value": "2",
            "acknowledged": "0",
            "ns": "0",
            "value_changed": "0"
        },
        {
            "eventid": "20618",
            "source": "0",
            "object": "0",
            "objectid": "14284",
            "clock": "1350477815",
            "value": "2",
            "acknowledged": "0",
            "ns": "0",
            "value_changed": "0"
        }
    ],
    "id": 1
}

```

See also

- [Alert](#)
- [Item](#)
- [Host](#)
- [Trigger](#)

Source

CEvent::get() in *frontends/php/api/classes/CEvent.php*.

Graph

This class is designed to work with items.

Object references:

- [Graph](#)

Available methods:

- [graph.create](#) - creating new graphs
- [graph.delete](#) - deleting graphs
- [graph.exists](#) - checking if graphs exists
- [graph.get](#) - retrieving graphs
- [graph.getobjects](#) - retrieving graphs by filters
- [graph.update](#) - updating graphs

> Graph object

The following objects are directly related to the `graph` API.

Graph

The graph object has the following properties.

Property	Type	Description
<code>graphid</code>	string	<i>(readonly)</i> ID of the graph.
<code>height</code> (required)	integer	Height of the graph in pixels.
<code>name</code> (required)	string	Name of the graph
<code>width</code> (required)	integer	Width of the graph in pixels.
<code>flags</code>	integer	<i>(readonly)</i> Origin of the graph. Possible values are: 0 - <i>(default)</i> a plain graph; 4 - a discovered graph.
<code>graphtype</code>	integer	Graph's layout type. Possible values: 0 - <i>(default)</i> normal; 1 - stacked; 2 - pie; 3 - exploded.
<code>percent_left</code>	float	Left percentile.
<code>percent_right</code>	float	Default: 0. Right percentile.
<code>show_3d</code>	integer	Default: 0. Whether to show pie and exploded graphs in 3D.
<code>show_legend</code>	integer	Possible values: 0 - <i>(default)</i> show in 2D; 1 - show in 3D. Whether to show the legend on the graph. Possible values: 0 - hide; 1 - <i>(default)</i> show.

Property	Type	Description
show_work_period	integer	Whether to show the working time on the graph. Possible values: 0 - hide; 1 - <i>(default)</i> show.
templateid	string	<i>(readonly)</i> ID of the parent template graph.
yaxismax	float	The fixed maximum value for the Y axis.
yaxismin	float	The fixed minimum value for the Y axis.
ymax_itemid	string	ID of the item that is used as the maximum value for the Y axis.
ymax_type	integer	Maximum value calculation method for the Y axis. Possible values: 0 - <i>(default)</i> calculated; 1 - fixed; 2 - item.
ymin_itemid	string	ID of the item that is used as the minimum value for the Y axis.
ymin_type	integer	Minimum value calculation method for the Y axis. Possible values: 0 - <i>(default)</i> calculated; 1 - fixed; 2 - item.

graph.create

Description

object graph.create(object/array graphs)

This method allows to create new graphs.

Parameters

(object/array) Graphs to create.

Additionally to the **standard graph properties**, the method accepts the following parameters.

Parameter	Type	Description
gitems (required)	array	Graph items to be created for the graph.

Return values

(object) Returns an object containing the IDs of the created graphs under the **graphids** property. The order of the returned IDs matches the order of the passed graphs.

Examples

Creating a graph

Create a graph with two items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graph.create",
  "params": {
    "name": "MySQL bandwidth",
    "width": 900,
    "height": 200,
    "gitems": [
```

```

        {
            "itemid": "22828",
            "color": "00AA00"
        },
        {
            "itemid": "22829",
            "color": "3333FF"
        }
    ]
},
"auth": "038e1d7b1735c6a5436ee9eae095879e",
"id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [
            "652"
        ]
    },
    "id": 1
}

```

See also

- [Graph item](#)

Source

CGraph::create() in *frontends/php/api/classes/CGraph.php*.

graph.delete

Description

object graph.delete(array graphIds)

This method allows to delete graphs.

Parameters

(array) IDs of the graphs to delete.

Return values

(object) Returns an object containing the IDs of the deleted graphs under the graphids property.

Examples

Deleting multiple graphs

Delete two graphs.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "graph.delete",
    "params": [
        "652",
        "653"
    ],
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "graphids": [
      "652",
      "653"
    ]
  },
  "id": 1
}
```

Source

CGraph::delete() in *frontends/php/api/classes/CGraph.php*.

graph.exists

Description

boolean graph.exists(object filter)

This method checks if at least one graph that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
host	string/array	Technical names of the hosts that the graphs belong to.
hostids	string/array	IDs of the hosts that the graphs belong to.
name	string/array	Names of the graphs.
node	string	Name of the node the graphs must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the graphs must belong to.

Return values

(boolean) Returns true if at least one graph that matches the given filter criteria exists.

Examples

Checking graph by name

Check if a graph named "CPU utilization" already exists on host "Zabbix server".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graph.exists",
  "params": {
    "name": "CPU utilization",
    "host": "Zabbix server"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
}
```

```

    "id": 1
}

```

Source

CGraph::exists() in *frontends/php/api/classes/CGraph.php*.

graph.get

Description

integer/array graph.get(object parameters)

The method allows to retrieve graphs according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
graphids	string/array	Return only graphs with the given IDs.
groupids	string/array	Return only graphs that belong to hosts in the given host groups.
templateids	string/array	Return only graph that belong to the given templates.
hostids	string/array	Return only graphs that belong to the given hosts.
itemids	string/array	Return only graphs that contain the given items.
templated	boolean	If set to <code>true</code> return only graphs that belong to templates.
inherited	boolean	If set to <code>true</code> return only graphs inherited from a template.
selectGroups	query	Return the host groups that the graph belongs to in the <code>groups</code> property.
selectTemplates	query	Does not support <code>shorten</code> or arrays of properties. Return the templates that the graph belongs to in the <code>templates</code> property.
selectHosts	query	Does not support <code>shorten</code> or arrays of properties. Return the hosts that the graph belongs to in the <code>hosts</code> property.
selectItems	query	Does not support <code>shorten</code> or arrays of properties. Return the items used in the graph in the <code>items</code> property.
selectGraphItems	query	Does not support <code>shorten</code> or arrays of properties. Return an object containing the graph items used in the graph in the <code>gitems</code> property.
selectDiscoveryRule	query	Does not support <code>shorten</code> or arrays of properties. Return the low-level discovery rule that created the graph in the <code>discoveryRule</code> property.
filter	object	Return only those results that exactly match the given filter.

Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against.

Supports additional filters:

`host` - technical name of the host that the graph belongs to;

`hostid` - ID of the host that the graph belongs to.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: graphid, name and graphtype. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving graphs from hosts

Retrieve all graphs from host "10107" and sort them by name.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graph.get",
  "params": {
    "output": "extend",
    "hostids": 10107,
    "sortfield": "name"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "graphid": "612",
      "name": "CPU jumps",
      "width": "900",
      "height": "200",
      "yaxismin": "0.0000",
      "yaxismax": "100.0000",
      "templateid": "439",
      "show_work_period": "1",
      "show_triggers": "1",
      "graphtype": "0",
      "show_legend": "1",
      "show_3d": "0",
      "percent_left": "0.0000",
      "percent_right": "0.0000",

```

```

    "ymin_type": "0",
    "ymax_type": "0",
    "ymin_itemid": "0",
    "ymax_itemid": "0",
    "flags": "0"
},
{
    "graphid": "613",
    "name": "CPU load",
    "width": "900",
    "height": "200",
    "yaxismin": "0.0000",
    "yaxismax": "100.0000",
    "templateid": "433",
    "show_work_period": "1",
    "show_triggers": "1",
    "graphtype": "0",
    "show_legend": "1",
    "show_3d": "0",
    "percent_left": "0.0000",
    "percent_right": "0.0000",
    "ymin_type": "1",
    "ymax_type": "0",
    "ymin_itemid": "0",
    "ymax_itemid": "0",
    "flags": "0"
},
{
    "graphid": "614",
    "name": "CPU utilization",
    "width": "900",
    "height": "200",
    "yaxismin": "0.0000",
    "yaxismax": "100.0000",
    "templateid": "387",
    "show_work_period": "1",
    "show_triggers": "0",
    "graphtype": "1",
    "show_legend": "1",
    "show_3d": "0",
    "percent_left": "0.0000",
    "percent_right": "0.0000",
    "ymin_type": "1",
    "ymax_type": "1",
    "ymin_itemid": "0",
    "ymax_itemid": "0",
    "flags": "0"
},
{
    "graphid": "645",
    "name": "Disk space usage /",
    "width": "600",
    "height": "340",
    "yaxismin": "0.0000",
    "yaxismax": "0.0000",
    "templateid": "0",
    "show_work_period": "0",
    "show_triggers": "0",
    "graphtype": "2",
    "show_legend": "1",
    "show_3d": "1",
    "percent_left": "0.0000",

```

```

        "percent_right": "0.0000",
        "ymin_type": "0",
        "ymax_type": "0",
        "ymin_itemid": "0",
        "ymax_itemid": "0",
        "flags": "4"
    }
],
    "id": 1
}

```

See also

- [graph.getobjects](#)
- [Discovery rule](#)
- [Graph item](#)
- [Item](#)
- [Host](#)
- [Host group](#)
- [Template](#)

Source

CGraph::get() in *frontends/php/api/classes/CGraph.php*.

graph.getobjects

Description

array graph.getobjects(object filter)

This method allows to retrieve graphs that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard graph properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the graphs must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the graphs must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving graphs from a host

Retrieve all graphs from host "Zabbix server".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "graph.getobjects",
    "params": {
        "host": "Zabbix server"
    },
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": [
    {
      "graphid": "612",
      "name": "CPU jumps",
      "width": "900",
      "height": "200",
      "yaxismin": "0.0000",
      "yaxismax": "100.0000",
      "templateid": "439",
      "show_work_period": "1",
      "show_triggers": "1",
      "graphtype": "0",
      "show_legend": "1",
      "show_3d": "0",
      "percent_left": "0.0000",
      "percent_right": "0.0000",
      "ymin_type": "0",
      "ymax_type": "0",
      "ymin_itemid": "0",
      "ymax_itemid": "0",
      "flags": "0"
    },
    {
      "graphid": "613",
      "name": "CPU load",
      "width": "900",
      "height": "200",
      "yaxismin": "0.0000",
      "yaxismax": "100.0000",
      "templateid": "433",
      "show_work_period": "1",
      "show_triggers": "1",
      "graphtype": "0",
      "show_legend": "1",
      "show_3d": "0",
      "percent_left": "0.0000",
      "percent_right": "0.0000",
      "ymin_type": "1",
      "ymax_type": "0",
      "ymin_itemid": "0",
      "ymax_itemid": "0",
      "flags": "0"
    },
    {
      "graphid": "614",
      "name": "CPU utilization",
      "width": "900",
      "height": "200",
      "yaxismin": "0.0000",
      "yaxismax": "100.0000",
      "templateid": "387",
      "show_work_period": "1",
      "show_triggers": "0",
      "graphtype": "1",
      "show_legend": "1",
      "show_3d": "0",
      "percent_left": "0.0000",
      "percent_right": "0.0000",
      "ymin_type": "1",
      "ymax_type": "1",
    }
  ]
}

```



```

        "ymin_itemid": "0",
        "ymax_itemid": "0",
        "flags": "0"
    }
],
    "id": 1
}

```

See also

- [graph.get](#)

Source

CGraph::getObject() in *frontends/php/api/classes/CGraph.php*.

graph.update

Description

object graph.update(object/array graphs)

This method allows to update existing graphs.

Parameters

(object/array) Graph properties to be updated.

The graphid property must be defined for each graph, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard graph properties](#) the method accepts the following parameters.

Parameter	Type	Description
gitems	array	Graph items to replace existing graph items. If a graph item has the gitemid property defined it will be updated, otherwise a new graph item will be created.

Return values

(object) Returns an object containing the IDs of the updated graphs under the graphids property.

Examples

Setting the maximum for the Y scale

Set the the maximum of the Y scale to a fixed value of 100.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "graph.update",
    "params": {
        "graphid": "439",
        "ymax_type": 1,
        "yaxismax": 100
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [

```

```

    ],
    "id": 1
}

```

Source

CGraph::update() in *frontends/php/api/classes/CGraph.php*.

Graph item

This class is designed to work with hosts.

Object references:

- [Graph item](#)

Available methods:

- [graphitem.get](#) - retrieving graph items

> Graph item object

The following objects are directly related to the `graphitem` API.

Graph item

Note:

Graph items can only be modified via the `graph` and `graphprototype` APIs.

The graph item object has the following properties.

Property	Type	Description
gitemid	string	(<i>readonly</i>) ID of the graph item.
color (required)	string	Graph item's draw color as a hexadecimal color code.
itemid (required)	string	ID of the item.
calc_fnc	integer	Value of the item that will be displayed. Possible values: 1 - minimum value; 2 - (<i>default</i>) average value; 4 - maximum value; 7 - all values; 9 - last value, used only by pie and exploded graphs.
drawtype	integer	Draw style of the graph item. Possible values: 0 - (<i>default</i>) line; 1 - filled region; 2 - bold line; 3 - dot; 4 - dashed line; 5 - gradient line.
graphid	string	ID of the graph that the graph item belongs to.
sortorder	integer	Position of the item in the graph. Default: 0.

Property	Type	Description
type	integer	Type of graph item. Possible values: 0 - <i>(default)</i> simple; 2 - graph sum, used only by pie and exploded graphs.
yaxisside	integer	Side of the graph where the graph item's Y scale will be drawn. Possible values: 0 - left side; 1 - <i>(default)</i> right side.

graphitem.get

Description

`integer/array graphitem.get(object parameters)`

The method allows to retrieve graph items according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
gitemids	string/array	Return only graph items with the given IDs.
graphids	string/array	Return only graph items that belong to the given graphs.
itemids	string/array	Return only graph items with the given item IDs.
type	integer	Return only graph items with the given type.
expandData	flag	Refer to the graph item object page for a list of supported graph item types. Return additional data about the item and the host.
selectGraphs	query	Adds the following properties to each graph item: <code>key_</code> - <i>(string)</i> key of the item; <code>hostid</code> - <i>(string)</i> ID of the host; <code>flags</code> - <i>(string)</i> origin of the item; <code>host</code> - <i>(string)</i> technical name of the host. Return the graph that the item belongs to as an array in the <code>graphs</code> property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>gitemid</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page .
editable	boolean	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
sortorder	string/array	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving graph items from a graph

Retrieve all graph items used in a graph with additional information about the item and the host.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graphitem.get",
  "params": {
    "output": "extend",
    "expandData": 1,
    "graphids": "387"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "gitemid": "1242",
      "graphid": "387",
      "itemid": "22665",
      "drawtype": "1",
      "sortorder": "1",
      "color": "FF5555",
      "yaxisside": "0",
      "calc_fnc": "2",
      "type": "0",
      "key_": "system.cpu.util[,steal]",
      "hostid": "10001",
      "flags": "0",
      "host": "Template OS Linux"
    },
    {
      "gitemid": "1243",
      "graphid": "387",
      "itemid": "22668",
      "drawtype": "1",
      "sortorder": "2",
      "color": "55FF55",
      "yaxisside": "0",
      "calc_fnc": "2",
      "type": "0",
      "key_": "system.cpu.util[,softirq]",
      "hostid": "10001",
      "flags": "0",
      "host": "Template OS Linux"
    },
    {
      "gitemid": "1244",
      "graphid": "387",
      "itemid": "22671",
      "drawtype": "1",
      "sortorder": "3",
      "color": "009999",
      "yaxisside": "0",

```

```

        "calc_fnc": "2",
        "type": "0",
        "key_": "system.cpu.util[,interrupt]",
        "hostid": "10001",
        "flags": "0",
        "host": "Template OS Linux"
    }
],
    "id": 1
}

```

See also

- [Graph](#)

Source

CGraphItem::get() in *frontends/php/api/classes/CGraphItem.php*.

Graph prototype

This class is designed to work with graph prototypes.

Object references:

- [Graph prototype](#)

Available methods:

- [graphprototype.create](#) - creating new graph prototypes
- [graphprototype.delete](#) - deleting graph prototypes
- [graphprototype.exists](#) - checking if graph prototypes exist
- [graphprototype.get](#) - retrieving graph prototypes
- [graphprototype.getobjects](#) - retrieving graph objects by filters
- [graphprototype.update](#) - updating graph prototypes

> Graph prototype object

The following objects are directly related to the graphprototype API.

Graph prototype

The graph prototype object has the following properties.

Property	Type	Description
graphid	string	(<i>readonly</i>) ID of the graph prototype.
height (required)	integer	Height of the graph prototype in pixels.
name (required)	string	Name of the graph prototype.
width (required)	integer	Width of the graph prototype in pixels.
graphtype	integer	Graph prototypes's layout type. Possible values: 0 - (<i>default</i>) normal; 1 - stacked; 2 - pie; 3 - exploded.
percent_left	float	Left percentile. Default: 0.

Property	Type	Description
percent_right	float	Right percentile.
show_3d	integer	Default: 0. Whether to show discovered pie and exploded graphs in 3D.
show_legend	integer	Possible values: 0 - (default) show in 2D; 1 - show in 3D. Whether to show the legend on the discovered graph.
show_work_period	integer	Possible values: 0 - hide; 1 - (default) show. Whether to show the working time on the discovered graph.
templateid	string	Possible values: 0 - hide; 1 - (default) show. (readonly) ID of the parent template graph prototype.
yaxismax	float	The fixed maximum value for the Y axis.
yaxismin	float	The fixed minimum value for the Y axis.
ymax_itemid	string	ID of the item that is used as the maximum value for the Y axis.
ymax_type	integer	Maximum value calculation method for the Y axis. Possible values: 0 - (default) calculated; 1 - fixed; 2 - item.
ymin_itemid	string	ID of the item that is used as the minimum value for the Y axis.
ymin_type	integer	Minimum value calculation method for the Y axis. Possible values: 0 - (default) calculated; 1 - fixed; 2 - item.

graphprototype.create

Description

`object graphprototype.create(object/array graphPrototypes)`

This method allows to create new graph prototypes.

Parameters

(object/array) Graph prototypes to create.

Additionally to the [standard graph prototype properties](#), the method accepts the following parameters.

Parameter	Type	Description
gitems (required)	array	Graph items to be created for the graph prototypes. Graph items can reference both items and item prototypes, but at least one item prototype must be present.

Return values

(object) Returns an object containing the IDs of the created graph prototypes under the `graphids` property. The order of the returned IDs matches the order of the passed graph prototypes.

Examples

Creating a graph prototype

Create a graph prototype with two items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graphprototype.create",
  "params": {
    "name": "Disk space usage {#FSNAME}",
    "width": 900,
    "height": 200,
    "gitems": [
      {
        "itemid": "22828",
        "color": "00AA00"
      },
      {
        "itemid": "22829",
        "color": "3333FF"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "graphids": [
      "652"
    ]
  },
  "id": 1
}
```

See also

- [Graph item](#)

Source

`CGraphPrototype::create()` in *frontends/php/api/classes/CGraphPrototype.php*.

graphprototype.delete

Description

`object graphprototype.delete(array graphPrototypeIds)`

This method allows to delete graph prototypes.

Parameters

(array) IDs of the graph prototypes to delete.

Return values

(object) Returns an object containing the IDs of the deleted graph prototypes under the `graphids` property.

Examples

Deleting multiple graph prototypes

Delete two graph prototypes.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graphprototype.delete",
  "params": [
    "652",
    "653"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "graphids": [
      "652",
      "653"
    ]
  },
  "id": 1
}
```

Source

CGraphPrototype::delete() in *frontends/php/api/classes/CGraphPrototype.php*.

graphprototype.exists

Description

boolean graphprototype.exists(object filter)

This method checks if at least one graph prototype that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
host	string/array	Technical names of the hosts that the graph prototypes belong to.
hostids	string/array	IDs of the hosts that the graph prototypes belong to.
name	string/array	Names of the graph prototypes.
node	string	Name of the node the graph prototypes must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the graph prototypes must belong to.

Return values

(boolean) Returns true if at least one graph prototype that matches the given filter criteria exists.

Examples

Checking a graph prototype on a host

Check if graph prototype "Disk space usage {#FSNAME}" exists on host "Zabbix server".

Request:


```
{
  "jsonrpc": "2.0",
  "method": "graphprototype.exists",
  "params": {
    "name": "Disk space usage {#FSNAME}",
    "host": "Zabbix server"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

CGraphPrototype::exists() in *frontends/php/api/classes/CGraphPrototype.php*.

graphprototype.get

Description

`integer/array graphprototype.get(object parameters)`

The method allows to retrieve graph prototypes according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
discoveryids	string/array	Return only graph prototypes that belong to the given discovery rules.
graphids	string/array	Return only graph prototypes with the given IDs.
groupids	string/array	Return only graph prototypes that belong to hosts in the given host groups.
hostids	string/array	Return only graph prototypes that belong to the given hosts.
inherited	boolean	If set to <code>true</code> return only graph prototypes inherited from a template.
itemids	string/array	Return only graph prototypes that contain the given item prototypes.
templated	boolean	If set to <code>true</code> return only graph prototypes that belong to templates.
templateids	string/array	Return only graph prototypes that belong to the given templates.
selectDiscoveryRule	query	Return the LLD rule that the graph prototype belongs to in the <code>discoveryRule</code> property.
selectGraphItems	query	Return an object containing graph items used in the graph prototype in the <code>gitems</code> property.
selectGroups	query	Does not support <code>shorten</code> or arrays of properties. Return the host groups that the graph prototype belongs to in the <code>groups</code> property.

Does not support `shorten` or arrays of properties.

Parameter	Type	Description
selectHosts	query	Return the hosts that the graph prototype belongs to in the <code>hosts</code> property.
selectItems	query	Does not support <code>shorten</code> or arrays of properties. Return the items and item prototypes used in the graph prototype in the <code>items</code> property.
selectTemplates	query	Does not support <code>shorten</code> or arrays of properties. Return the templates that the graph prototype belongs to in the <code>templates</code> property.
filter	object	Does not support <code>shorten</code> or arrays of properties. Return only those results that exactly match the given filter. Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Supports additional filters: <code>host</code> - technical name of the host that the graph prototype belongs to; <code>hostid</code> - ID of the host that the graph prototype belongs to.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: <code>graphid</code> , <code>name</code> and <code>graphtype</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving graph prototypes from a LLD rule

Retrieve all graph prototypes from an LLD rule.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graphprototype.get",
  "params": {
    "output": "extend",
    "discoveryids": "27426"
  },
}
```

```

    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": [
    {
      "graphid": "1017",
      "parent_itemid": "27426",
      "name": "Disk space usage {#FSNAME}",
      "width": "600",
      "height": "340",
      "yaxismin": "0.0000",
      "yaxismax": "0.0000",
      "templateid": "442",
      "show_work_period": "0",
      "show_triggers": "0",
      "graphtype": "2",
      "show_legend": "1",
      "show_3d": "1",
      "percent_left": "0.0000",
      "percent_right": "0.0000",
      "ymin_type": "0",
      "ymax_type": "0",
      "ymin_itemid": "0",
      "ymax_itemid": "0"
    }
  ],
  "id": 1
}

```

See also

- [graphprototype.getobjects](#)
- [Discovery rule](#)
- [Graph item](#)
- [Item](#)
- [Host](#)
- [Host group](#)
- [Template](#)

Source

`CGraphPrototype::get()` in *frontends/php/api/classes/CGraphPrototype.php*.

graphprototype.getobjects

Description

`array graphprototype.getobjects(object filter)`

This method allows to retrieve graph prototypes that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard graph prototype properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the graph prototypes must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the graph prototypes must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving graph prototypes from a host

Retrieve all graph prototypes from host "Zabbix server".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graphprototype.getobjects",
  "params": {
    "host": "Zabbix server"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "graphid": "1017",
      "name": "Disk space usage {#FSNAME}",
      "width": "600",
      "height": "340",
      "yaxismin": "0.0000",
      "yaxismax": "0.0000",
      "templateid": "442",
      "show_work_period": "0",
      "show_triggers": "0",
      "graphtype": "2",
      "show_legend": "1",
      "show_3d": "1",
      "percent_left": "0.0000",
      "percent_right": "0.0000",
      "ymin_type": "0",
      "ymax_type": "0",
      "ymin_itemid": "0",
      "ymax_itemid": "0"
    }
  ],
  "id": 1
}
```

See also

- [graphprototype.get](#)

Source

CGraphPrototype::getObject() in *frontends/php/api/classes/CGraphPrototype.php*.

graphprototype.update

Description

object graphprototype.update(object/array graphPrototypes)

This method allows to update existing graph prototypes.

Parameters

(object/array) Graph prototype properties to be updated.

The graphid property must be defined for each graph prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the **standard graph prototype properties**, the method accepts the following parameters.

Parameter	Type	Description
gitems	array	Graph items to replace existing graph items. If a graph item has the gitemid property defined it will be updated, otherwise a new graph item will be created.

Return values

(object) Returns an object containing the IDs of the updated graph prototypes under the graphids property.

Examples

Changing the size of a graph prototype

Change the size of a graph prototype to 1100 to 400 pixels.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "graphprototype.update",
  "params": {
    "graphid": "439",
    "width": 1100,
    "height": 400
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "graphids": [
      "439"
    ]
  },
  "id": 1
}
```

Source

CGraphPrototype::update() in *frontends/php/api/classes/CGraphPrototype.php*.

History

This class is designed to work with history data.

Object references:

- **History**

Available methods:

- **history.get** - retrieving history data.

> History object

The following objects are directly related to the `history` API.

Note:

History objects differ depending on the item's type of information. They are created by the Zabbix server and cannot be modified via the API.

Float history

The float history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	string	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	float	Received value.

Integer history

The integer history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	string	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	integer	Received value.

String history

The string history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	string	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	string	Received value.

Text history

The text history object has the following properties.

Property	Type	Description
id	string	ID of the history entry.
clock	timestamp	Time when that value was received.
itemid	string	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	string	Received value.

Log history

The log history object has the following properties.

Property	Type	Description
id	string	ID of the history entry.
clock	timestamp	Time when that value was received.
itemid	string	ID of the related item.
logeventid	integer	Windows event log entry ID.
ns	integer	Nanoseconds when the value was received.
severity	integer	Windows event log entry level.
source	string	Windows event log entry source.
timestamp	timestamp	Windows event log entry time.
value	string	Received value.

history.get

Description

`integer/array history.get(object parameters)`

The method allows to retrieve history data according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
history	integer	History object types to return. Possible values: 0 - numeric float; 1 - character; 2 - log; 3 - numeric unsigned; 4 - text. Default: 3.
hostids	string/array	Return only history from the given hosts.
itemids	string/array	Return only history from the given items.
time_from	timestamp	Return only values that have been received after or at the given time.
time_till	timestamp	Return only values that have been received before or at the given time.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: <code>itemid</code> and <code>clock</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving item history data

Return 10 latest values received from a numeric(float) item.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "history.get",
  "params": {
    "output": "extend",
    "history": 0,
    "itemids": "23296",
    "sortfield": "clock",
    "sortorder": "DESC",
    "limit": 10
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "23296",
      "clock": "1351090996",
      "value": "0.0850",
      "ns": "563157632"
    },
    {
      "itemid": "23296",
      "clock": "1351090936",
      "value": "0.1600",
      "ns": "549216402"
    },
    {
      "itemid": "23296",
      "clock": "1351090876",
      "value": "0.1800",
      "ns": "537418114"
    },
    {
      "itemid": "23296",
      "clock": "1351090816",
      "value": "0.2100",
      "ns": "522659528"
    },
    {
      "itemid": "23296",
      "clock": "1351090756",
      "value": "0.2150",
      "ns": "507809457"
    },
    {
      "itemid": "23296",
      "clock": "1351090696",
      "value": "0.2550",
      "ns": "495509699"
    }
  ]
}
```



```

    },
    {
        "itemid": "23296",
        "clock": "1351090636",
        "value": "0.3600",
        "ns": "477708209"
    },
    {
        "itemid": "23296",
        "clock": "1351090576",
        "value": "0.3750",
        "ns": "463251343"
    },
    {
        "itemid": "23296",
        "clock": "1351090516",
        "value": "0.3150",
        "ns": "447947017"
    },
    {
        "itemid": "23296",
        "clock": "1351090456",
        "value": "0.2750",
        "ns": "435307141"
    }
],
    "id": 1
}

```

Source

CHistory::get() in *frontends/php/api/classes/CHistory.php*.

Host

This class is designed to work with hosts.

Object references:

- [Host](#)
- [Host inventory](#)

Available methods:

- [host.create](#) - creating new hosts
- [host.delete](#) - deleting hosts
- [host.exists](#) - checking if a host exists
- [host.get](#) - retrieving hosts
- [host.isreadable](#) - checking if hosts are readable
- [host.iswritable](#) - checking if hosts are writable
- [host.massadd](#) - adding related objects to hosts
- [host.massremove](#) - removing related objects from hosts
- [host.massupdate](#) - replacing or removing related objects from hosts
- [host.update](#) - updating hosts

> Host object

The following objects are directly related to the host API.

Host

The host object has the following properties.

Property	Type	Description
hostid	string	<i>(readonly)</i> ID of the host.
host	string	Technical name of the host.
(required)		
available	integer	<i>(readonly)</i> Availability of Zabbix agent. Possible values are: 0 - <i>(default)</i> unknown; 1 - available; 2 - unavailable.
disable_until	timestamp	<i>(readonly)</i> The next polling time of an unavailable Zabbix agent.
error	string	<i>(readonly)</i> Error text if Zabbix agent is unavailable.
errors_from	timestamp	<i>(readonly)</i> Time when Zabbix agent became unavailable.
inventory_mode	integer	Host inventory population mode. Possible values are: -1 - disabled; 0 - <i>(default)</i> manual; 1 - automatic.
ipmi_authtype	integer	IPMI authentication algorithm. Possible values are: -1 - <i>(default)</i> default; 0 - none; 1 - MD2; 2 - MD5 4 - straight; 5 - OEM; 6 - RMCP+.
ipmi_available	integer	<i>(readonly)</i> Availability of IPMI agent. Possible values are: 0 - <i>(default)</i> unknown; 1 - available; 2 - unavailable.
ipmi_disable_until	timestamp	<i>(readonly)</i> The next polling time of an unavailable IPMI agent.
ipmi_error	string	<i>(readonly)</i> Error text if IPMI agent is unavailable.
ipmi_errors_from	timestamp	<i>(readonly)</i> Time when IPMI agent became unavailable.
ipmi_password	string	IPMI password.
ipmi_privilege	integer	IPMI privilege level. Possible values are: 1 - callback; 2 - <i>(default)</i> user; 3 - operator; 4 - admin; 5 - OEM.
ipmi_username	string	IPMI username.
jmx_available	integer	<i>(readonly)</i> Availability of JMX agent. Possible values are: 0 - <i>(default)</i> unknown; 1 - available; 2 - unavailable.
jmx_disable_until	timestamp	<i>(readonly)</i> The next polling time of an unavailable JMX agent.
jmx_error	string	<i>(readonly)</i> Error text if JMX agent is unavailable.
jmx_errors_from	timestamp	<i>(readonly)</i> Time when JMX agent became unavailable.
maintenance_from	timestamp	<i>(readonly)</i> Starting time of the effective maintenance.

Property	Type	Description
maintenance_status	integer	(readonly) Effective maintenance status. Possible values are: 0 - (default) no maintenance; 1 - maintenance in effect.
maintenance_type	integer	(readonly) Effective maintenance type. Possible values are: 0 - (default) maintenance with data collection; 1 - maintenance without data collection.
maintenanceid	string	(readonly) ID of the maintenance that is currently in effect on the host.
name	string	Visible name of the host.
proxy_hostid	string	Default: host property value. ID of the proxy that is used to monitor the host.
snmp_available	integer	(readonly) Availability of SNMP agent. Possible values are: 0 - (default) unknown; 1 - available; 2 - unavailable.
snmp_disable_until	timestamp	(readonly) The next polling time of an unavailable SNMP agent.
snmp_error	string	(readonly) Error text if SNMP agent is unavailable.
snmp_errors_from	timestamp	(readonly) Time when SNMP agent became unavailable.
status	integer	Status and function of the host. Possible values are: 0 - (default) monitored host; 1 - unmonitored host.

Host inventory

The host inventory object has the following properties.

Note:

Each property has it's own unique ID number, which is used to associate host inventory fields with items.

ID	Property	Type	Description
4	alias	string	Alias.
11	asset_tag	string	Asset tag.
28	chassis	string	Chassis.
23	contact	string	Contact person.
32	contract_number	string	Contract number.
47	date_hw_decomm	string	HW decommissioning date.
46	date_hw_expiry	string	HW maintenance expiry date.
45	date_hw_install	string	HW installation date.
44	date_hw_purchase	string	HW purchase date.
34	deployment_status	string	Deployment status.
14	hardware	string	Hardware.
15	hardware_full	string	Detailed hardware.
39	host_netmask	string	Host subnet mask.
38	host_networks	string	Host networks.
40	host_router	string	Host router.
30	hw_arch	string	HW architecture.
33	installer_name	string	Installer name.
24	location	string	Location.
25	location_lat	string	Location latitude.
26	location_lon	string	Location longitude.

ID	Property	Type	Description
12	macaddress_a	string	MAC address A.
13	macaddress_b	string	MAC address B.
29	model	string	Model.
3	name	string	Name.
27	notes	string	Notes.
41	oob_ip	string	OOB IP address.
42	oob_netmask	string	OOB host subnet mask.
43	oob_router	string	OOB router.
5	os	string	OS name.
6	os_full	string	Detailed OS name.
7	os_short	string	Short OS name.
61	poc_1_cell	string	Primary POC mobile number.
58	poc_1_email	string	Primary email.
57	poc_1_name	string	Primary POC name.
63	poc_1_notes	string	Primary POC notes.
59	poc_1_phone_a	string	Primary POC phone A.
60	poc_1_phone_b	string	Primary POC phone B.
62	poc_1_screen	string	Primary POC screen name.
68	poc_2_cell	string	Secondary POC mobile number.
65	poc_2_email	string	Secondary POC email.
64	poc_2_name	string	Secondary POC name.
70	poc_2_notes	string	Secondary POC notes.
66	poc_2_phone_a	string	Secondary POC phone A.
67	poc_2_phone_b	string	Secondary POC phone B.
69	poc_2_screen	string	Secondary POC screen name.
8	serialno_a	string	Serial number A.
9	serialno_b	string	Serial number B.
48	site_address_a	string	Site address A.
49	site_address_b	string	Site address B.
50	site_address_c	string	Site address C.
51	site_city	string	Site city.
53	site_country	string	Site country.
56	site_notes	string	Site notes.
55	site_rack	string	Site rack location.
52	site_state	string	Site state.
54	site_zip	string	Site ZIP/postal code.
16	software	string	Software.
18	software_app_a	string	Software application A.
19	software_app_b	string	Software application B.
20	software_app_c	string	Software application C.
21	software_app_d	string	Software application D.
22	software_app_e	string	Software application E.
17	software_full	string	Software details.
10	tag	string	Tag.
1	type	string	Type.
2	type_full	string	Type details.
35	url_a	string	URL A.
36	url_b	string	URL B.
37	url_c	string	URL C.
31	vendor	string	Vendor.

host.create

Description

object `host.create(object/array hosts)`

This method allows to create new hosts.

Parameters

(object/array) Hosts to create.

Additionally to the **standard host properties**, the method accepts the following parameters.

Parameter	Type	Description
groups (required)	object/array	Host groups to add the host to. The host groups must have the <code>groupid</code> property defined.
interfaces (required)	object/array	Interfaces to be created for the host.
templates	object/array	Templates to be linked to the host. The templates must have the <code>templateid</code> property defined.
macros	object/array	User macros to be created for the host.
inventory	object	Host inventory properties.

Return values

(object) Returns an object containing the IDs of the created hosts under the `hostids` property. The order of the returned IDs matches the order of the passed hosts.

Examples

Creating a host

Create a host called "Linux server" with an IP interface, add it to a group, link a template to it and set the MAC addresses in the host inventory.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.create",
  "params": {
    "host": "Linux server",
    "interfaces": [
      {
        "type": 1,
        "main": 1,
        "useip": 1,
        "ip": "192.168.3.1",
        "dns": "",
        "port": "10050"
      }
    ],
    "groups": [
      {
        "groupid": "50"
      }
    ],
    "templates": [
      {
        "templateid": "20045"
      }
    ],
    "inventory_mode": 0,
    "inventory": {
      "macaddress_a": "01234",
      "macaddress_b": "56768"
    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "107819"
    ]
  },
  "id": 1
}
```

See also

- [Host group](#)
- [Template](#)
- [User macro](#)
- [Host interface](#)
- [Host inventory](#)

Source

`CHost::create()` in *frontends/php/api/classes/CHost.php*.

host.delete

Description

`object host.delete(array hosts)`

This method allows to delete hosts.

Parameters

(array) Hosts to delete.

The hosts must have the `hostid` property defined.

Return values

(object) Returns an object containing the IDs of the deleted hosts under the `hostids` property.

Examples

Deleting multiple hosts

Delete two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.delete",
  "params": [
    {"hostid": "13"},
    {"hostid": "32"}
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "13",
      "32"
    ]
  },
}
```

```
    "id": 1
}
```

Source

`CHost::delete()` in *frontends/php/api/classes/CHost.php*.

host.exists

Description

`boolean host.exists(object filter)`

This method checks if at least one host that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
hostid	string/array	Host IDs.
host	string/array	Technical names of the hosts.
name	string/array	Visible names of the hosts.
node	string	Name of the node the hosts must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the node the hosts must belong to.

Return values

(boolean) Returns `true` if at least one host that matches the given filter criteria exists.

Examples

Check host on a node

Check if a host with the technical name "Zabbix Server" exists on the node with ID 1.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.exists",
  "params": {
    "host": "Zabbix Server",
    "nodeids": [
      "1"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [host.isreadable](#)
- [host.iswritable](#)

Source

CHost::exists() in *frontends/php/api/classes/CHost.php*.

host.get

Description

integer/array host.get(object parameters)

The method allows to retrieve hosts according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	string/array	Return only hosts that belong to the given groups.
applicationids	string/array	Return only hosts that have the given applications.
dserviceids	string/array	Return only hosts that are related to the given discovered services.
graphids	string/array	Return only hosts that have the given graphs.
hostids	string/array	Return only hosts with the given host IDs.
httptestids	string/array	Return only hosts that have the given web checks.
interfaceids	string/array	Return only hosts that use the given interfaces.
itemids	string/array	Return only hosts that have the given items.
maintenanceids	string/array	Return only hosts that are affected by the given maintenances.
monitored_hosts	flag	Return only monitored hosts.
proxy_hosts	flag	Return only proxies.
proxyids	string/array	Return only hosts that are monitored by the given proxies.
templated_hosts	flag	Return both hosts and templates.
templateids	string/array	Return only hosts that are linked to the given templates.
triggerids	string/array	Return only hosts that have the given triggers.
with_items	flag	Return only hosts that have items.
		Overrides the with_monitored_items, with_historical_items and with_simple_graph_items parameters.
with_applications	flag	Return only hosts that have applications.
with_graphs	flag	Return only hosts that have graphs.
with_historical_items	flag	Return only hosts that have enabled or not supported items with historical data.
		Overrides the with_simple_graph_items parameter.
with_httptests	flag	Return only hosts that have web checks.
		Overrides the with_monitored_httptests parameter.
with_monitored_httptests	flag	Return only hosts that have enabled web checks.
with_monitored_items	flag	Return only hosts that have enabled items.
		Overrides the with_historical_items and with_simple_graph_items parameters.
with_monitored_triggers	flag	Return only hosts that have enabled triggers. All of the items used in the trigger must also be enabled.
with_simple_graph_items	flag	Return only hosts that have items with numeric type of information.

Parameter	Type	Description
with_triggers	flag	Return only hosts that have triggers.
withInventory	flag	Overrides the with_monitored_triggers parameter. Return only hosts that have inventory data.
selectGroups	query	Return the host groups that the host belongs to in the groups property.
selectApplications	query	Does not support shorten or arrays of properties. Return the applications from the host in the applications property.
selectDiscoveries	query	Supports count; does not support shorten. Return the low level discoveries from the host in the discoveries property.
selectGraphs	query	Supports count; does not support shorten. Return the graphs from the host in the graphs property.
selectInterfaces	query	Supports count; does not support shorten. Return the host interfaces in the interfaces property.
selectInventory	boolean/array	Supports count; does not support shorten. Return the host inventory from the host in the inventory property.
selectItems	query	Possible values are true to return all of the data, or an array of property names to return only specific properties. Return the items from the host in the items property.
selectMacros	query	Supports count; does not support shorten. Return the macros from the host in the macros property.
selectParentTemplates	query	Does not support shorten or arrays of properties. Return the templates that the host is linked to in the parentTemplates property.
selectScreens	query	Supports count; does not support shorten. Return the screens from the host in the screens property.
selectTriggers	query	Supports count; does not support shorten. Return the triggers from the host in the triggers property.
filter	object	Supports count; does not support shorten. Return only those results that exactly match the given filter. Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Allows filtering by interface properties.

Parameter	Type	Description
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: selectParentTemplates - results will be sorted by host; selectInterfaces; selectItems - sorted by name; selectDiscoveries - sorted by name; selectTriggers - sorted by description; selectGraphs - sorted by name; selectApplications - sorted by name; selectScreens - sorted by name.
search	object	Return results that match the given wildcard search. Accepts an array, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a LIKE "%...%" search.
sortfield	string/array	Allows searching by interface properties. Works only with text fields. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>hostid</code> , <code>host</code> , <code>name</code> , <code>status</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving data by name

Retrieve all data about two hosts named "Zabbix server" and "Linux server".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.get",
  "params": {
    "output": "extend",
    "filter": {
      "host": [
        "Zabbix server",
        "Linux server"
      ]
    }
  }
}
```

```

    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": [
    {
      "maintenances": [],
      "hostid": "10160",
      "proxy_hostid": "0",
      "host": "Zabbix server",
      "status": "0",
      "disable_until": "0",
      "error": "",
      "available": "0",
      "errors_from": "0",
      "lastaccess": "0",
      "ipmi_authtype": "-1",
      "ipmi_privilege": "2",
      "ipmi_username": "",
      "ipmi_password": "",
      "ipmi_disable_until": "0",
      "ipmi_available": "0",
      "snmp_disable_until": "0",
      "snmp_available": "0",
      "maintenanceid": "0",
      "maintenance_status": "0",
      "maintenance_type": "0",
      "maintenance_from": "0",
      "ipmi_errors_from": "0",
      "snmp_errors_from": "0",
      "ipmi_error": "",
      "snmp_error": "",
      "jmx_disable_until": "0",
      "jmx_available": "0",
      "jmx_errors_from": "0",
      "jmx_error": "",
      "name": "Zabbix server"
    },
    {
      "maintenances": [],
      "hostid": "10167",
      "proxy_hostid": "0",
      "host": "Linux server",
      "status": "0",
      "disable_until": "0",
      "error": "",
      "available": "0",
      "errors_from": "0",
      "lastaccess": "0",
      "ipmi_authtype": "-1",
      "ipmi_privilege": "2",
      "ipmi_username": "",
      "ipmi_password": "",
      "ipmi_disable_until": "0",
      "ipmi_available": "0",
      "snmp_disable_until": "0",
      "snmp_available": "0",

```

```

        "maintenanceid": "0",
        "maintenance_status": "0",
        "maintenance_type": "0",
        "maintenance_from": "0",
        "ipmi_errors_from": "0",
        "snmp_errors_from": "0",
        "ipmi_error": "",
        "snmp_error": "",
        "jmx_disable_until": "0",
        "jmx_available": "0",
        "jmx_errors_from": "0",
        "jmx_error": "",
        "name": "Linux server"
    }
],
    "id": 1
}

```

Retrieving host groups

Retrieve names of the groups host "Zabbix server" is member of, but no host details themselves.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid"],
        "selectGroups": "extend",
        "filter": {
            "host": [
                "Zabbix server"
            ]
        }
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 2
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": [
        {
            "groups": [
                {
                    "groupid": "2",
                    "name": "Linux servers",
                    "internal": "0"
                },
                {
                    "groupid": "4",
                    "name": "Zabbix servers",
                    "internal": "0"
                }
            ],
            "hostid": "10085"
        }
    ],
    "id": 2
}

```

Retrieving linked templates

Retrieve the IDs and names of templates linked to host "10084".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.get",
  "params": {
    "output": ["hostid"],
    "selectParentTemplates": [
      "templateid",
      "name"
    ],
    "hostids": "10084"
  },
  "id": 1,
  "auth": "70785d2b494a7302309b48afcdb3a401"
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "parentTemplates": [
        {
          "name": "Template OS Linux",
          "templateid": "10001"
        },
        {
          "name": "Template App Zabbix Server",
          "templateid": "10047"
        }
      ],
      "hostid": "10084"
    }
  ],
  "id": 1
}
```

See also

- [Host group](#)
- [Template](#)
- [User macro](#)
- [Host interface](#)

Source

`CHost::get()` in *frontends/php/api/classes/CHost.php*.

host.getobjects

Description

array `host.getobjects(object filter)`

This method allows to retrieve hosts that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard **standard host properties** the following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the hosts must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the hosts must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving a host by name

Retrieve the host with the technical name "Zabbix server".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.getobjects",
  "params": {
    "name": "Zabbix server"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "maintenances": [],
      "hostid": "10084",
      "proxy_hostid": "0",
      "host": "Zabbix server",
      "status": "0",
      "disable_until": "0",
      "error": "",
      "available": "1",
      "errors_from": "0",
      "lastaccess": "0",
      "ipmi_authtype": "-1",
      "ipmi_privilege": "2",
      "ipmi_username": "",
      "ipmi_password": "",
      "ipmi_disable_until": "0",
      "ipmi_available": "0",
      "snmp_disable_until": "0",
      "snmp_available": "0",
      "maintenanceid": "0",
      "maintenance_status": "0",
      "maintenance_type": "0",
      "maintenance_from": "0",
      "ipmi_errors_from": "0",
      "snmp_errors_from": "0",
      "ipmi_error": "",
      "snmp_error": "",
      "jmx_disable_until": "0",
      "jmx_available": "0",
      "jmx_errors_from": "0",
      "jmx_error": "",
      "name": "Zabbix server"
    }
  ]
}
```

```

    }
  ],
  "id": 1
}

```

See also

- [host.get](#)

Source

`CHost::getObject()` in *frontends/php/api/classes/CHost.php*.

host.isreadable

Description

`boolean host.isreadable(array hostIds)`

This method checks if the given hosts are available for reading.

Parameters

(array) IDs of the hosts to check.

Return values

(boolean) Returns `true` if the given hosts are available for reading.

Examples

Check multiple hosts

Check if the two hosts are readable.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "host.isreadable",
  "params": [
    "143",
    "943"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}

```

See also

- [host.exists](#)
- [host.iswritable](#)

Source

`CHost::isReadable()` in *frontends/php/api/classes/CHost.php*.

host.iswritable

Description

`boolean host.iswritable(array hostIds)`

This method checks if the given hosts are available for writing.

Parameters

(array) IDs of the hosts to check.

Return values

(boolean) Returns true if the given hosts are available for writing.

Examples

Check multiple hosts

Check if the two hosts are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.iswritable",
  "params": [
    "143",
    "943"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [host.isreadable](#)
- [host.exists](#)

Source

CHost::isWritable() in *frontends/php/api/classes/CHost.php*.

host.massadd

Description

object host.massadd(object parameters)

This method allows to simultaneously add multiple related objects to all the given hosts.

Parameters

(object) Parameters containing the IDs of the hosts to update and the objects to add to all the hosts.

The method accepts the following parameters.

Parameter	Type	Description
hosts (required)	object/array	Hosts to be updated.
groups	object/array	The hosts must have the <code>hostid</code> property defined. Host groups to add to the given hosts.
interfaces	object/array	The host groups must have the <code>groupid</code> property defined. Host interfaces to be created for the given hosts.
macros	object/array	User macros to be created for the given hosts.

Parameter	Type	Description
templates	object/array	Templates to link to the given hosts. The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated hosts under the `hostids` property.

Examples

Adding macros

Add two new macros to two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.massadd",
  "params": {
    "hosts": [
      {
        "hostid": "10160"
      },
      {
        "hostid": "10167"
      }
    ],
    "macros": [
      {
        "macro": "${TEST1}",
        "value": "MACROTEST1"
      },
      {
        "macro": "${TEST2}",
        "value": "MACROTEST2"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "10160",
      "10167"
    ]
  },
  "id": 1
}
```

See also

- [host.update](#)
- [Host group](#)
- [Template](#)
- [User macro](#)
- [Host interface](#)

Source

`CHost::massAdd()` in *frontends/php/api/classes/CHost.php*.

host.massremove

Description

`object host.massremove(object parameters)`

This method allows to remove related objects from multiple hosts.

Parameters

(object) Parameters containing the IDs of the hosts to update and the objects that should be removed.

Parameter	Type	Description
hostids (required)	string/array	IDs of the hosts to be updated.
groupids	string/array	Host groups to remove the given hosts from.
interfaces	object/array	Host interfaces to remove from the given hosts. The host interface object must have the ip, dns and port properties defined.
macros	string/array	User macros to delete from the given hosts.
templateids	string/array	Templates to unlink from the given hosts.
templateids_clear	string/array	Templates to unlink and clear from the given hosts.

Return values

(object) Returns an object containing the IDs of the updated hosts under the `hostids` property.

Examples

Unlinking templates

Unlink a template from two hosts and delete all of the templated entities.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.massremove",
  "params": {
    "hostids": ["69665", "69666"],
    "templateids_clear": "325"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "69665",
      "69666"
    ]
  },
  "id": 1
}
```

See also

- [host.update](#)
- [User macro](#)
- [Host interface](#)

Source

CHost::massRemove() in *frontends/php/api/classes/CHost.php*.

host.massupdate

Description

object host.massupdate(object parameters)

This method allows to simultaneously replace or remove related objects and update properties on multiple hosts.

Parameters

(object) Parameters containing the IDs of the hosts to update and the properties that should be updated.

Additionally to the **standard host properties**, the method accepts the following parameters.

Parameter	Type	Description
hosts (required)	object/array	Hosts to be updated.
groups	object/array	The hosts must have the <code>hostid</code> property defined. Host groups to replace the current host groups the hosts belong to.
interfaces	object/array	The host groups must have the <code>groupid</code> property defined. Host interfaces to replace the current host interfaces on the given hosts.
macros	object/array	User macros to replace the current user macros on the given hosts.
templates	object/array	Templates to replace the currently linked templates on the given hosts.
templates_clear	object/array	The templates must have the <code>templateid</code> property defined. Templates to unlink and clear from the given hosts.
		The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated hosts under the `hostids` property.

Examples

Enabling multiple hosts

Enable monitoring of two hosts, i.e., set their status to 0.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.massupdate",
  "params": {
    "hosts": [
      {
        "hostid": "69665"
      },
      {
        "hostid": "69666"
      }
    ],
    "status": 0
  },
}
```

```

    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "69665",
            "69666"
        ]
    },
    "id": 1
}

```

See also

- [host.update](#)
- [host.massadd](#)
- [host.massremove](#)
- [Host group](#)
- [Template](#)
- [User macro](#)
- [Host interface](#)

Source

`CHost::massUpdate()` in *frontends/php/api/classes/CHost.php*.

host.update

Description

`object host.update(object/array hosts)`

This method allows to update existing hosts.

Parameters

(object/array) Host properties to be updated.

The `hostid` property must be defined for each host, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Additionally to the [standard host properties](#), the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to replace the current host groups the host belongs to. The host groups must have the <code>groupid</code> property defined.
interfaces	object/array	Host interfaces to replace the current host interfaces. <i>The method accepts a single interface object starting from 2.0.2.</i>
inventory	object	Host inventory properties.
macros	object/array	User macros to replace the current user macros.
templates	object/array	Templates to replace the currently linked templates. Templates that are not passed are only unlinked. The templates must have the <code>templateid</code> property defined.

Parameter	Type	Description
templates_clear	object/array	Templates to unlink and clear from the host. The templates must have the templateid property defined.

Note:

As opposed to the Zabbix frontend, when name is the same as host, updating host will not automatically update name. Both properties need to be updated explicitly.

Return values

(object) Returns an object containing the IDs of the updated hosts under the hostids property.

Examples

Enabling a host

Enable host monitoring, i.e. set its status to 0.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.update",
  "params": {
    "hostid": "10126",
    "status": 0
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "10126"
    ]
  },
  "id": 1
}
```

Unlinking templates

Unlink and clear two templates from host.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.update",
  "params": {
    "hostid": "10126",
    "templates_clear": [
      {
        "templateid": "10124"
      },
      {
        "templateid": "10125"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "10126"
    ]
  },
  "id": 1
}
```

Updating host macros

Replace all host macros with two new ones.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.update",
  "params": {
    "hostid": "10126",
    "macros": [
      {
        "macro": "${PASS}",
        "value": "password"
      },
      {
        "macro": "${DISC}",
        "value": "sda"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "10126"
    ]
  },
  "id": 1
}
```

Updating host inventory

Change inventory mode and add location

Request:

```
{
  "jsonrpc": "2.0",
  "method": "host.update",
  "params": {
    "hostid": "10387",
    "inventory_mode": 0,
    "inventory": {
      "location": "Latvia, Riga"
    }
  }
}
```

```

    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "hostids": [
      "10387"
    ]
  },
  "id": 2
}

```

See also

- [host.massadd](#)
- [host.massupdate](#)
- [host.massremove](#)
- [Host group](#)
- [Template](#)
- [User macro](#)
- [Host interface](#)
- [Host inventory](#)

Source

`CHost::update()` in *frontends/php/api/classes/CHost.php*.

Host group

This class is designed to work with host groups.

Object references:

- [Host group](#)

Available methods:

- [hostgroup.create](#) - creating new host groups
- [hostgroup.delete](#) - deleting host groups
- [hostgroup.exists](#) - checking if a host group exists
- [hostgroup.get](#) - retrieving host groups
- [hostgroup.getobjects](#) - retrieving host groups by filters
- [hostgroup.isreadable](#) - checking if host groups are readable
- [hostgroup.iswritable](#) - checking if host groups are writable
- [hostgroup.massadd](#) - adding related objects to host groups
- [hostgroup.massremove](#) - removing related objects from host groups
- [hostgroup.massupdate](#) - replacing or removing related objects from host groups
- [hostgroup.update](#) - updating host groups

> Host group object

The following objects are directly related to the `hostgroup` API.

Host group

The host group object has the following properties.

Property	Type	Description
<code>groupid</code>	string	(<i>readonly</i>) ID of the host group.

Property	Type	Description
name (required)	string	Name of the host group.
internal	integer	<i>(readonly)</i> Whether the group is used internally by the system. An internal group cannot be deleted. Possible values: 0 - <i>(default)</i> not internal; 1 - internal.

hostgroup.create

Description

`object hostgroup.create(object/array hostGroups)`

This method allows to create new host groups.

Parameters

(object/array) Host groups to create. The method accepts host groups with the **standard host group properties**.

Return values

(object) Returns an object containing the IDs of the created host groups under the `groupids` property. The order of the returned IDs matches the order of the passed host groups.

Examples

Creating a host group

Create a host group called "Linux servers".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.create",
  "params": {
    "name": "Linux servers"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "groupids": [
      "107819"
    ]
  },
  "id": 1
}
```

Source

`CHostGroup::create()` in `frontends/php/api/classes/CHostGroup.php`.

hostgroup.delete

Description

`object hostgroup.delete(array hostGroupIds)`

This method allows to delete host groups.

A host group can not be deleted if:

- it contains hosts that belong to this group only;
- it's marked as internal;
- it is used in a global script.

Parameters

(array) IDs of the host groups to delete.

Return values

(object) Returns an object containing the IDs of the deleted host groups under the `groupids` property.

Examples

Deleting multiple host groups

Delete two host groups.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.delete",
  "params": [
    "107824",
    "107825"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "groupids": [
      "107824",
      "107825"
    ]
  },
  "id": 1
}
```

Source

`CHostGroup::delete()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.exists

Description

`boolean hostgroup.exists(object filter)`

This method checks if at least one host group that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
<code>groupid</code>	string/array	Host group IDs.
<code>name</code>	string/array	Names of the host groups.
<code>node</code>	string	Name of the node the host groups must belong to.

This will override the `nodeids` parameter.

Parameter	Type	Description
nodeids	string/array	IDs of the nodes the host groups must belong to.

Return values

(boolean) Returns true if at least one host group that matches the given filter criteria exists.

Examples

Check host group on a node

Check if a host group called "Zabbix servers" exists on the node with ID 1.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.exists",
  "params": {
    "name": "Linux servers",
    "nodeids": [
      "1"
    ]
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [hostgroup.isreadable](#)
- [hostgroup.iswritable](#)

Source

CHostGroup::exists() in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.get

Description

integer/array hostgroup.get(object parameters)

The method allows to retrieve host groups according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
graphids	string/array	Return only host groups that contain hosts or templates with the given graphs.
groupids	string/array	Return only host groups with the given host group IDs.
hostids	string/array	Return only host groups that contain the given hosts.
maintenanceids	string/array	Return only host groups that are affected by the given maintenances.
monitored_hosts	flag	Return only host groups that contain monitored hosts.
not_proxy_hosts	flag	Return only host groups that do not contain proxies.

Parameter	Type	Description
real_hosts	flag	Return only host groups that contain hosts.
templated_hosts	flag	Return only host groups that contain templates.
templateids	string/array	Return only host groups that contain the given templates.
triggerids	string/array	Return only host groups that contain hosts or templates with the given triggers.
with_applications	flag	Return only host groups that contain hosts with applications.
with_graphs	flag	Return only host groups that contain hosts with graphs.
with_historical_items	flag	Return only host groups that contain hosts with enabled or not supported items with historical data.
with_hosts_and_templates	flag	Overrides the with_simple_graph_items parameter. Return only host groups that contain hosts or templates.
with_httptests	flag	Return only host groups that contain hosts with web checks.
with_items	flag	Overrides the with_monitored_httptests parameter. Return only host groups that contain hosts or templates with items.
with_monitored_httptests	flag	Overrides the with_monitored_items, with_historical_items and with_simple_graph_items parameters. Return only host groups that contain hosts with enabled web checks.
with_monitored_items	flag	Return only host groups that contain hosts or templates with enabled items.
with_monitored_triggers	flag	Overrides the with_historical_items and with_simple_graph_items parameters. Return only host groups that contain hosts with enabled triggers. All of the items used in the trigger must also be enabled.
with_simple_graph_items	flag	Return only host groups that contain hosts with numeric items.
with_triggers	flag	Return only host groups that contain hosts with triggers.
selectHosts	query	Overrides the with_monitored_triggers parameter. Return the hosts that belong to the host group in the hosts property.
selectTemplates	query	Supports count; does not support shorten. Return the templates that belong to the host group in the templates property.
limitSelects	integer	Supports count; does not support shorten. Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: selectHosts - results will be sorted by host; selectTemplates - results will be sorted by host. Sort the result by the given properties. Possible values are: groupid, name.

Parameter	Type	Description
countOutput	flag	These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving data by name

Retrieve all data about two host groups named "Zabbix servers" and "Linux servers".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.get",
  "params": {
    "output": "extend",
    "filter": {
      "name": [
        "Zabbix servers",
        "Linux servers"
      ]
    }
  },
  "auth": "6f38cddc44cfbb6c1bd186f9a220b5a0",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "groupid": "2",
      "name": "Linux servers",
      "internal": "0"
    },
    {
      "groupid": "4",
      "name": "Zabbix servers",
      "internal": "0"
    }
  ],
  "id": 1
}
```

See also

- [hostgroup.getobjects](#)
- [Host](#)
- [Template](#)

Source

`CHostGroup::get()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.getobjects

Description

`array hostgroup.getobjects(object filter)`

This method allows to retrieve host groups that match the given filter criteria.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
groupid	string/array	Host group IDs.
name	string/array	Names of the host groups.
node	string	Name of the node the host groups must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the host groups must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieve a host group by name

Retrieve a host group called "Zabbix servers."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.getobjects",
  "params": {
    "name": "Linux servers"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "groupid": "2",
      "name": "Linux servers",
      "internal": "0"
    }
  ],
  "id": 16
}
```

See also

- [hostgroup.get](#)

Source

`CHostGroup::getObject()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.isreadable

Description

`boolean hostgroup.isreadable(array hostGroupIds)`

This method checks if the given host groups are available for reading.

Parameters

(array) IDs of the host groups to check.

Return values

(boolean) Returns true if the given host groups are available for reading.

Examples

Check multiple host groups

Check if the two host groups are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.isreadable",
  "params": [
    "5",
    "7"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [hostgroup.exists](#)
- [hostgroup.iswritable](#)

Source

`CHostGroup::isReadable()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.iswritable

Description

`boolean hostgroup.iswritable(array hostGroupIds)`

This method checks if the given host groups are available for writing.

Parameters

(array) IDs of the host groups to check.

Return values

(boolean) Returns true if the given host groups are available for writing.

Examples

Check multiple host groups

Check if the two host groups are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.iswritable",
  "params": [
    "5",
    "7"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [hostgroup.isreadable](#)
- [hostgroup.exists](#)

Source

CHostGroup::isWritable() in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.massadd

Description

object hostgroup.massadd(object parameters)

This method allows to simultaneously add multiple related objects to all the given host groups.

Parameters

(object) Parameters containing the IDs of the host groups to update and the objects to add to all the host groups.

The method accepts the following parameters.

Parameter	Type	Description
groups (required)	object/array	Host groups to be updated. The host groups must have the <code>groupid</code> property defined.
hosts	object/array	Hosts to add to all host groups. The hosts must have the <code>hostid</code> property defined.
templates	object/array	Templates to add to all host groups. The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Adding hosts to host groups

Add two hosts to host groups with IDs 5 and 6.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.massadd",
  "params": {
    "groups": [
      {
        "groupid": "5"
      },
      {
        "groupid": "6"
      }
    ],
    "hosts": [
      {
        "hostid": "30050"
      },
      {
        "hostid": "30001"
      }
    ]
  },
  "auth": "f223adf833b2bf2ff38574a67bba6372",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "groupids": [
      "5",
      "6"
    ]
  },
  "id": 1
}
```

See also

- [Host](#)
- [Template](#)

Source

`CHostGroup::massAdd()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.massremove

Description

`object hostgroup.massremove(object parameters)`

This method allows to remove related objects from multiple host groups.

Parameters

(object) Parameters containing the IDs of the host groups to update and the objects that should be removed.

Parameter	Type	Description
groupids (required)	string/array	IDs of the host groups to be updated.
hostids	string/array	Hosts to remove from all host groups.

Parameter	Type	Description
templateids	string/array	Templates to remove from all host groups.

Return values

(object) Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Removing hosts from host groups

Remove two hosts from the given host groups.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.massremove",
  "params": {
    "groupids": [
      "5",
      "6"
    ],
    "hostids": [
      "30050",
      "30001"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "groupids": [
      "5",
      "6"
    ]
  },
  "id": 1
}
```

Source

`CHostGroup::massRemove()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.massupdate

Description

object `hostgroup.massupdate(object parameters)`

This method allows to simultaneously replace or remove related objects for multiple host groups.

Parameters

(object) Parameters containing the IDs of the host groups to update and the objects that should be updated.

Parameter	Type	Description
groups (required)	object/array	Host groups to be updated. The host groups must have the <code>groupid</code> property defined.

Parameter	Type	Description
hosts	object/array	Hosts to replace the current hosts on the given host groups.
templates	object/array	The hosts must have the <code>hostid</code> property defined. Templates to replace the current templates on the given host groups. The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Replacing hosts in a host group

Replace all hosts in the host group with ID.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.massupdate",
  "params": {
    "groups": [
      {
        "groupid": "6"
      }
    ],
    "hosts": [
      {
        "hostid": "30050"
      }
    ]
  },
  "auth": "f223adf833b2bf2ff38574a67bba6372",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "groupids": [
      "6",
    ]
  },
  "id": 1
}
```

See also

- [hostgroup.update](#)
- [hostgroup.massadd](#)
- [Host](#)
- [Template](#)

Source

`CHostGroup::massUpdate()` in *frontends/php/api/classes/CHostGroup.php*.

hostgroup.update

Description

`object hostgroup.update(object/array hostGroups)`

This method allows to update existing hosts groups.

Parameters

(object/array) **Host group properties** to be updated.

The `groupid` property must be defined for each host group, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Renaming a host group

Rename a host group to "Linux hosts."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostgroup.update",
  "params": {
    "groupid": "7",
    "name": "Linux hosts"
  },
  "auth": "700ca65537074ec963db7efabda78259",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "groupids": [
      "7"
    ]
  },
  "id": 1
}
```

Source

`CHostGroup::update()` in *frontends/php/api/classes/CHostGroup.php*.

Host interface

This class is designed to work with host interfaces.

Object references:

- **Host interface**

Available methods:

- **hostinterface.create** - creating new host interfaces
- **hostinterface.delete** - deleting host interfaces
- **hostinterface.exists** - checking if a host interface exists
- **hostinterface.get** - retrieving host interfaces
- **hostinterface.massadd** - adding host interfaces to hosts
- **hostinterface.massremove** - removing host interfaces from hosts
- **hostinterface.replacehostinterfaces** - replacing host interfaces on a host
- **hostinterface.update** - updating host interfaces

> Host interface object

The following objects are directly related to the `hostinterface` API.

Host interface

The host interface object has the following properties.

Attention:

Note that both IP and DNS are required. If you do not want to use DNS, set it to an empty string.

Property	Type	Description
interfaceid	string	(<i>readonly</i>) ID of the interface.
dns (required)	string	DNS name used by the interface.
hostid (required)	string	Can be empty if the connection is made via IP. ID of the host the interface belongs to.
ip (required)	string	IP address used by the interface.
main (required)	integer	Can be empty if the connection is made via DNS. Whether the interface is used as default on the host. Only one interface of some type can be set as default on a host.
port (required)	string	Possible values are: 0 - not default; 1 - default. Port number used by the interface. Can contain user macros.
type (required)	integer	Interface type. Possible values are: 1 - agent; 2 - SNMP; 3 - IPMI; 4 - JMX.
useip (required)	integer	Whether the connection should be made via IP. Possible values are: 0 - connect using host DNS name; 1 - connect using host IP address.

hostinterface.create

Description

`object hostinterface.create(object/array hostInterfaces)`

This method allows to create new host interfaces.

Parameters

(object/array) Host interfaces to create. The method accepts host interfaces with the **standard host interface properties**.

Return values

(object) Returns an object containing the IDs of the created host interfaces under the `interfaceids` property. The order of the returned IDs matches the order of the passed host interfaces.

Examples

Create a new interface

Create a secondary IP agent interface on host "30052."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostinterface.create",
  "params": {
    "hostid": "30052",
    "dns": "",
    "ip": "127.0.0.1",
    "main": 0,
    "port": "10050",
    "type": 1,
    "useip": 1
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "interfaceids": [
      "30062"
    ]
  },
  "id": 1
}
```

See also

- [hostinterface.massadd](#)
- [host.massadd](#)

Source

`CHostInterface::create()` in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.delete

Description

object `hostinterface.delete(array hostInterfaceIds)`

This method allows to delete host interfaces.

Parameters

(array) IDs of the host interfaces to delete.

Return values

(object) Returns an object containing the IDs of the deleted host interfaces under the `interfaceids` property.

Examples

Delete a host interface

Delete the host interface with ID 30062.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostinterface.delete",
  "params": [
    "30062"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "interfaceids": [
      "30062"
    ]
  },
  "id": 1
}
```

See also

- [hostinterface.massremove](#)
- [host.massremove](#)

Source

`CHostInterface::delete()` in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.exists

Description

`boolean hostinterface.exists(object filter)`

This method checks if at least one host interface that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
dns	string/array	DNS of the host interfaces.
hostid	string/array	IDs of the hosts that the host interfaces must belong to.
interfaceid	string/array	Host interface IDs.
ip	string/array	IPs of the host interfaces.
node	string	Name of the node the host interfaces must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the node the host interfaces must belong to.

Return values

(boolean) Returns true if at least one host interface that matches the given filter criteria exists.

Examples

Check interface on host

Check if a host interface with IP 127.0.0.1 exists on host 30037.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostinterface.exists",
  "params": {
    "hostid": "30037",
    "ip": "127.0.0.1"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
}
```

```

    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}

```

Source

CHostInterface::exists() in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.get

Description

integer/array hostinterface.get(object parameters)

The method allows to retrieve host interfaces according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
hostids	string/array	Return only host interfaces used by the given hosts.
interfaceids	string/array	Return only host interfaces with the given IDs.
itemids	string/array	Return only host interfaces used by the given items.
triggerids	string/array	Return only host interfaces used by items in the given triggers.
selectItems	query	Return the items that use the interface in the <code>items</code> property.
selectHosts	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return the host that uses the interface as an array in the <code>hosts</code> property.
limitSelects	integer	Does not support <code>shorten</code> . Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: <code>selectItems</code> . Sort the result by the given properties.
countOutput	flag	Possible values are: <code>interfaceid</code> , <code>dns</code> , <code>ip</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve host interfaces

Retrieve all data about the interfaces used by host "30057."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostinterface.get",
  "params": {
    "output": "extend",
    "hostids": "30057"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "interfaceid": "30050",
      "hostid": "30057",
      "main": "1",
      "type": "1",
      "useip": "1",
      "ip": "127.0.0.1",
      "dns": "",
      "port": "10050"
    },
    {
      "interfaceid": "30067",
      "hostid": "30057",
      "main": "0",
      "type": "1",
      "useip": "0",
      "ip": "",
      "dns": "localhost",
      "port": "10050"
    },
    {
      "interfaceid": "30068",
      "hostid": "30057",
      "main": "1",
      "type": "2",
      "useip": "1",
      "ip": "127.0.0.1",
      "dns": "",
      "port": "161"
    }
  ],
  "id": 1
}
```

See also

- [Host](#)

- [Item](#)

Source

`CHostInterface::get()` in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.massadd

Description

`object hostinterface.massadd(object parameters)`

This method allows to simultaneously add host interfaces to multiple hosts.

Parameters

(object) Parameters containing the host interfaces to be created on the given hosts.

The method accepts the following parameters.

Parameter	Type	Description
hosts (required)	object/array	Hosts to be updated.
interfaces (required)	object/array	The hosts must have the <code>hostid</code> property defined. Host interfaces to create on the given hosts.

Return values

(object) Returns an object containing the IDs of the created host interfaces under the `interfaceids` property.

Examples

Creating interfaces

Create an interface on two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostinterface.massadd",
  "params": {
    "hosts": [
      {
        "hostid": "30050"
      },
      {
        "hostid": "30052"
      }
    ],
    "interfaces": {
      "dns": "",
      "ip": "127.0.0.1",
      "main": 0,
      "port": "10050",
      "type": 1,
      "useip": 1
    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "interfaceids": [
      "30069",
      "30070"
    ]
  },
  "id": 1
}
```

See also

- [hostinterface.create](#)
- [host.massadd](#)
- [Host](#)

Source

`CHostInterface::massAdd()` in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.massremove

Description

`object hostinterface.massremove(object parameters)`

This method allows to remove host interfaces from the given hosts.

Parameters

(object) Parameters containing the IDs of the hosts to be updated and the interfaces to be removed.

Parameter	Type	Description
hostids (required)	string/array	IDs of the hosts to be updated.
interfaces (required)	object/array	Host interfaces to remove from the given hosts. The host interface object must have the ip, dns and port properties defined

Return values

(object) Returns an object containing the IDs of the deleted host interfaces under the `interfaceids` property.

Examples

Removing interfaces

Remove the "127.0.0.1" SNMP interface from two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "hostinterface.massremove",
  "params": {
    "hostids": [
      "30050",
      "30052"
    ],
    "interfaces": {
      "dns": "",
      "ip": "127.0.0.1",
      "port": "161"
    }
  },
}
```

```

    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "interfaceids": [
      "30069",
      "30070"
    ]
  },
  "id": 1
}

```

See also

- [hostinterface.delete](#)
- [host.massremove](#)

Source

CHostInterface::massRemove() in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.replacehostinterfaces

Description

object `hostinterface.replacehostinterfaces(object parameters)`

This method allows to replace all host interfaces on a given host.

Parameters

(object) Parameters containing the ID of the host to be updated and the new host interfaces.

Parameter	Type	Description
hostid (required)	string	ID of the host to be updated.
interfaces (required)	object/array	Host interfaces to replace the current host interfaces with.

Return values

(object) Returns an object containing the IDs of the created host interfaces under the `interfaceids` property.

Examples

Replacing host interfaces

Replace all host interfaces with a single agent interface.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "hostinterface.replacehostinterfaces",
  "params": {
    "hostid": "30052",
    "interfaces": {
      "dns": "",
      "ip": "127.0.0.1",
      "main": 1,
      "port": "10050",
      "type": 1,
      "useip": 1
    }
  }
}

```

```

    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "interfaceids": [
      "30081"
    ]
  },
  "id": 1
}

```

See also

- [host.massupdate](#)

Source

`CHostInterface::replaceHostInterfaces()` in *frontends/php/api/classes/CHostInterface.php*.

hostinterface.update

Description

`object hostinterface.update(object/array hostInterfaces)`

This method allows to update existing host interfaces.

Parameters

(object/array) **Host interface properties** to be updated.

The `interfaceid` property must be defined for each host interface, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated host interfaces under the `interfaceids` property.

Examples

Changing a host interface port

Change the port of a host interface.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "hostinterface.update",
  "params": {
    "interfaceid": "30048",
    "port": "30050"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "interfaceids": [
      "30048"
    ]
  }
}

```

```

    ],
    },
    "id": 1
}

```

Source

CHostInterface::update() in *frontends/php/api/classes/CHostInterface.php*.

Icon map

This class is designed to work with icon maps.

Object references:

- [Icon map](#)
- [Icon mapping](#)

Available methods:

- [iconmap.create](#) - create new icon maps
- [iconmap.delete](#) - delete icon maps
- [iconmap.get](#) - retrieve icon maps
- [iconmap.isreadable](#) - check if an icon map is readable
- [iconmap.iswritable](#) - check if an icon map is writable
- [iconmap.update](#) - update icon maps

> Icon map object

The following objects are directly related to the `iconmap` API.

Icon map

The icon map object has the following properties.

Property	Type	Description
<code>iconmapid</code>	string	(<i>readonly</i>) ID of the icon map.
<code>default_iconid</code> (required)	string	ID of the default icon.
<code>name</code> (required)	string	Name of the icon map.

Icon mapping

The icon mapping object defines a specific icon to be used for hosts with a certain inventory field value. It has the following properties.

Property	Type	Description
<code>iconmappingid</code>	string	(<i>readonly</i>) ID of the icon map.
<code>iconid</code> (required)	string	ID of the icon used by the icon mapping.
<code>expression</code> (required)	string	Expression to match the inventory field against.
<code>inventory_link</code> (required)	integer	ID of the host inventory field. Refer to the host inventory object for a list of supported inventory fields.
<code>iconmapid</code>	string	(<i>readonly</i>) ID of the icon map that the icon mapping belongs to.

Property	Type	Description
sortorder	integer	Position of the icon mapping in the icon map.
		Default: 0.

iconmap.create

Description

object iconmap.create(object/array iconMaps)

This method allows to create new icon maps.

Parameters

(object/array) Icon maps to create.

Additionally to the **standard icon map properties**, the method accepts the following parameters.

Parameter	Type	Description
mappings (required)	array	Icon mappings to be created for the icon map.

Return values

(object) Returns an object containing the IDs of the created icon maps under the `iconmapids` property. The order of the returned IDs matches the order of the passed icon maps.

Examples

Create an icon map

Create an icon map to display hosts of different types.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "iconmap.create",
  "params": {
    "name": "Type icons",
    "default_iconid": "2",
    "mappings": [
      {
        "inventory_link": 1,
        "expression": "server",
        "iconid": "3"
      },
      {
        "inventory_link": 1,
        "expression": "switch",
        "iconid": "4"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "iconmapids": [
      "2"
    ]
  }
}
```

```

    ]
  },
  "id": 1
}

```

See also

- [Icon mapping](#)

Source

ClconMap::create() in *frontends/php/api/classes/ClconMap.php*.

iconmap.delete

Description

object iconmap.delete(array iconMapIds)

This method allows to delete icon maps.

Parameters

(array) IDs of the icon maps to delete.

Return values

(object) Returns an object containing the IDs of the deleted icon maps under the `iconmapids` property.

Examples

Delete multiple icon maps

Delete two icon maps.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "iconmap.delete",
  "params": [
    "2",
    "5"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "iconmapids": [
      "2",
      "5"
    ]
  },
  "id": 1
}

```

Source

ClconMap::delete() in *frontends/php/api/classes/ClconMap.php*.

iconmap.get

Description

integer/array iconmap.get(object parameters)

The method allows to retrieve icon maps according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
iconmapids	string/array	Return only icon maps with the given IDs.
sysmapids	string/array	Return only icon maps that are used in the given maps.
selectMappings	query	Return an object containing the used icon mappings in the mappings property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: iconmapid and name. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve an icon map

Retrieve all data about icon map "3".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "iconmap.get",
  "params": {
    "iconmapids": "3",
    "output": "extend",
    "selectMappings": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
```



```

{
    "mappings": {
        "3": {
            "iconmappingid": "3",
            "iconmapid": "3",
            "iconid": "6",
            "inventory_link": "1",
            "expression": "server",
            "sortorder": "0"
        },
        "4": {
            "iconmappingid": "4",
            "iconmapid": "3",
            "iconid": "10",
            "inventory_link": "1",
            "expression": "switch",
            "sortorder": "1"
        }
    },
    "iconmapid": "3",
    "name": "Host type icons",
    "default_iconid": "2"
}
],
"id": 1
}

```

See also

- [Icon mapping](#)

Source

ClconMap::get() in *frontends/php/api/classes/ClconMap.php*.

iconmap.isreadable

Description

boolean iconmap.isreadable(array iconMapIds)

This method checks if the given icon maps are available for reading.

Parameters

(array) IDs of the icon maps to check.

Return values

(boolean) Returns true if the given icon maps are available for reading.

Examples

Check multiple icon maps

Check if the two icon maps are readable.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "iconmap.isreadable",
    "params": [
        "4", "6"
    ],
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [iconmap.iswritable](#)

Source

ClconMap::isReadable() in *frontends/php/api/classes/ClconMap.php*.

iconmap.iswritable

Description

boolean iconmap.iswritable(array iconMapIds)

This method checks if the given icon maps are available for writing.

Parameters

(array) IDs of the icon maps to check.

Return values

(boolean) Returns true if the given icon maps are available for writing.

Examples

Check multiple icon maps

Check if the two icon maps are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "iconmap.iswritable",
  "params": [
    "4", "6"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [iconmap.isreadable](#)

Source

ClconMap::isWritable() in *frontends/php/api/classes/ClconMap.php*.

iconmap.update

Description

object iconmap.update(object/array iconMaps)

This method allows to update existing icon maps.

Parameters

(object/array) Icon map properties to be updated.

The `iconmapid` property must be defined for each icon map, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard icon map properties](#), the method accepts the following parameters.

Parameter	Type	Description
mappings	array	Icon mappings to replace the existing icon mappings.

Return values

(object) Returns an object containing the IDs of the updated icon maps under the `iconmapids` property.

Examples

Rename icon map

Rename an icon map to "OS icons".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "iconmap.update",
  "params": {
    "iconmapid": "1",
    "name": "OS icons"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "iconmapids": [
      "1"
    ]
  },
  "id": 1
}
```

See also

- [Icon mapping](#)

Source

`CIconMap::update()` in *frontends/php/api/classes/CIconMap.php*.

Image

This class is designed to work with images.

Object references:

- [Image](#)

Available methods:

- [image.create](#) - create new images
- [image.delete](#) - delete images
- [image.exists](#) - check if an image exists
- [image.get](#) - retrieve images
- [image.getobjects](#) - retrieve images by filters

- **image.update** - update images

> Image object

The following objects are directly related to the `image` API.

Image

The image object has the following properties.

Property	Type	Description
<code>imageid</code>	string	(<i>readonly</i>) ID of the image.
<code>name</code> (required)	string	Name of the image.
<code>imagetype</code>	integer	Type of image. Possible values: 1 - (<i>default</i>) icon; 2 - background image.

image.create

Description

`object image.create(object/array images)`

This method allows to create new images.

Parameters

(object/array) Images to create.

Additionally to the **standard image properties**, the method accepts the following parameters.

Parameter	Type	Description
<code>image</code> (required)	string	Base64 encoded image. The maximum size of the encoded image is 1 MB.

Return values

(object) Returns an object containing the IDs of the created images under the `imageids` property. The order of the returned IDs matches the order of the passed images.

Examples

Create an image

Create a cloud icon.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "image.create",
  "params": {
    "imagetype": 1,
    "name": "Cloud_(24)",
    "image": "iVBORwOKGgoAAAANSUhEUgAAABgAAAANCAYAAACzbK7QAAAABHNCSVQICAgIfAhkiAAAAAlwSFlzAAACmAAApgE
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "imageids": [
      "188"
    ]
  },
  "id": 1
}
```

Source

CImage::create() in *frontends/php/api/classes/CImage.php*.

image.delete

Description

object image.delete(array imageIds)

This method allows to delete images.

Parameters

(array) IDs of the images to delete.

Return values

(object) Returns an object containing the IDs of the deleted images under the `imageids` property.

Examples

Delete multiple images

Delete two images.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "image.delete",
  "params": [
    "188",
    "192"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "imageids": [
      "188",
      "192"
    ]
  },
  "id": 1
}
```

Source

CImage::delete() in *frontends/php/api/classes/CImage.php*.

image.exists

Description

`boolean image.exists(object filter)`

This method checks if at least one image that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the images must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the images must belong to.
imageid	string/array	IDs of images.
imagetype	integer/array	Types of images.
		Refer to the image "imagetype" property for a list of supported types.
name	string/array	Names of images.

Return values

(boolean) Returns true if at least one image that matches the given filter criteria exists.

Examples

Check image by name

Check if an image called "Cloud_(96)" exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "image.exists",
  "params": {
    "name": "Cloud_(96)"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

`CImage::exists()` in *frontends/php/api/classes/CImage.php*.

image.get

Description

`integer/array image.get(object parameters)`

The method allows to retrieve images according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
imageids	string/array	Return only images with the given IDs.
sysmapids	string/array	Return images that are used on the given maps.
select_image	flag	Return the Base64 encoded image in the <code>image</code> property.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: <code>imageid</code> and <code>name</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve an image

Retrieve all data for image with ID "2".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "image.get",
  "params": {
    "output": "extend",
    "select_image": true,
    "imageids": "2"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "imageid": "2",
      "imagetype": "1",
      "name": "Cloud_(24)",
      "image": "iVBORwOKGgoAAAANSUhEUgAAABgAAAANCAYAAACzbK7QAAAABHNCSVQICAgIfAhkiAAAAA1wSF1zAAACmAA"
    }
  ],
  "id": 1
}
```

See also

- [image.getobjects](#)

Source

CImage::get() in *frontends/php/api/classes/CImage.php*.

image.getobjects

Description

array image.getobjects(object filter)

This method allows to retrieve images that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard image properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the images must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the images must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieve image by name

Retrieve image called "Cloud_(24)".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "image.getobjects",
  "params": {
    "name": "Cloud_(24)"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "imageid": "2",
      "imagetype": "1",
      "name": "Cloud_(24)"
    }
  ],
  "id": 1
}
```

See also

- [image.get](#)

Source

CImage::getObject() in *frontends/php/api/classes/CImage.php*.

image.update

Description

object image.update(object/array images)

This method allows to update existing images.

Parameters

(object/array) Image properties to be updated.

The `imageid` property must be defined for each image, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the **standard image properties**, the method accepts the following parameters.

Parameter	Type	Description
image	string	Base64 encoded image. The maximum size of the encoded image is 1 MB.

Return values

(object) Returns an object containing the IDs of the updated images under the `imageids` property.

Examples

Rename image

Rename image to "Cloud icon".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "image.update",
  "params": {
    "imageid": "2",
    "name": "Cloud icon"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "imageids": [
      "2"
    ]
  },
  "id": 1
}
```

Source

CIImage::update() in *frontends/php/api/classes/CIImage.php*.

Item

This class is designed to work with items.

Object references:

- **Item**

Available methods:

- **item.create** - creating new items
- **item.delete** - deleting items
- **item.exists** - checking if items exists
- **item.get** - retrieving items
- **item.getobjects** - retrieving items by filters
- **item.isreadable** - checking if items are readable
- **item.iswritable** - checking if items are writable
- **item.update** - updating items

> Item object

The following objects are directly related to the `item` API.

Item

Note:

Web items cannot be directly created, updated or deleted via the Zabbix API.

The item object has the following properties.

Property	Type	Description
itemid	string	(<i>readonly</i>) ID of the item.
delay (required)	integer	Update interval of the item in seconds.
hostid (required)	string	ID of the host or template that the item belongs to.
interfaceid (required)	string	ID of the item's host interface.
		Not required for template items. Optional for Zabbix agent (active), Zabbix internal, Zabbix trapper, Zabbix aggregate, database monitor and calculated items.
key_ (required)	string	Item key.
name (required)	string	Name of the item.
type (required)	integer	Type of the item. Possible values: 0 - Zabbix agent; 1 - SNMPv1 agent; 2 - Zabbix trapper; 3 - simple check; 4 - SNMPv2 agent; 5 - Zabbix internal; 6 - SNMPv3 agent; 7 - Zabbix agent (active); 8 - Zabbix aggregate; 9 - web item; 10 - external check; 11 - database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 15 - calculated; 16 - JMX agent.

Property	Type	Description
value_type (required)	integer	Type of information of the item. Possible values: 0 - numeric float; 1 - character; 2 - log; 3 - numeric unsigned; 4 - text.
authtype	integer	SSH authentication method. Used only by SSH agent items. Possible values: 0 - <i>(default)</i> password; 1 - public key.
data_type	integer	Data type of the item. Possible values: 0 - <i>(default)</i> decimal; 1 - octal; 2 - hexadecimal; 3 - boolean.
delay_flex	string	Flexible intervals as a serialized string. Each serialized flexible interval consists of an update interval and a time period separated by a forward slash. Multiple intervals are separated by a colon.
delta	integer	Value that will be stored. Possible values: 0 - <i>(default)</i> as is; 1 - Delta, speed per second; 2 - Delta, simple change.
description	string	Description of the item.
error	string	<i>(readonly)</i> Error text if there are problems updating the item.
flags	integer	<i>(readonly)</i> Origin of the item. Possible values: 0 - a plain item; 4 - a discovered item.
formula	integer/float	Custom multiplier.
history	integer	Default: 1. Number of days to keep item's history data.
inventory_link	integer	Default: 90. ID of the host inventory field that is populated by the item. Refer to the host inventory page for a list of supported host inventory fields and their IDs.
ipmi_sensor	string	Default: 0. IPMI sensor. Used only by IPMI items.
lastclock	timestamp	<i>(readonly)</i> Time when the item was last updated.
lastns	integer	<i>(readonly)</i> Nanoseconds when the item was last updated.
lastvalue	string	<i>(readonly)</i> Last value of the item.
logtimefmt	string	Format of the time in log entries. Used only by log items.
mtime	timestamp	Time when the monitored log file was last updated. Used only by log items.

Property	Type	Description
multiplier	integer	Whether to use a custom multiplier.
params	string	Additional parameters depending on the type of the item: - executed script for SSH and telnet items; - additional parameters for database monitor items; - formula for calculated items.
password	string	Password for authentication. Used only by SSH, telnet and JMX items.
port	string	Port monitored by the item. Used only by SNMP items.
prevorgvalue	string	<i>(readonly)</i> Raw previous value of the item.
prevvalue	string	<i>(readonly)</i> Previous value of the item.
privatekey	string	Name of the private key file.
publickey	string	Name of the public key file.
snmp_community	string	SNMP community.
snmp_oid	string	SNMP OID.
snmpv3_authpassphrase	string	SNMPv3 auth passphrase. Used only by SNMPv3 items.
snmpv3_privpassphrase	string	SNMPv3 priv passphrase. Used only by SNMPv3 items.
snmpv3_securitylevel	integer	SNMPv3 security level. Used only by SNMPv3 items. Possible values: 0 - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv.
snmpv3_securityname	string	SNMPv3 security name. Used only by SNMPv3 items.
status	integer	Status of the item. Possible values: 0 - <i>(default)</i> enabled item; 1 - disabled item; 3 - unsupported item.
templateid	string	<i>(readonly)</i> ID of the parent template item.
trapper_hosts	string	<i>Hint:</i> Use the <code>hostid</code> property to specify the template that the item belongs to. Allowed hosts. Used only by trapper items.
trends	integer	Number of days to keep item's trends data.
units	string	Default: 365. Value units.
username	string	Username for authentication. Used only by SSH, telnet and JMX items.
valuemapid	string	Required by SSH and telnet items. ID of the associated value map.

item.create

Description

object `item.create(object/array items)`

This method allows to create new items.

Note:

Web items cannot be created via the Zabbix API.

Parameters

(object/array) Items to create.

Additionally to the **standard item properties**, the method accepts the following parameters.

Parameter	Type	Description
applications	array	IDs of the applications to add the item to.

Return values

(object) Returns an object containing the IDs of the created items under the `itemids` property. The order of the returned IDs matches the order of the passed items.

Examples

Creating an item

Create a numeric Zabbix agent item to monitor free disk space on host with ID "30074" and add it to two applications.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.create",
  "params": {
    "name": "Free disk space on $1",
    "key_": "vfs.fs.size[/home/joe/,free]",
    "hostid": "30074",
    "type": 0,
    "value_type": 3,
    "interfaceid": "30084",
    "applications": [
      "609",
      "610"
    ],
    "delay": 30
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "24758"
    ]
  },
  "id": 1
}
```

Creating a host inventory item

Create a Zabbix agent item to populate the host's "OS" inventory field.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.create",
  "params": {
    "name": "uname",
    "key_": "system.uname",
    "hostid": "30021",
    "type": 0,
    "interfaceid": "30007",
    "value_type": 1,
    "delay": 10,
    "inventory_link": 5
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

```
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "24759"
    ]
  },
  "id": 1
}
```

Source

CItem::create() in *frontends/php/api/classes/CItem.php*.

item.delete

Description

object item.delete(array itemIds)

This method allows to delete items.

Note:

Web items cannot be deleted via the Zabbix API.

Parameters

(array) IDs of the items to delete.

Return values

(object) Returns an object containing the IDs of the deleted items under the `itemids` property.

Examples

Deleting multiple items

Delete two items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.delete",
  "params": [
    "22982",
    "22986"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "22982",
      "22986"
    ]
  },
}
```

```
    "id": 1
}
```

Source

CItem::delete() in *frontends/php/api/classes/CItem.php*.

item.exists

Description

boolean item.exists(object filter)

This method checks if at least one item that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
key_ (required)	string/array	Keys of the items.
host	string/array	Names of the hosts that the items must belong to.
hostid	string/array	IDs of the hosts that the items must belong to.
node	string	Name of the node the items must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the items must belong to.

Return values

(boolean) Returns true if at least one item that matches the given filter criteria exists.

Examples

Check item by key

Check if an item with key "vm.memory.size[available]" exists on the host "Linux Server."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.exists",
  "params": {
    "host": "Linux Server",
    "key_": "vm.memory.size[available]"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [item.isreadable](#)
- [item.iswritable](#)

Source

CItem::exists() in *frontends/php/api/classes/CItem.php*.

item.get

Description

integer/array item.get(object parameters)

The method allows to retrieve items according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
itemids	string/array	Return only items with the given IDs.
groupids	string/array	Return only items that belong to the hosts from the given groups.
templateids	string/array	Return only items that belong to the given templates.
hostids	string/array	Return only items that belong to the given hosts.
proxyids	string/array	Return only items that are monitored by the given proxies.
interfaceids	string/array	Return only items that use the given host interfaces.
graphids	string/array	Return only items that are used in the given graphs.
triggerids	string/array	Return only items that are used in the given triggers.
applicationids	string/array	Return only items that belong to the given applications.
webitems	flag	Include web items in the result.
inherited	boolean	If set to true return only items inherited from a template.
templated	boolean	If set to true return only items that belong to templates.
monitored	boolean	If set to true return only enabled items that belong to monitored hosts.
group	string	Return only items that belong to a group with the given name.
host	string	Return only items that belong to a host with the given name.
application	string	Return only items that belong to an application with the given name.
with_triggers	boolean	If set to true return only items that are used in triggers.
selectHosts	query	Returns the host that the item belongs to as an array in the <code>hosts</code> property.
selectInterfaces	query	Does not support <code>shorten</code> . Returns the host interface used by the item as an array in the <code>interfaces</code> property.
selectTriggers	query	Does not support <code>shorten</code> . Return triggers that the item is used in in the <code>triggers</code> property.
selectGraphs	query	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Return graphs that contain the item in the <code>graphs</code> property.
selectApplications	query	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Return the applications that the item belongs to in the <code>applications</code> property.
		Does not support <code>shorten</code> or arrays of properties.

Parameter	Type	Description
selectDiscoveryRule	query	Return the LLD rule that created the item in the <code>discoveryRule</code> property.
selectItemDiscovery	query	Does not support <code>shorten</code> . Return the item discovery object in the <code>itemDiscovery</code> property. The item discovery objects links the item to an item prototype and has the following properties: <code>itemdiscoveryid</code> - (<i>string</i>) ID of the item discovery; <code>itemid</code> - (<i>string</i>) ID of the discovered item; <code>parent_itemid</code> - (<i>string</i>) ID of the item prototype from which the item has been created; <code>key_</code> - (<i>string</i>) key of the item prototype; <code>lastcheck</code> - (<i>timestamp</i>) time when the item was last discovered; <code>ts_delete</code> - (<i>timestamp</i>) time when an item that is no longer discovered will be deleted.
filter	object	Return only those results that exactly match the given filter. Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Supports additional filters: <code>host</code> - technical name of the host that the item belongs to.
limitSelects	integer	Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: <code>selectGraphs</code> - results will be sorted by name; <code>selectTriggers</code> - results will be sorted by description. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>itemid</code> , <code>name</code> , <code>key_</code> , <code>delay</code> , <code>history</code> , <code>trends</code> , <code>type</code> and <code>status</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(*integer/array*) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Finding items by key

Retrieve all items from host with ID "10084" that have the word "system" in the key and sort them by name.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.get",
  "params": {
    "output": "extend",
    "hostids": "10084",
    "search": {
      "key_": "system"
    },
    "sortfield": "name"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "23298",
      "type": "0",
      "snmp_community": "",
      "snmp_oid": "",
      "hostid": "10084",
      "name": "Context switches per second",
      "key_": "system.cpu.switches",
      "delay": "60",
      "history": "7",
      "trends": "365",
      "lastvalue": "2552",
      "lastclock": "1351090998",
      "prevvalue": "2641",
      "status": "0",
      "value_type": "3",
      "trapper_hosts": "",
      "units": "sps",
      "multiplier": "0",
      "delta": "1",
      "prevorgvalue": "67773258",
      "snmpv3_securityname": "",
      "snmpv3_securitylevel": "0",
      "snmpv3_authpassphrase": "",
      "snmpv3_privpassphrase": "",
      "formula": "1",
      "error": "",
      "lastlogsize": "0",
      "logtimefmt": "",
      "templateid": "22680",
      "valuemapid": "0",
      "delay_flex": "",
      "params": "",
      "ipmi_sensor": "",
      "data_type": "0",
      "authtype": "0",
      "username": "",
      "password": "",
      "publickey": "",
      "privatekey": "",
      "mtime": "0",
    }
  ]
}
```

```

    "lastns": "564054253",
    "flags": "0",
    "filter": "",
    "interfaceid": "1",
    "port": "",
    "description": "",
    "inventory_link": "0",
    "lifetime": "0"
  },
  {
    "itemid": "23299",
    "type": "0",
    "snmp_community": "",
    "snmp_oid": "",
    "hostid": "10084",
    "name": "CPU $2 time",
    "key_": "system.cpu.util[,idle]",
    "delay": "60",
    "history": "7",
    "trends": "365",
    "lastvalue": "86.031879",
    "lastclock": "1351090999",
    "prevvalue": "85.306944",
    "status": "0",
    "value_type": "0",
    "trapper_hosts": "",
    "units": "%",
    "multiplier": "0",
    "delta": "0",
    "prevorgvalue": "0",
    "snmpv3_securityname": "",
    "snmpv3_securitylevel": "0",
    "snmpv3_authpassphrase": "",
    "snmpv3_privpassphrase": "",
    "formula": "1",
    "error": "",
    "lastlogsize": "0",
    "logtimefmt": "",
    "templateid": "17354",
    "valuemapid": "0",
    "delay_flex": "",
    "params": "",
    "ipmi_sensor": "",
    "data_type": "0",
    "authtype": "0",
    "username": "",
    "password": "",
    "publickey": "",
    "privatekey": "",
    "mtime": "0",
    "lastns": "564256864",
    "flags": "0",
    "filter": "",
    "interfaceid": "1",
    "port": "",
    "description": "The time the CPU has spent doing nothing.",
    "inventory_link": "0",
    "lifetime": "0"
  },
  {
    "itemid": "23300",
    "type": "0",

```

```

        "snmp_community": "",
        "snmp_oid": "",
        "hostid": "10084",
        "name": "CPU $2 time",
        "key_": "system.cpu.util[,interrupt]",
        "delay": "60",
        "history": "7",
        "trends": "365",
        "lastvalue": "0.008389",
        "lastclock": "1351091000",
        "prevvalue": "0.000000",
        "status": "0",
        "value_type": "0",
        "trapper_hosts": "",
        "units": "%",
        "multiplier": "0",
        "delta": "0",
        "prevorgvalue": "0",
        "snmpv3_securityname": "",
        "snmpv3_securitylevel": "0",
        "snmpv3_authpassphrase": "",
        "snmpv3_privpassphrase": "",
        "formula": "1",
        "error": "",
        "lastlogsize": "0",
        "logtimefmt": "",
        "templateid": "22671",
        "valuemapid": "0",
        "delay_flex": "",
        "params": "",
        "ipmi_sensor": "",
        "data_type": "0",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "mtime": "0",
        "lastns": "564661387",
        "flags": "0",
        "filter": "",
        "interfaceid": "1",
        "port": "",
        "description": "The amount of time the CPU has been servicing hardware interrupts.",
        "inventory_link": "0",
        "lifetime": "0"
    }
],
    "id": 1
}

```

See also

- [item.getobjects](#)
- [Application](#)
- [Discovery rule](#)
- [Graph](#)
- [Host](#)
- [Host interface](#)
- [Trigger](#)

Source

`CIitem::get()` in *frontends/php/api/classes/CIitem.php*.

item.getobjects

Description

array item.getobjects(object filter)

This method allows to retrieve items that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard **standard item properties** the following parameters are supported as search criteria.

Parameter	Type	Description
host	string/array	Technical name of the host that the item belongs to.
node	string	Name of the node the items must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the items must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving items from a host

Retrieve all items from the host "Zabbix server."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.getobjects",
  "params": {
    "host": "Zabbix server"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "23327",
      "type": "0",
      "snmp_community": "",
      "snmp_oid": "",
      "hostid": "10084",
      "name": "Host name of zabbix_agentd running",
      "key_": "agent.hostname",
      "delay": "3600",
      "history": "7",
      "trends": "365",
      "lastvalue": "trapper-host",
      "lastclock": "1351088927",
      "prevvalue": "0",
      "status": "0",
      "value_type": "1",
      "trapper_hosts": "",
      "units": "",
      "multiplier": "0",
    }
  ]
}
```

```

    "delta": "0",
    "prevorgvalue": "0",
    "snmpv3_securityname": "",
    "snmpv3_securitylevel": "0",
    "snmpv3_authpassphrase": "",
    "snmpv3_privpassphrase": "",
    "formula": "1",
    "error": "",
    "lastlogsize": "0",
    "logtimefmt": "",
    "templateid": "23319",
    "valuemapid": "0",
    "delay_flex": "",
    "params": "",
    "ipmi_sensor": "",
    "data_type": "0",
    "authtype": "0",
    "username": "",
    "password": "",
    "publickey": "",
    "privatekey": "",
    "mtime": "0",
    "lastns": "40510111",
    "flags": "0",
    "filter": "",
    "interfaceid": "1",
    "port": "",
    "description": "",
    "inventory_link": "0",
    "lifetime": "30"
  },
  {
    "itemid": "23287",
    "type": "0",
    "snmp_community": "",
    "snmp_oid": "",
    "hostid": "10084",
    "name": "Agent ping",
    "key_": "agent.ping",
    "delay": "60",
    "history": "7",
    "trends": "365",
    "lastvalue": "1",
    "lastclock": "1351090987",
    "prevvalue": "1",
    "status": "0",
    "value_type": "3",
    "trapper_hosts": "",
    "units": "",
    "multiplier": "0",
    "delta": "0",
    "prevorgvalue": "0",
    "snmpv3_securityname": "",
    "snmpv3_securitylevel": "0",
    "snmpv3_authpassphrase": "",
    "snmpv3_privpassphrase": "",
    "formula": "1",
    "error": "",
    "lastlogsize": "0",
    "logtimefmt": "",
    "templateid": "10020",
    "valuemapid": "10",

```

```

    "delay_flex": "",
    "params": "",
    "ipmi_sensor": "",
    "data_type": "0",
    "authtype": "0",
    "username": "",
    "password": "",
    "publickey": "",
    "privatekey": "",
    "mtime": "0",
    "lastns": "560794191",
    "flags": "0",
    "filter": "",
    "interfaceid": "1",
    "port": "",
    "description": "The agent always returns 1 for this item. It could be used in combination with",
    "inventory_link": "0",
    "lifetime": "0"
  },
  {
    "itemid": "23288",
    "type": "0",
    "snmp_community": "",
    "snmp_oid": "",
    "hostid": "10084",
    "name": "Version of zabbix_agent(d) running",
    "key_": "agent.version",
    "delay": "3600",
    "history": "7",
    "trends": "365",
    "lastvalue": "2.0.0",
    "lastclock": "1351088888",
    "prevvalue": "0",
    "status": "0",
    "value_type": "1",
    "trapper_hosts": "",
    "units": "",
    "multiplier": "0",
    "delta": "0",
    "prevorgvalue": "0",
    "snmpv3_securityname": "",
    "snmpv3_securitylevel": "0",
    "snmpv3_authpassphrase": "",
    "snmpv3_privpassphrase": "",
    "formula": "1",
    "error": "",
    "lastlogsize": "0",
    "logtimefmt": "",
    "templateid": "10059",
    "valuemapid": "0",
    "delay_flex": "",
    "params": "",
    "ipmi_sensor": "",
    "data_type": "0",
    "authtype": "0",
    "username": "",
    "password": "",
    "publickey": "",
    "privatekey": "",
    "mtime": "0",
    "lastns": "8826267",
    "flags": "0",

```

```

        "filter": "",
        "interfaceid": "1",
        "port": "",
        "description": "",
        "inventory_link": "0",
        "lifetime": "0"
    }
],
    "id": 1
}

```

See also

- [item.get](#)

Source

CItem::getObject() in *frontends/php/api/classes/CItem.php*.

item.isreadable

Description

boolean item.isreadable(array itemIds)

This method checks if the given items are available for reading.

Parameters

(array) IDs of the items to check.

Return values

(boolean) Returns true if the given items are available for reading.

Examples

Check multiple items

Check if the two items are readable.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "item.isreadable",
    "params": [
        "23298",
        "23323"
    ],
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}

```

See also

- [item.exists](#)
- [item.iswritable](#)

Source

CItem::isReadable() in *frontends/php/api/classes/CItem.php*.

item.iswritable

Description

`boolean item.iswritable(array itemIds)`

This method checks if the given items are available for writing.

Parameters

(array) IDs of the items to check.

Return values

(boolean) Returns true if the given items are available for writing.

Examples

Check multiple items

Check if the two items are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.iswritable",
  "params": [
    "23298",
    "23323"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [item.isreadable](#)
- [item.exists](#)

Source

`CItem::isWritable()` in *frontends/php/api/classes/CItem.php*.

item.update

Description

`object item.update(object/array items)`

This method allows to update existing items.

Note:

Web items cannot be updated via the Zabbix API.

Parameters

(object/array) Item properties to be updated.

The `itemid` property must be defined for each item, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard item properties](#), the method accepts the following parameters.

Parameter	Type	Description
applications	array	IDs of the applications to replace the current applications.

Return values

(object) Returns an object containing the IDs of the updated items under the `itemids` property.

Examples

Enabling an item

Enable an item, that is, set its status to "0".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "item.update",
  "params": {
    "itemid": "10092",
    "status": 0
  },
  "auth": "700ca65537074ec963db7efabda78259",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "10092"
    ]
  },
  "id": 1
}
```

Source

`CItem::update()` in *frontends/php/api/classes/CItem.php*.

Item prototype

This class is designed to work with item prototypes.

Object references:

- [Item prototype](#)

Available methods:

- [itemprototype.create](#) - creating new item prototypes
- [itemprototype.delete](#) - deleting item prototypes
- [itemprototype.exists](#) - checking if item prototypes exist
- [itemprototype.get](#) - retrieving item prototypes
- [itemprototype.isreadable](#) - checking if item prototypes are readable
- [itemprototype.iswritable](#) - checking if item prototypes are writable
- [itemprototype.update](#) - updating item prototypes

> Item prototype object

The following objects are directly related to the `itemprototype` API.

Item prototype

The item prototype object has the following properties.

Property	Type	Description
itemid	string	<i>(readonly)</i> ID of the item prototype.
delay (required)	integer	Update interval of the item prototype in seconds.
hostid (required)	string	ID of the host that the item prototype belongs to.
interfaceid (required)	string	ID of the item prototype's host interface. Used only for host item prototypes.
key_ (required)	string	Optional for Zabbix agent (active), Zabbix internal, Zabbix trapper, Zabbix aggregate, database monitor and calculated item prototypes. Item prototype key.
name (required)	string	Name of the item prototype.
type (required)	integer	Type of the item prototype. Possible values: 0 - Zabbix agent; 1 - SNMPv1 agent; 2 - Zabbix trapper; 3 - simple check; 4 - SNMPv2 agent; 5 - Zabbix internal; 6 - SNMPv3 agent; 7 - Zabbix agent (active); 8 - Zabbix aggregate; 10 - external check; 11 - database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 15 - calculated; 16 - JMX agent.
value_type (required)	integer	Type of information of the item prototype. Possible values: 0 - numeric float; 1 - character; 2 - log; 3 - numeric unsigned; 4 - text.
authtype	integer	SSH authentication method. Used only by SSH agent item prototypes. Possible values: 0 - <i>(default)</i> password; 1 - public key.
data_type	integer	Data type of the item prototype. Possible values: 0 - <i>(default)</i> decimal; 1 - octal; 2 - hexadecimal; 3 - boolean.

Property	Type	Description
delay_flex	string	Flexible intervals as a serialized string.
delta	integer	Each serialized flexible interval consists of an update interval and a time period separated by a forward slash. Multiple intervals are separated by a colon. Value that will be stored.
description	string	Possible values: 0 - <i>(default)</i> as is; 1 - Delta, speed per second; 2 - Delta, simple change. Description of the item prototype.
formula	integer/float	Custom multiplier.
history	integer	Default: 1. Number of days to keep item prototype's history data.
ipmi_sensor	string	Default: 90. IPMI sensor. Used only by IPMI item prototypes.
logtimefmt	string	Format of the time in log entries. Used only by log item prototypes.
multiplier	integer	Whether to use a custom multiplier.
params	string	Additional parameters depending on the type of the item prototype: - executed script for SSH and telnet item prototypes; - additional parameters for database monitor item prototypes; - formula for calculated item prototypes.
password	string	Password for authentication. Used only by SSH, telnet and JMX item prototypes.
port	string	Port monitored by the item prototype. Used only by SNMP items prototype.
privatekey	string	Name of the private key file.
publickey	string	Name of the public key file.
snmp_community	string	SNMP community.
snmp_oid	string	SNMP OID.
snmpv3_authpassphrase	string	SNMPv3 auth passphrase. Used only by SNMPv3 item prototypes.
snmpv3_privpassphrase	string	SNMPv3 priv passphrase. Used only by SNMPv3 item prototypes.
snmpv3_securitylevel	integer	SNMPv3 security level. Used only by SNMPv3 item prototypes.
snmpv3_securityname	string	Possible values: 0 - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv. SNMPv3 security name. Used only by SNMPv3 item prototypes.
status	integer	Status of the item prototype.
templateid	string	Possible values: 0 - <i>(default)</i> enabled item prototype; 1 - disabled item prototype; 3 - unsupported item prototype. (readonly) ID of the parent template item prototype.
trapper_hosts	string	Allowed hosts. Used only by trapper item prototypes.
trends	integer	Number of days to keep item prototype's trends data.
units	string	Default: 365. Value units.

Property	Type	Description
username	string	Username for authentication. Used by simple check, SSH, Telnet, database monitor and JMX item prototypes.
valuemapid	string	Required by SSH and Telnet item prototypes. ID of the associated value map.

itemprototype.create

Description

object itemprototype.create(object/array itemPrototypes)

This method allows to create new item prototypes.

Parameters

(object/array) Item prototype to create.

Additionally to the [standard item prototype properties](#), the method accepts the following parameters.

Parameter	Type	Description
ruleid (required)	string	ID of the LLD rule that the item belongs to.
applications	array	IDs of applications to be assigned to the discovered items.

Return values

(object) Returns an object containing the IDs of the created item prototypes under the `itemids` property. The order of the returned IDs matches the order of the passed item prototypes.

Examples

Creating an item prototype

Create an item prototype to monitor free disc space on a discovered file system. Discovered items should be numeric Zabbix agent items updated every 30 seconds.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "itemprototype.create",
  "params": {
    "name": "Free disk space on $1",
    "key_": "vfs.fs.size[#{FSNAME},free]",
    "hostid": "10197",
    "ruleid": "27665",
    "type": 0,
    "value_type": 3,
    "interfaceid": "112",
    "delay": 30
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "27666"
    ]
  }
}
```

```

    },
    "id": 1
}

```

Source

CItemPrototype::create() in *frontends/php/api/classes/CItemPrototype.php*.

itemprototype.delete

Description

object itemprototype.delete(array itemPrototypeIds)

This method allows to delete item prototypes.

Parameters

(array) IDs of the item prototypes to delete.

Return values

(object) Returns an object containing the IDs of the deleted item prototypes under the `prototypeids` property.

Examples

Deleting multiple item prototypes

Delete two item prototypes.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "itemprototype.delete",
  "params": [
    "27352",
    "27356"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "prototypeids": [
      "27352",
      "27356"
    ]
  },
  "id": 1
}

```

Source

CItemPrototype::delete() in *frontends/php/api/classes/CItemPrototype.php*.

itemprototype.exists

Description

boolean itemprototype.exists(object filter)

This method checks if at least one item prototype that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
key_ (required)	string/array	Keys of the item prototypes.
host	string/array	Names of the hosts that the item prototypes must belong to.
hostid	string/array	IDs of the hosts that the item prototypes must belong to.
node	string	Name of the node the item prototypes must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the item prototypes must belong to.

Return values

(boolean) Returns true if at least one item prototype that matches the given filter criteria exists.

Examples

Checking if an item prototype exists on a host

Check if item prototype with key "net.if.in[#{IFNAME}]" exists on host "Zabbix server".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "itemprototype.exists",
  "params": {
    "host": "Zabbix server",
    "key_": "net.if.in[#{IFNAME}]"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [itemprototype.isreadable](#)
- [itemprototype.iswritable](#)

Source

`CItemPrototype::exists()` in *frontends/php/api/classes/CItemPrototype.php*.

itemprototype.get

Description

integer/array `itemprototype.get(object parameters)`

The method allows to retrieve item prototypes according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
discoveryids	string/array	Return only item prototypes that belong to the given LLD rules.
graphids	string/array	Return only item prototypes that are used in the given graph prototypes.
hostids	string/array	Return only item prototypes that belong to the given hosts.
inherited	boolean	If set to <code>true</code> return only item prototypes inherited from a template.
itemids	string/array	Return only item prototypes with the given IDs.
monitored	boolean	If set to <code>true</code> return only enabled item prototypes that belong to monitored hosts.
templated	boolean	If set to <code>true</code> return only item prototypes that belong to templates.
templateids	string/array	Return only item prototypes that belong to the given templates.
triggerids	string/array	Return only item prototypes that are used in the given trigger prototypes.
selectApplications	query	Return applications that the item prototype belongs to in the <code>applications</code> property.
selectGraphs	query	Does not support <code>shorten</code> or arrays of properties. Return graph prototypes that the item prototype is used in in the <code>graphs</code> property.
selectHosts	query	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Returns the host that the item prototype belongs to as an array in the <code>hosts</code> property.
selectTriggers	query	Does not support <code>shorten</code> . Return trigger prototypes that the item prototype is used in in the <code>triggers</code> property.
filter	object	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Return only those results that exactly match the given filter.
limitSelects	integer	Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Supports additional filters: <code>host</code> - technical name of the host that the item prototype belongs to. Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: <code>selectGraphs</code> - results will be sorted by <code>name</code> ; <code>selectTriggers</code> - results will be sorted by <code>description</code> . Sort the result by the given properties.
countOutput	flag	Possible values are: <code>itemid</code> , <code>name</code> , <code>key_</code> , <code>delay</code> , <code>type</code> and <code>status</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	

Parameter	Type	Description
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving item prototypes from an LLD rule

Retrieve all item prototypes from an LLD rule.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "itemprototype.get",
  "params": {
    "output": "extend",
    "discoveryids": "27426"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "27427",
      "type": "0",
      "snmp_community": "",
      "snmp_oid": "",
      "hostid": "10202",
      "name": "Incoming network traffic on $1 23",
      "key_": "2net.if.in[#{IFNAME}]",
      "delay": "60",
      "history": "7",
      "trends": "365",
      "status": "0",
      "value_type": "3",
      "trapper_hosts": "",
      "units": "bps",
      "multiplier": "1",
      "delta": "1",
      "snmpv3_securityname": "",
      "snmpv3_securitylevel": "0",
      "snmpv3_authpassphrase": "",
      "snmpv3_privpassphrase": "",
      "formula": "8",
      "logtimefmt": "",
      "templateid": "23881",
      "valuemapid": "0",

```

```

        "delay_flex": "",
        "params": "",
        "ipmi_sensor": "",
        "data_type": "0",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "mtime": "0",
        "filter": "",
        "interfaceid": "119",
        "port": "",
        "description": ""
    },
    {
        "itemid": "27428",
        "type": "0",
        "snmp_community": "",
        "snmp_oid": "",
        "hostid": "10202",
        "name": "Incoming network traffic on $1",
        "key_": "net.if.in[#{IFNAME}]",
        "delay": "60",
        "history": "7",
        "trends": "365",
        "status": "0",
        "value_type": "3",
        "trapper_hosts": "",
        "units": "bps",
        "multiplier": "1",
        "delta": "1",
        "snmpv3_securityname": "",
        "snmpv3_securitylevel": "0",
        "snmpv3_authpassphrase": "",
        "snmpv3_privpassphrase": "",
        "formula": "8",
        "logtimefmt": "",
        "templateid": "22446",
        "valuemapid": "0",
        "delay_flex": "",
        "params": "",
        "ipmi_sensor": "",
        "data_type": "0",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "mtime": "0",
        "filter": "",
        "interfaceid": "119",
        "port": "",
        "description": ""
    }
],
    "id": 1
}

```

See also

- [Application](#)
- [Host](#)

- [Graph prototype](#)
- [Trigger prototype](#)

Source

CItemPrototype::get() in *frontends/php/api/classes/CItemPrototype.php*.

itemprototype.isreadable

Description

boolean itemprototype.isreadable(array itemPrototypeIds)

This method checks if the given item prototypes are available for reading.

Parameters

(array) IDs of the item prototypes to check.

Return values

(boolean) Returns true if the given item prototypes are available for reading.

Examples

Check multiple item prototypes

Check if the two item prototypes are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "itemprototype.isreadable",
  "params": [
    "27352",
    "27356"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [itemprototype.exists](#)
- [itemprototype.iswritable](#)

Source

CItemPrototype::isReadable() in *frontends/php/api/classes/CItemPrototype.php*.

itemprototype.iswritable

Description

boolean itemprototype.iswritable(array itemPrototypeIds)

This method checks if the given item prototypes are available for writing.

Parameters

(array) IDs of the item prototypes to check.

Return values

(boolean) Returns true if the given item prototypes are available for writing.

Examples

Check multiple item prototypes

Check if the two item prototypes are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "itemprototype.iswritable",
  "params": [
    "27352",
    "27356"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [itemprototype.isreadable](#)
- [itemprototype.exists](#)

Source

CItemPrototype::isWritable() in *frontends/php/api/classes/CItemPrototype.php*.

itemprototype.update

Description

object itemprototype.update(object/array itemPrototypes)

This method allows to update existing item prototypes.

Parameters

(object/array) Item prototype properties to be updated.

The `itemid` property must be defined for each item prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard item prototype properties](#), the method accepts the following parameters.

Parameter	Type	Description
applications	array	IDs of the applications to replace the current applications.

Return values

(object) Returns an object containing the IDs of the updated item prototypes under the `itemids` property.

Examples

Changing the interface of an item prototype

Change the host interface that will be used by discovered items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "itemprototype.update",
  "params": {
    "itemid": "27428",
    "interfaceid": "132"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "27428"
    ]
  },
  "id": 1
}
```

Source

CIItemPrototype::update() in *frontends/php/api/classes/CIItemPrototype.php*.

IT service

This class is designed to work with IT services.

Object references:

- [IT service](#)
- [Service time](#)
- [Service dependency](#)
- [Service alarm](#)

Available methods:

- [service.adddependencies](#) - adding dependencies between IT services
- [service.addtimes](#) - adding service times
- [service.create](#) - creating new IT services
- [service.delete](#) - deleting IT services
- [service.deletedependencies](#) - deleting dependencies between IT services
- [service.deletetimes](#) - deleting service times
- [service.get](#) - retrieving IT services
- [service.getsla](#) - retrieving availability information about IT services
- [service.isreadable](#) - checking if IT services are readable
- [service.iswritable](#) - checking if IT services are writable
- [service.update](#) - updating IT services

> IT Service object

The following objects are directly related to the `service` API.

IT Service

The IT service object has the following properties.

Property	Type	Description
serviceid	string	(<i>readonly</i>) ID of the IT service.

Property	Type	Description
algorithm (required)	integer	Algorithm used to calculate the state of the IT service. Possible values: 0 - do not calculate; 1 - problem, if at least one child has a problem; 2 - problem, if all children have problems.
name (required)	string	Name of the IT service.
showsla (required)	integer	Whether SLA should be calculated. Possible values: 0 - do not calculate; 1 - calculate.
sortorder (required)	integer	Position of the IT service used for sorting.
goodsla	float	Minimum acceptable SLA value. If the SLA drops lower, the IT service is considered to be in problem state.
status	integer	Default: 99.9. (<i>readonly</i>) Whether the IT service is in OK or problem state. If the IT service is in problem state, status is equal either to: - the priority of the linked trigger if it is set to 2, "Warning" or higher (priorities 0, "Not classified" and 1, "Information" are ignored); - the highest status of a child IT service in problem state.
triggerid	string	If the IT service is in OK state, status is equal to 0. Trigger associated with the IT service. Can only be set for IT services that don't have children. Default: 0

Service time

The service time object defines periods, when an IT service is scheduled to be up or down. It has the following properties.

Property	Type	Description
timeid	string	(<i>readonly</i>) ID of the service time.
serviceid (required)	string	ID of the IT service.
ts_from (required)	integer	Cannot be updated. Time when the service time comes into effect. For onetime downtimes ts_from must be set as a Unix timestamp, for other types - as a specific time in a week, in seconds, for example, 90000 for Tue, 2:00 AM.
ts_to (required)	integer	Time when the service time ends. For onetime uptimes ts_to must be set as a Unix timestamp, for other types - as a specific time in a week, in seconds, for example, 90000 for Tue, 2:00 AM.
type (required)	integer	Service time type. Possible values: 0 - planned uptime, repeated every week; 1 - planned downtime, repeated every week; 2 - one-time downtime.

Property	Type	Description
note	string	Additional information about the service time.

Service dependency

The service dependency object represents a dependency between IT services. It has the following properties.

Property	Type	Description
linkid	string	<i>(readonly)</i> ID of the service dependency.
servicedownid (required)	string	ID of the IT service, that a service depends on, that is, the child service. An IT service can have multiple children.
serviceupid (required)	string	ID of the IT service, that is dependent on a service, that is, the parent service. An IT service can have multiple parents forming a directed graph.
soft (required)	integer	Type of dependency between IT services. Possible values: 0 - hard dependency; 1 - soft dependency. An IT service can have only one hard-dependent parent. This attribute has no effect on status or SLA calculation and is only used to create a core IT service tree. Additional parents can be added as soft dependencies forming a graph. An IT service can not be deleted if it has hard-dependent children.

Service alarm

Note:

Service alarms cannot be directly created, updated or deleted via the Zabbix API.

The service alarm objects represents an IT service's state change. It has the following properties.

Property	Type	Description
servicealarmid	string	ID of the service alarm.
serviceid	string	ID of the IT service.
clock	timestamp	Time when the IT service state change has happened.
value	integer	Status of the IT service. Refer the the IT service status property for a list of possible values.

service.adddependencies

Description

```
object service.adddependencies(object/array serviceDependencies)
```

This method allows to create dependencies between IT services.

Parameters

(object/array) Service dependencies to create.

Each service dependency has the following parameters.

Parameter	Type	Description
serviceid	string	ID of the IT service that depends on a service, that is, the parent service.
dependsOnServiceid	string	ID of the IT service that a service depends on, that is, the child service.
soft	string	Type of dependency.
Refer to the service dependency object page for more information on dependency types.		

Return values

(object) Returns an object containing the IDs of the affected parent IT services under the `serviceids` property.

Examples

Creating a hard dependency

Make IT service "2" a hard-dependent child of service "3".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.adddependencies",
  "params": {
    "serviceid": "3",
    "dependsOnServiceid": "2",
    "soft": 0
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "serviceids": [
      "3"
    ]
  },
  "id": 1
}
```

See also

- [service.update](#)

Source

CService::addDependencies() in *frontends/php/api/classes/CService.php*.

service.addtimes

Description

object `service.addtimes(object/array serviceTimes)`

This method allows to create new service times.

Parameters

(object/array) Service times to create.

The method accepts service times with the [standard service time properties](#).

Return values

(object) Returns an object containing the IDs of the affected IT services under the `serviceids` property.

Examples

Adding a scheduled downtime

Add a downtime for IT service "2" scheduled weekly from Monday 22:00 till Tuesday 10:00.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.addtimes",
  "params": {
    "serviceid": "4",
    "type": 1,
    "ts_from": 165600,
    "ts_to": 201600
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "serviceids": [
      "4"
    ]
  },
  "id": 1
}
```

See also

- [service.update](#)

Source

CService::addTimes() in *frontends/php/api/classes/CService.php*.

service.create

Description

object `service.create(object/array $itServices)`

This method allows to create new IT services.

Parameters

(object/array) IT services to create.

Additionally to the [standard IT service properties](#), the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Service dependencies. Each service dependency has the following parameters: - dependsOnServiceid - (<i>string</i>) ID of an IT service the service depends on, that is, the child IT service. - soft - (<i>integer</i>) type of service dependency; refer to the service dependency object page for more information on dependency types.
parentid	string	ID of a hard-linked parent IT service.
times	array	Service times to be created for the IT service.

Return values

(object) Returns an object containing the IDs of the created IT services under the `serviceids` property. The order of the returned IDs matches the order of the passed IT services.

Examples

Creating an IT service

Create an IT service that will be switched to problem state, if at least one child has a problem. SLA calculation will be on and the minimum acceptable SLA is 99.99%.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.create",
  "params": {
    "name": "Server 1",
    "algorithm": 1,
    "showsla": 1,
    "goodsla": 99.99,
    "sortorder": 1
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "serviceids": [
      "5"
    ]
  },
  "id": 1
}
```

Source

CService::create() in *frontends/php/api/classes/CService.php*.

service.delete

Description

object service.delete(array itServiceIds)

This method allows to delete IT services.

IT services with hard-dependent child services cannot be deleted.

Parameters

(array) IDs of the IT services to delete.

Return values

(object) Returns an object containing the IDs of the deleted IT services under the `serviceids` property.

Examples

Deleting multiple IT services

Delete two IT services.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.delete",
```

```

    "params": [
        "4",
        "5"
    ],
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "serviceids": [
            "4",
            "5"
        ]
    },
    "id": 1
}

```

Source

CService::delete() in *frontends/php/api/classes/CService.php*.

service.deletedependencies

Description

object service.deletedependencies(string/array serviceIds)

This method allows to delete all dependencies from IT services.

Parameters

(string/array) IDs of the IT services to delete all dependencies from.

Return values

(object) Returns an object containing the IDs of the affected IT services under the `serviceids` property.

Examples

Deleting dependencies from an IT service

Delete all dependencies from IT service "2".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "service.deletedependencies",
    "params": [
        "2"
    ],
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "serviceids": [
            "2"
        ]
    },
    "id": 1
}

```

See also

- [service.update](#)

Source

CService::delete() in *frontends/php/api/classes/CService.php*.

service.deletetimes

Description

object service.deletetimes(string/array serviceIds)

This method allows to delete all service times from IT services.

Parameters

(string/array) IDs of the IT services to delete all service times from.

Return values

(object) Returns an object containing the IDs of the affected IT services under the `serviceids` property.

Examples

Deleting service times from an IT service

Delete all service times from IT service "2".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.deletetimes",
  "params": [
    "2"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "serviceids": [
      "2"
    ]
  },
  "id": 1
}
```

See also

- [service.update](#)

Source

CService::delete() in *frontends/php/api/classes/CService.php*.

service.get

Description

integer/array service.get(object parameters)

The method allows to retrieve IT services according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
serviceids	string/array	Return only IT services with the given IDs.
parentids	string/array	Return only IT services with the given hard-dependent parent IT services.
childids	string/array	Return only IT services that are hard-dependent on the given child IT services.
selectParent	query	Return the hard-dependent parent IT service in the <code>parent</code> property.
selectDependencies	query	Return child service dependencies in the <code>dependencies</code> property.
selectParentDependencies	query	Return parent service dependencies in the <code>parentDependencies</code> property.
selectTimes	query	Return service times in the <code>times</code> property.
selectAlarms	query	Return service alarms in the <code>alarms</code> property.
selectTrigger	query	Return the associated trigger in the <code>trigger</code> property.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: <code>name</code> and <code>sortorder</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving all IT services

Retrieve all data about all IT services and their dependencies.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.get",
  "params": {
    "output": "extend",
    "selectDependencies": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```

{
  "jsonrpc": "2.0",
  "result": [
    {
      "serviceid": "2",
      "name": "Server 1",
      "status": "0",
      "algorithm": "1",
      "triggerid": "0",
      "showsla": "1",
      "goodsla": "99.9000",
      "sortorder": "0",
      "dependencies": []
    },
    {
      "serviceid": "3",
      "name": "Data center 1",
      "status": "0",
      "algorithm": "1",
      "triggerid": "0",
      "showsla": "1",
      "goodsla": "99.9000",
      "sortorder": "0",
      "dependencies": [
        {
          "linkid": "11",
          "serviceupid": "3",
          "servicedownid": "2",
          "soft": "0",
          "sortorder": "0",
          "serviceid": "2"
        },
        {
          "linkid": "10",
          "serviceupid": "3",
          "servicedownid": "5",
          "soft": "0",
          "sortorder": "1",
          "serviceid": "5"
        }
      ]
    },
    {
      "serviceid": "5",
      "name": "Server 2",
      "status": "0",
      "algorithm": "1",
      "triggerid": "0",
      "showsla": "1",
      "goodsla": "99.9900",
      "sortorder": "1",
      "dependencies": []
    }
  ],
  "id": 1
}

```

Source

CService::get() in *frontends/php/api/classes/CService.php*.

service.getsla

Description

`object service.getsla(object parameters)`

This method allows to calculate availability information about IT services.

Parameters

(object) Parameters containing the IDs of the IT services and time intervals to calculate SLA.

Parameter	Type	Description
serviceids	string/array	IDs of IT services to return availability information for.
intervals	array	Time intervals to return service layer availability information about. Each time interval must have the following parameters: - <code>from</code> - (<i>timestamp</i>) interval start time; - <code>to</code> - (<i>timestamp</i>) interval end time.

Return values

(object) Returns the following availability information about each IT service under the corresponding service ID.

Property	Type	Description
status	integer	Current status of the IT service. Refer to the IT service object page for more information on service statuses.
problems	array	Triggers that are currently in problem state and are linked either to the IT service or one of its descendants.
sla	array	SLA data about each time period. Each SLA object has the following properties: - <code>from</code> - (<i>timestamp</i>) interval start time; - <code>to</code> - (<i>timestamp</i>) interval end time; - <code>sla</code> - (<i>float</i>) SLA for the given time interval; - <code>okTime</code> - (<i>integer</i>) time the service was in OK state, in seconds; - <code>problemTime</code> - (<i>integer</i>) time the service was in problem state, in seconds; - <code>downtimeTime</code> - (<i>integer</i>) time the service was in scheduled downtime, in seconds.

Examples

Retrieving availability information for an IT service

Retrieve availability information about a service during a week.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.getsla",
  "params": {
    "serviceids": "2",
    "intervals": [
      {
        "from": 1352452201,
        "to": 1353057001
      }
    ]
  }
},
```

```

    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "2": {
      "status": "3",
      "problems": {
        "13904": {
          "triggerid": "13904",
          "expression": "{13359}=0",
          "description": "Service unavailable",
          "url": "",
          "status": "0",
          "value": "1",
          "priority": "3",
          "lastchange": "1352967420",
          "comments": "",
          "error": "",
          "templateid": "0",
          "type": "0",
          "value_flags": "0",
          "flags": "0"
        }
      },
      "sla": [
        {
          "from": 1352452201,
          "to": 1353057001,
          "sla": 97.046296296296,
          "okTime": 586936,
          "problemTime": 17864,
          "downtimeTime": 0
        }
      ]
    }
  },
  "id": 1
}

```

See also

- [Trigger](#)

Source

CService::getSla() in *frontends/php/api/classes/CService.php*.

service.isreadable

Description

boolean service.isreadable(array serviceIds)

This method checks if the given IT services are available for reading.

Parameters

(array) IDs of the IT services to check.

Return values

(boolean) Returns true if the given IT services are available for reading.

Examples

Check multiple IT services

Check if the two IT services are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.isreadable",
  "params": [
    "3", "4"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [service.iswritable](#)

Source

CService::isReadable() in *frontends/php/api/classes/CService.php*.

service.iswritable

Description

`boolean service.iswritable(array serviceIds)`

This method checks if the given IT services are available for writing.

Parameters

(array) IDs of the IT services to check.

Return values

(boolean) Returns true if the given IT services are available for writing.

Examples

Check multiple IT services

Check if the two IT services are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "service.iswritable",
  "params": [
    "3", "4"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
}
```

```
}
    "id": 1
}
```

See also

- [service.isreadable](#)

Source

CService::isWritable() in *frontends/php/api/classes/CService.php*.

service.update

Description

`object service.update(object/array itServices)`

This method allows to update existing IT services.

Parameters

(object/array) IT service properties to be updated.

The `serviceid` property must be defined for each IT service, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard IT service properties](#), the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Service dependencies to replace the current service dependencies. Each service dependency has the following parameters: - <code>dependsOnServiceid</code> - (<i>string</i>) ID of an IT service the service depends on, that is, the child IT service. - <code>soft</code> - (<i>integer</i>) type of service dependency; refer to the service dependency object page for more information on dependency types.
parentid	string	ID of a hard-linked parent IT service.
times	array	Service times to replace the current service times.

Return values

(object) Returns an object containing the IDs of the updated IT services under the `serviceids` property.

Examples

Setting the parent of an IT service

Make IT service "3" the hard-linked parent of service "5".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "service.update",
    "params": {
        "serviceid": "5",
        "parentid": "3"
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "serviceids": [
      "5"
    ]
  },
  "id": 1
}
```

See also

- [service.adddependencies](#)
- [service.addtimes](#)
- [service.deletedependencies](#)
- [service.deletetimes](#)

Source

CService::update() in *frontends/php/api/classes/CService.php*.

LLD rule

This class is designed to work with low level discovery rules.

Object references:

- [LLD rule](#)

Available methods:

- [discoveryrule.copy](#) - copying LLD rules
- [discoveryrule.create](#) - creating new LLD rules
- [discoveryrule.delete](#) - deleting LLD rules
- [discoveryrule.exists](#) - checking if LLD rules exist
- [discoveryrule.get](#) - retrieving LLD rules
- [discoveryrule.isreadable](#) - checking if LLD rules are readable
- [discoveryrule.iswritable](#) - checking if LLD rules are writable
- [discoveryrule.update](#) - updating LLD rules

> LLD rule object

The following objects are directly related to the `discoveryrule` API.

LLD rule

The low-level discovery rule object has the following properties.

Property	Type	Description
<code>itemid</code>	string	<i>(readonly)</i> ID of the LLD rule.
<code>delay</code> (required)	integer	Update interval of the LLD rule in seconds.
<code>hostid</code> (required)	string	ID of the host that the LLD rule belongs to.
<code>interfaceid</code> (required)	string	ID of the LLD rule's host interface. Used only for host LLD rules.
<code>key_</code> (required)	string	Optional for Zabbix agent (active), Zabbix internal, Zabbix trapper and database monitor LLD rules. LLD rule key.
<code>name</code> (required)	string	Name of the LLD rule.

Property	Type	Description
type (required)	integer	Type of the LLD rule. Possible values: 0 - Zabbix agent; 1 - SNMPv1 agent; 2 - Zabbix trapper; 3 - simple check; 4 - SNMPv2 agent; 5 - Zabbix internal; 6 - SNMPv3 agent; 7 - Zabbix agent (active); 8 - Zabbix aggregate; 10 - external check; 11 - database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 15 - calculated; 16 - JMX agent.
authtype	integer	SSH authentication method. Used only by SSH agent LLD rules. Possible values: 0 - (default) password; 1 - public key.
delay_flex	string	Flexible intervals as a serialized string. Each serialized flexible interval consists of an update interval and a time period separated by a forward slash. Multiple intervals are separated by a colon.
description	string	Description of the LLD rule.
error	string	(readonly) Error text if there are problems updating the LLD rule.
filter	string	LLD rule filter containing the macro to filter by and the regexp to be used for filtering separated by a colon. For example {#IFNAME}:@Network interfaces for discovery.
ipmi_sensor	string	IPMI sensor. Used only by IPMI LLD rules.
lifetime	integer	Time period after which items that are no longer discovered will be deleted, in days.
params	string	Default: 30. Additional parameters depending on the type of the LLD rule: - executed script for SSH and telnet LLD rules; - additional parameters for database monitor LLD rules; - formula for calculated LLD rules.
password	string	Password for authentication. Used only by SSH, telnet and JMX LLD rules.
port	string	Port used by the LLD rule. Used only by SNMP LLD rules.
privatekey	string	Name of the private key file.
publickey	string	Name of the public key file.
snmp_community	string	SNMP community.
snmp_oid	string	SNMP OID.
snmpv3_authpassphrase	string	SNMPv3 auth passphrase. Used only by SNMPv3 LLD rules.
snmpv3_privpassphrase	string	SNMPv3 priv passphrase. Used only by SNMPv3 LLD rules.

Property	Type	Description
snmpv3_securitylevel	integer	SNMPv3 security level. Used only by SNMPv3 LLD rules. Possible values: 0 - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv.
snmpv3_securityname	string	SNMPv3 security name. Used only by SNMPv3 LLD rules.
status	integer	Status of the LLD rule. Possible values: 0 - <i>(default)</i> enabled LLD rule; 1 - disabled LLD rule; 3 - unsupported LLD rule.
templateid	string	(readonly) ID of the parent template LLD rule.
trapper_hosts	string	Allowed hosts. Used only by trapper LLD rules.
username	string	Username for authentication. Used only by SSH, telnet and JMX LLD rules.
		Required by SSH and telnet LLD rules.

discoveryrule.copy

Description

`object discoveryrule.copy(object parameters)`

This method allows to copy LLD rules with all of the prototypes to the given hosts.

Parameters

(object) Parameters defining the LLD rules to copy and the target hosts.

Parameter	Type	Description
discoveryids	array	IDs of the LLD rules to be copied.
hostids	array	IDs of the hosts to copy the LLD rules to.

Return values

(boolean) Returns true if the copying was successful.

Examples

Copy an LLD rule to multiple hosts

Copy an LLD rule to two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.copy",
  "params": {
    "discoveryids": [
      "27426"
    ],
    "hostids": [
      "10196",
      "10197"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

CDiscoveryrule::copy() in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.create

Description

object discoveryrule.create(object/array lldRules)

This method allows to create new LLD rules.

Parameters

(object/array) LLD rules to create.

The method accepts LLD rules with the **standard LLD rule properties**.

Return values

(object) Returns an object containing the IDs of the created LLD rules under the `itemids` property. The order of the returned IDs matches the order of the passed LLD rules.

Examples

Creating an LLD rule

Create a Zabbix agent LLD rule to discover mounted file systems. Discovered items will be updated every 30 seconds.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.create",
  "params": {
    "name": "Mounted filesystem discovery",
    "key_": "vfs.fs.discovery",
    "hostid": "10197",
    "type": "0",
    "interfaceid": "112",
    "delay": 30
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "27665"
    ]
  },
  "id": 1
}
```

Source

CDiscoveryRule::create() in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.delete

Description

object `discoveryrule.delete(array lldRuleIds)`

This method allows to delete LLD rules.

Parameters

(array) IDs of the LLD rules to delete.

Return values

(object) Returns an object containing the IDs of the deleted LLD rules under the `itemids` property.

Examples

Deleting multiple LLD rules

Delete two LLD rules.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.delete",
  "params": [
    "27665",
    "27668"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "ruleids": [
      "27665",
      "27668"
    ]
  },
  "id": 1
}
```

Source

`CDiscoveryRule::delete()` in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.exists

Description

boolean `discoveryrule.exists(object filter)`

This method checks if at least one LLD rule that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
key_ (required)	string/array	Keys of the LLD rules.
host	string/array	Names of the hosts that the LLD rules must belong to.
hostid	string/array	IDs of the hosts that the LLD rules must belong to.

Parameter	Type	Description
node	string	Name of the node the LLD rules must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the LLD rules must belong to.

Return values

(boolean) Returns `true` if at least one LLD rule that matches the given filter criteria exists.

Examples

Checking if an LLD rule exists on a host

Check if the LLD rule with the key `"vfs.fs.discovery"` exists on host `"Zabbix server"`.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.exists",
  "params": {
    "host": "Zabbix server",
    "key_": "vfs.fs.discovery"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [discoveryrule.isreadable](#)
- [discoveryrule.iswritable](#)

Source

`CDiscoveryRule::exists()` in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.get

Description

`integer/array discoveryrule.get(object parameters)`

The method allows to retrieve LLD rules according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
itemids	string/array	Return only LLD rules with the given IDs.
hostids	string/array	Return only LLD rules that belong to the given hosts.
inherited	boolean	If set to <code>true</code> return only LLD rules inherited from a template.
interfaceids	string/array	Return only LLD rules use the given host interfaces.
monitored	boolean	If set to <code>true</code> return only enabled LLD rules that belong to monitored hosts.

Parameter	Type	Description
templated	boolean	If set to <code>true</code> return only LLD rules that belong to templates.
templateids	string/array	Return only LLD rules that belong to the given templates.
selectHosts	query	Returns the host that the LLD rule belongs to as an array in the <code>hosts</code> property.
selectGraphs	query	Does not support <code>shorten</code> . Returns graph prototypes that belong to the LLD rule in the <code>graphs</code> property.
selectItems	query	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Returns item prototypes that belong to the LLD rule in the <code>items</code> property.
selectTriggers	query	Supports <code>count</code> ; does not support <code>shorten</code> . Returns trigger prototypes that belong to the LLD rule in the <code>triggers</code> property.
filter	object	Supports <code>count</code> ; does not support <code>shorten</code> or arrays of properties. Return only those results that exactly match the given filter.
limitSelects	integer	Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Supports additional filters: <code>host</code> - technical name of the host that the LLD rule belongs to. Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: <code>selectItems</code> ; <code>selectGraphs</code> ; <code>selectTriggers</code> . Sort the result by the given properties.
countOutput	flag	Possible values are: <code>itemid</code> , <code>name</code> , <code>key_</code> , <code>delay</code> , <code>type</code> and <code>status</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(`integer/array`) Returns either:

- an array of objects;

- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving discovery rules from a host

Retrieve all discovery rules from host "10202".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.get",
  "params": {
    "output": "extend",
    "hostids": "10202"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "27425",
      "type": "0",
      "snmp_community": "",
      "snmp_oid": "",
      "hostid": "10202",
      "name": "Network interface discovery",
      "key_": "net.if.discovery",
      "delay": "3600",
      "status": "0",
      "trapper_hosts": "",
      "snmpv3_securityname": "",
      "snmpv3_securitylevel": "0",
      "snmpv3_authpassphrase": "",
      "snmpv3_privpassphrase": "",
      "error": "",
      "templateid": "22444",
      "delay_flex": "",
      "params": "",
      "ipmi_sensor": "",
      "authtype": "0",
      "username": "",
      "password": "",
      "publickey": "",
      "privatekey": "",
      "filter": "{#IFNAME}:@Network interfaces for discovery",
      "interfaceid": "119",
      "port": "",
      "description": "Discovery of network interfaces as defined in global regular expression \\"Netw",
      "lifetime": "30"
    },
    {
      "itemid": "27426",
      "type": "0",
      "snmp_community": "",
      "snmp_oid": "",
      "hostid": "10202",
      "name": "Mounted filesystem discovery",
      "key_": "vfs.fs.discovery",
      "delay": "3600",

```

```

        "status": "0",
        "trapper_hosts": "",
        "snmpv3_securityname": "",
        "snmpv3_securitylevel": "0",
        "snmpv3_authpassphrase": "",
        "snmpv3_privpassphrase": "",
        "error": "",
        "templateid": "22450",
        "delay_flex": "",
        "params": "",
        "ipmi_sensor": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "filter": "{#FSTYPE}:@File systems for discovery",
        "interfaceid": "119",
        "port": "",
        "description": "Discovery of file systems of different types as defined in global regular expressions",
        "lifetime": "30"
    }
],
    "id": 2
}

```

See also

- [Host](#)
- [Item prototype](#)
- [Graph prototype](#)
- [Trigger prototype](#)

Source

CDiscoveryRule::get() in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.isreadable

Description

boolean discoveryrule.isreadable(array lldRuleIds)

This method checks if the given LLD rules are available for reading.

Parameters

(array) IDs of the LLD rules to check.

Return values

(boolean) Returns true if the given LLD rules are available for reading.

Examples

Check multiple LLD rules

Check if the two LLD rules are readable.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "discoveryrule.isreadable",
    "params": [
        "27425",
        "27429"
    ],
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
}

```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [discoveryrule.exists](#)
- [discoveryrule.iswritable](#)

Source

CDiscoveryRule::isReadable() in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.iswritable

Description

boolean `discoveryrule.iswritable(array lldRuleIds)`

This method checks if the given LLD rules are available for writing.

Parameters

(array) IDs of the LLD rules to check.

Return values

(boolean) Returns true if the given LLD rules are available for writing.

Examples

Check multiple LLD rules

Check if the two LLD rules are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.iswritable",
  "params": [
    "27425",
    "27429"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [discoveryrule.isreadable](#)
- [discoveryrule.exists](#)

Source

CDiscoveryRule::isWritable() in *frontends/php/api/classes/CDiscoveryRule.php*.

discoveryrule.update

Description

object discoveryrule.update(object/array lldRules)

This method allows to update existing LLD rules.

Parameters

(object/array) **LLD rule properties** to be updated.

The itemid property must be defined for each LLD rule, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated LLD rules under the itemids property.

Examples

Adding a filter to an LLD rule

Add a filter so that the contents of the `{#FSTYPE}` macro would match the *@File systems for discovery* regexp.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "discoveryrule.update",
  "params": {
    "itemid": "22450",
    "filter": "{#FSTYPE}:@File systems for discovery"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
      "22450"
    ]
  },
  "id": 1
}
```

Source

CDiscoveryRule::update() in *frontends/php/api/classes/CDiscoveryRule.php*.

Maintenance

This class is designed to work with maintenances.

Object references:

- **Maintenance**
- **Time period**

Available methods:

- **maintenance.create** - creating new maintenances
- **maintenance.delete** - deleting maintenances
- **maintenance.exists** - checking if a maintenance exists
- **maintenance.get** - retrieving maintenances
- **maintenance.update** - updating maintenances

> Maintenance object

The following objects are directly related to the maintenance API.

Maintenance

The maintenance object has the following properties.

Property	Type	Description
maintenanceid	string	(<i>readonly</i>) ID of the maintenance.
name (required)	string	Name of the maintenance.
active_since	timestamp	Time when the maintenance becomes active.
active_till	timestamp	Default: current time. Time when the maintenance stops being active.
description	string	Default: the next day. Description of the maintenance.
maintenance_type	integer	Type of maintenance. Possible values: 0 - (<i>default</i>) with data collection; 1 - without data collection.

Time period

The time period object is used to define periods when the maintenance must come into effect. It has the following properties.

Property	Type	Description
timeperiodid	string	(<i>readonly</i>) ID of the maintenance.
day	integer	Day of the month when the maintenance must come into effect.
dayofweek	integer	Required only for monthly time periods. Days of the week when the maintenance must come into effect. Days are stored in binary form with each bit representing the corresponding day. For example, 4 equals 100 in binary and means, that maintenance will be enabled on Wednesday.
every	integer	Used for weekly and monthly time periods. Required only for weekly time periods. For daily and weekly periods every defines day or week intervals at which the maintenance must come into effect. For monthly periods every defines the week of the month when the maintenance must come into effect. Possible values: 1 - first week; 2 - second week; 3 - third week; 4 - fourth week; 5 - last week.

Property	Type	Description
month	integer	Months when the maintenance must come into effect. Months are stored in binary form with each bit representing the corresponding month. For example, 5 equals 101 in binary and means, that maintenance will be enabled in January and March.
period	integer	Required only for monthly time periods. Duration of the maintenance period in seconds.
start_date	timestamp	Default: 3600. Date when the maintenance period must come into effect. Required only for one time periods.
start_time	integer	Default: current date. Time of day when the maintenance starts in seconds.
timeperiod_type	integer	Required for daily, weekly and monthly periods. Type of time period. Possible values: 0 - (<i>default</i>) one time only; 2 - daily; 3 - weekly; 4 - monthly.

maintenance.create

Description

`object maintenance.create(object/array maintenances)`

This method allows to create new maintenances.

Parameters

(object/array) Maintenances to create.

Additionally to the **standard maintenance properties**, the method accepts the following parameters.

Parameter	Type	Description
groupids (required)	array	IDs of the host groups that will undergo maintenance.
hostids (required)	array	IDs of the hosts that will undergo maintenance.
timeperiods (required)	array	Maintenance time periods.

Attention:

At least one host or host group must be defined for each maintenance.

Return values

(object) Returns an object containing the IDs of the created maintenances under the `maintenanceids` property. The order of the returned IDs matches the order of the passed maintenances.

Examples

Creating a maintenance

Create a maintenance with data collection for host group "2". It must be active from 22.01.2013 till 22.01.2014, come in effect each Sunday at 18:00 and last for one hour.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "maintenance.create",
  "params": {
    "name": "Sunday maintenance",
    "active_since": 1358844540,
    "active_till": 1390466940,
    "groupids": [
      "2"
    ],
    "timeperiods": [
      {
        "timeperiod_type": 3,
        "every": 1,
        "dayofweek": 64,
        "start_time": 64800,
        "period": 3600
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "maintenanceids": [
      "3"
    ]
  },
  "id": 1
}
```

See also

- [Time period](#)

Source

CMaintenance::create() in *frontends/php/api/classes/CMaintenance.php*.

maintenance.delete

Description

object maintenance.delete(array maintenanceIds)

This method allows to delete maintenances.

Parameters

(array) IDs of the maintenances to delete.

Return values

(object) Returns an object containing the IDs of the deleted maintenances under the `maintenanceids` property.

Examples

Deleting multiple maintenances

Delete two maintenances.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "maintenance.delete",
  "params": [
    "3",
    "1"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "maintenanceids": [
      "3",
      "1"
    ]
  },
  "id": 1
}
```

Source

CMaintenance::delete() in *frontends/php/api/classes/CMaintenance.php*.

maintenance.exists

Description

boolean maintenance.exists(object filter)

This method checks if at least one maintenance that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
maintenanceid	string/array	IDs of the maintenances
name	string/array	Names of the maintenances

Return values

(boolean) Returns true if at least one maintenance that matches the given filter criteria exists.

Examples

Checking maintenance by name

Check if maintenance with the name "Sunday maintenance" already exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "maintenance.exists",
  "params": {
    "name": "Sunday maintenance"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

CMaintenance::exists() in *frontends/php/api/classes/CMaintenance.php*.

maintenance.get

Description

integer/array maintenance.get(object parameters)

The method allows to retrieve maintenances according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	string/array	Return only maintenances that are assigned to the given host groups.
hostids	string/array	Return only maintenances that are assigned to the given hosts.
maintenanceids	string/array	Return only maintenances with the given IDs.
selectGroups	query	Return host groups assigned to the maintenance in the groups property.
selectHosts	query	Does not support shorten. Return hosts assigned to the maintenance in the hosts property.
selectTimeperiods	query	Does not support shorten. (added in 2.0.3) Return the maintenance's time periods in the timeperiods property.
sortfield	string/array	Does not support shorten. Sort the result by the given properties.
countOutput	flag	Possible values are: maintenanceid, name and maintenance_type. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving maintenances

Retrieve all configured maintenances, and the data about the assigned host groups, hosts and defined time periods.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "maintenance.get",
  "params": {
    "output": "extend",
    "selectGroups": "extend",
    "selectTimeperiods": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "maintenanceid": "3",
      "name": "Sunday maintenance",
      "maintenance_type": "0",
      "description": "",
      "active_since": "1358844540",
      "active_till": "1390466940",
      "groups": [
        {
          "groupid": "4",
          "name": "Zabbix servers",
          "internal": "0"
        }
      ],
      "timeperiods": [
        {
          "timeperiodid": "4",
          "timeperiod_type": "3",
          "every": "1",
          "month": "0",
          "dayofweek": "1",
          "day": "0",
          "start_time": "64800",
          "period": "3600",
          "start_date": "2147483647"
        }
      ]
    }
  ],
  "id": 1
}
```

See also

- [Host](#)
- [Host group](#)
- [Time period](#)

Source

CMaintenance::get() in *frontends/php/api/classes/CMaintenance.php*.

maintenance.update

Description

object maintenance.update(object/array maintenances)

This method allows to update existing maintenances.

Parameters

(object/array) Maintenance properties to be updated.

The `maintenanceid` property must be defined for each maintenance, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Attention:

At this time, partial maintenance update is not supported, all parameters are mandatory. See [ZBX-6167](#) for current status.

Additionally to the **standard maintenance properties**, the method accepts the following parameters.

Parameter	Type	Description
groupids	array	IDs of the host groups to replace the current groups.
hostids	array	IDs of the hosts to replace the current hosts.
timeperiods	array	Maintenance time periods to replace the current periods.

Attention:

At least one host or host group must be defined for each maintenance.

Return values

(object) Returns an object containing the IDs of the updated maintenances under the `maintenanceids` property.

Examples

Assigning different hosts

Replace the hosts currently assigned to maintenance "3" with two different ones.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "maintenance.update",
  "params": {
    "maintenanceid": "3",
    "hostids": [
      "10085",
      "10084"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "maintenanceids": [
      "3"
    ]
  },
}
```

```
}  
  "id": 1  
}
```

See also

- [Time period](#)

Source

CMaintenance::update() in *frontends/php/api/classes/CMaintenance.php*.

Map

This class is designed to work with maps.

Object references:

- [Map](#)
- [Map element](#)
- [Map link](#)
- [Map URL](#)

Available methods:

- [map.create](#) - create new maps
- [map.delete](#) - delete maps
- [map.exists](#) - check if a map exists
- [map.get](#) - retrieve maps
- [map.getobjects](#) - retrieve maps by filters
- [map.isreadable](#) - check if maps are readable
- [map.iswritable](#) - check if maps are writable
- [map.update](#) - update maps

> Map object

The following objects are directly related to the map API.

Map

The map object has the following properties.

Property	Type	Description
sysmapid	string	(<i>readonly</i>) ID of the map.
height (required)	integer	Height of the map in pixels.
name (required)	string	Name of the map.
width (required)	integer	Width of the map in pixels.
backgroundid	string	ID of the image used as the background for the map.
expand_macros	integer	Whether to expand macros in labels when configuring the map. Possible values: 0 - (<i>default</i>) do not expand macros; 1 - expand macros.
expandproblem	integer	Whether the the problem trigger will be displayed for elements with a single problem. Possible values: 0 - always display the number of problems; 1 - (<i>default</i>) display the problem trigger if there's only one problem.

Property	Type	Description
grid_align	integer	Whether to enable grid aligning. Possible values: 0 - disable grid aligning; 1 - <i>(default)</i> enable grid aligning.
grid_show	integer	Whether to show the grid on the map. Possible values: 0 - do not show the grid; 1 - <i>(default)</i> show the grid.
grid_size	integer	Size of the map grid in pixels. Supported values: 20, 40, 50, 75 and 100.
highlight	integer	Default: 50. Whether icon highlighting is enabled. Possible values: 0 - highlighting disabled; 1 - <i>(default)</i> highlighting enabled.
iconmapid	string	ID of the icon map used on the map.
label_format	integer	Whether to enable advanced labels. Possible values: 0 - <i>(default)</i> disable advanced labels; 1 - enable advanced labels.
label_location	integer	Location of the map element label. Possible values: 0 - bottom; 1 - left; 2 - right; 3 - <i>(default)</i> top.
label_string_host	string	Custom label for host elements.
label_string_hostgroup	string	Required for maps with custom host label type. Custom label for host group elements.
label_string_image	string	Required for maps with custom host group label type. Custom label for image elements.
label_string_map	string	Required for maps with custom image label type. Custom label for map elements.
label_string_trigger	string	Required for maps with custom map label type. Custom label for trigger elements.
label_type	integer	Required for maps with custom trigger label type. Map element label type. Possible values: 0 - label; 1 - IP address; 2 - <i>(default)</i> element name; 3 - status only; 4 - nothing.

Property	Type	Description
label_type_host	integer	Label type for host elements. Possible values: 0 - label; 1 - IP address; 2 - <i>(default)</i> element name; 3 - status only; 4 - nothing; 5 - custom.
label_type_hostgroup	integer	Label type for host group elements. Possible values: 0 - label; 2 - <i>(default)</i> element name; 3 - status only; 4 - nothing; 5 - custom.
label_type_image	integer	Label type for host group elements. Possible values: 0 - label; 2 - <i>(default)</i> element name; 4 - nothing; 5 - custom.
label_type_map	integer	Label type for map elements. Possible values: 0 - label; 2 - <i>(default)</i> element name; 3 - status only; 4 - nothing; 5 - custom.
label_type_trigger	integer	Label type for trigger elements. Possible values: 0 - label; 2 - <i>(default)</i> element name; 3 - status only; 4 - nothing; 5 - custom.
markelements	integer	Whether to highlight map elements that have recently changed their status. Possible values: 0 - <i>(default)</i> do not highlight elements; 1 - highlight elements.
show_unack	integer	How problems should be displayed. Possible values: 0 - <i>(default)</i> display the count of all problems; 1 - display only the count of unacknowledged problems; 2 - display the count of acknowledged and unacknowledged problems separately.

Map element

The map element object defines an object displayed on a map. It has the following properties.

Property	Type	Description
selementid	string	<i>(readonly)</i> ID of the map element.

Property	Type	Description
elementid (required)	string	ID of the object that the map element represents.
elementtype (required)	integer	Required for host, host group, trigger and map type elements. Type of map element. Possible values: 0 - host; 1 - map; 2 - trigger; 3 - host group; 4 - image.
iconid_off (required)	string	ID of the image used to display the element in default state.
areatype	integer	How separate host group hosts should be displayed. Possible values: 0 - <i>(default)</i> the host group element will take up the whole map; 1 - the host group element will have a fixed size.
elementsubtype	integer	How a host group element should be displayed on a map. Possible values: 0 - <i>(default)</i> display the host group as a single element; 1 - display each host in the group separately.
height	integer	Height of the fixed size host group element in pixels. Default: 200.
iconid_disabled	string	ID of the image used to display disabled map elements. Unused for image elements.
iconid_maintenance	string	ID of the image used to display map elements in maintenance. Unused for image elements.
iconid_on	string	ID of the image used to display map elements with problems. Unused for image elements.
label	string	Label of the element.
label_location	integer	Location of the map element label. Possible values: 0 - <i>(default)</i> bottom; 1 - left; 2 - right; 3 - top.
sysmapid	string	<i>(readonly)</i> ID of the map that the element belongs to.
urls	array	Map element URLs.
use_iconmap	integer	The map element URL object is described in detail below . Whether icon mapping must be used for host elements. Possible values: 0 - do not use icon mapping; 1 - <i>(default)</i> use icon mapping.
viewtype	integer	Host group element placing algorithm. Possible values: 0 - <i>(default)</i> grid.
width	integer	Width of the fixed size host group element in pixels. Default: 200.
x	integer	X-coordinates of the element in pixels. Default: 0.

Property	Type	Description
y	integer	Y-coordinates of the element in pixels.
		Default: 0.

Map element URL

The map element URL object defines a clickable link that will be available for a specific map element. It has the following properties:

Property	Type	Description
sysmapelementurlid	string	(<i>readonly</i>) ID of the map element URL.
name (required)	string	Link caption.
url (required)	string	Link URL.
selementid	string	ID of the map element that the URL belongs to.

Map link

The map link object defines a link between two map elements. It has the following properties.

Property	Type	Description
linkid	string	(<i>readonly</i>) ID of the map link.
selementid1 (required)	string	ID of the first map element linked on one end.
selementid2 (required)	string	ID of the first map element linked on the other end.
color	string	Line color as a hexadecimal color code.
drawtype	integer	Default: 000000. Link line draw style. Possible values: 0 - (<i>default</i>) line; 2 - bold line; 3 - dotted line; 4 - dashed line.
label	string	Link label.
linktriggers	array	Map link triggers to use as link status indicators.
sysmapid	string	The map link trigger object is described in detail below . ID of the map the link belongs to.

Map link trigger

The map link trigger object defines a map link status indicator based on the state of a trigger. It has the following properties:

Property	Type	Description
linktriggerid	string	(<i>readonly</i>) ID of the map link trigger.
triggerid (required)	string	ID of the trigger used as a link indicator.
color	string	Indicator color as a hexadecimal color code.
		Default: DD0000.

Property	Type	Description
drawtype	integer	Indicator draw style. Possible values: 0 - <i>(default)</i> line; 2 - bold line; 3 - dotted line; 4 - dashed line.
linkid	string	ID of the map link that the link trigger belongs to.

Map URL

The map URL object defines a clickable link that will be available for all elements of a specific type on the map. It has the following properties:

Property	Type	Description
sysmapurlid	string	<i>(readonly)</i> ID of the map URL.
name (required)	string	Link caption.
url (required)	string	Link URL.
elementtype	integer	Type of map element for which the URL will be available. Refer to the map element "type" property for a list of supported types.
sysmapid	string	Default: 0. ID of the map that the URL belongs to.

map.create

Description

`object map.create(object/array maps)`

This method allows to create new maps.

Parameters

(object/array) Maps to create.

Additionally to the [standard map properties](#), the method accepts the following parameters.

Parameter	Type	Description
links	array	Map links to be created on the map.
selements	array	Map elements to be created on the map.
urls	array	Map URLs to be created on the map.

Note:

To create map links you'll need to set a map elements `selementid` to an arbitrary value and then use this value to reference this element in the links `selementid1` or `selementid2` properties. When the element is created, this value will be replaced with the correct ID generated by Zabbix. [See example.](#)

Return values

(object) Returns an object containing the IDs of the created maps under the `sysmapids` property. The order of the returned IDs matches the order of the passed maps.

Examples

Create an empty map

Create a map with no elements.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.create",
  "params": {
    "name": "Map",
    "width": 600,
    "height": 600
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "sysmapids": [
      "8"
    ]
  },
  "id": 1
}
```

Create a host map

Create a map with two host elements and a link between them. Note the use of temporary "selementid1" and "selementid2" values in the map link object to refer to map elements.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.create",
  "params": {
    "name": "Host map",
    "width": 600,
    "height": 600,
    "selements": [
      {
        "elementid": "1033",
        "selementid": "1",
        "elementtype": 0,
        "iconid_off": "2"
      },
      {
        "elementid": "1037",
        "selementid": "2",
        "elementtype": 0,
        "iconid_off": "2"
      }
    ],
    "links": [
      {
        "selementid1": "1",
        "selementid2": "2"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "sysmapids": [
      "9"
    ]
  },
  "id": 1
}
```

See also

- [Map element](#)
- [Map link](#)
- [Map URL](#)

Source

CMap::create() in *frontends/php/api/classes/CMap.php*.

map.delete

Description

object map.delete(array mapIds)

This method allows to delete maps.

Parameters

(array) IDs of the maps to delete.

Return values

(object) Returns an object containing the IDs of the deleted maps under the sysmapids property.

Examples

Delete multiple maps

Delete two maps.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.delete",
  "params": [
    "12",
    "34"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "sysmapids": [
      "12",
      "34"
    ]
  },
  "id": 1
}
```

Source

CMap::delete() in *frontends/php/api/classes/CMap.php*.

map.exists

Description

`boolean map.exists(object filter)`

This method checks if at least one map that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
name	string/array	Names of the maps.
node	string	Name of the node the maps must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the maps must belong to.
sysmapid	string/array	IDs of the maps.

Return values

(boolean) Returns true if at least one map that matches the given filter criteria exists.

Examples

Check a map by name

Check if map "Local network" exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.exists",
  "params": {
    "name": "Local network"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [map.isreadable](#)
- [map.iswritable](#)

Source

CMap::exists() in *frontends/php/api/classes/CMap.php*.

map.get

Description

`integer/array map.get(object parameters)`

The method allows to retrieve maps according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
sysmapids	string/array	Return only maps with the given IDs.
expandUrls	flag	Adds global map URLs to the corresponding map elements and expands macros in all map element URLs.
selectIconMap	query	Returns the icon map used on the map in the <code>iconmap</code> property.
selectLinks	query	Does not support <code>shorten</code> or arrays of properties. Returns map links between elements in the <code>links</code> property.
selectSelements	query	Does not support <code>shorten</code> or arrays of properties. Returns the map elements from the map in the <code>selements</code> property.
sortfield	string/array	Does not support <code>shorten</code> or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>name</code> , <code>width</code> and <code>height</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve a map

Retrieve all data about map "3".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.get",
  "params": {
    "output": "extend",
    "selectSelements": "extend",
    "selectLinks": "extend",
    "sysmapids": "3"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "selements": [
        {
          "selementid": "10",
          "sysmapid": "3",
          "elementid": "0",
          "elementtype": "4",
          "iconid_off": "1",
          "iconid_on": "0",
          "label": "Zabbix server",
          "label_location": "3",
          "x": "11",
          "y": "141",
          "iconid_disabled": "0",
          "iconid_maintenance": "0",
          "elementsubtype": "0",
          "areatype": "0",
          "width": "200",
          "height": "200",
          "viewtype": "0",
          "use_iconmap": "1",
          "urls": []
        },
        {
          "selementid": "11",
          "sysmapid": "3",
          "elementid": "0",
          "elementtype": "4",
          "iconid_off": "1",
          "iconid_on": "0",
          "label": "Web server",
          "label_location": "3",
          "x": "211",
          "y": "191",
          "iconid_disabled": "0",
          "iconid_maintenance": "0",
          "elementsubtype": "0",
          "areatype": "0",
          "width": "200",
          "height": "200",
          "viewtype": "0",
          "use_iconmap": "1",
          "urls": []
        }
      ],
      "links": [
        {
          "linkid": "23",
          "sysmapid": "3",
          "selementid1": "10",
          "selementid2": "11",
          "drawtype": "0",
          "color": "00CC00",
          "label": "",
          "linktriggers": []
        }
      ]
    }
  ]
}
```

```

    }
],
"urls": [],
"sysmapid": "3",
"name": "Local network",
"width": "400",
"height": "400",
"backgroundid": "0",
"label_type": "2",
"label_location": "3",
"highlight": "1",
"expandproblem": "1",
"markelements": "0",
"show_unack": "0",
"grid_size": "50",
"grid_show": "1",
"grid_align": "1",
"label_format": "0",
"label_type_host": "2",
"label_type_hostgroup": "2",
"label_type_trigger": "2",
"label_type_map": "2",
"label_type_image": "2",
"label_string_host": "",
"label_string_hostgroup": "",
"label_string_trigger": "",
"label_string_map": "",
"label_string_image": "",
"iconmapid": "0",
"expand_macros": "0"
}
],
"id": 1
}

```

See also

- [map.getobjects](#)
- [Icon map](#)
- [Map element](#)
- [Map link](#)
- [Map URL](#)

Source

CMap::get() in *frontends/php/api/classes/CMap.php*.

map.getobjects

Description

array map.getobjects(object filter)

This method allows to retrieve maps that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard map properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the maps must belong to.
		This will override the <code>nodeids</code> parameter.

Parameter	Type	Description
nodeids	string/array	IDs of the nodes the maps must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieve a map by name

Retrieve a map called "Local network".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.getobjects",
  "params": {
    "name": "Local network"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "urls": [],
      "sysmapid": "3",
      "name": "Local network",
      "width": "400",
      "height": "400",
      "backgroundid": "0",
      "label_type": "2",
      "label_location": "3",
      "highlight": "1",
      "expandproblem": "1",
      "markelements": "0",
      "show_unack": "0",
      "grid_size": "50",
      "grid_show": "1",
      "grid_align": "1",
      "label_format": "0",
      "label_type_host": "2",
      "label_type_hostgroup": "2",
      "label_type_trigger": "2",
      "label_type_map": "2",
      "label_type_image": "2",
      "label_string_host": "",
      "label_string_hostgroup": "",
      "label_string_trigger": "",
      "label_string_map": "",
      "label_string_image": "",
      "iconmapid": "0",
      "expand_macros": "0"
    }
  ],
  "id": 1
}
```

See also

- [map.get](#)

Source

CMap::getObject() in *frontends/php/api/classes/CMap.php*.

map.isreadable

Description

boolean map.isreadable(array sysmapIds)

This method checks if the given maps are available for reading.

Parameters

(array) IDs of the maps to check.

Return values

(boolean) Returns true if the given maps are available for reading.

Examples

Check multiple maps

Check if the two maps are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.isreadable",
  "params": [
    "32", "6"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [map.exists](#)
- [map.iswritable](#)

Source

CMap::isReadable() in *frontends/php/api/classes/CMap.php*.

map.iswritable

Description

boolean map.iswritable(array sysmapIds)

This method checks if the given maps are available for writing.

Parameters

(array) IDs of the maps to check.

Return values

(boolean) Returns true if the given maps are available for writing.

Examples

Check multiple maps

Check if the two maps are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.iswritable",
  "params": [
    "32", "7"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [map.isreadable](#)
- [map.exists](#)

Source

CMap::isWritable() in *frontends/php/api/classes/CMap.php*.

map.update

Description

object map.update(object/array maps)

This method allows to update existing maps.

Parameters

(object/array) Map properties to be updated.

The mapid property must be defined for each map, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard map properties](#), the method accepts the following parameters.

Parameter	Type	Description
links	array	Map links to replace the existing links.
selements	array	Map elements to replace the existing elements.
urls	array	Map URLs to replace the existing URLs.

Note:

To create map links between new map elements you'll need to set an elements selementid to an arbitrary value and then use this value to reference this element in the links selementid1 or selementid2 properties. When the element is created, this value will be replaced with the correct ID generated by Zabbix. [See example for map.create.](#)

Return values

(object) Returns an object containing the IDs of the updated maps under the sysmapids property.

Examples

Resize a map

Change the size of the map to 1200x1200 pixels.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "map.update",
  "params": {
    "sysmapid": "8",
    "width": 1200,
    "height": 1200
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "sysmapids": [
      "8"
    ]
  },
  "id": 1
}
```

See also

- [Map element](#)
- [Map link](#)
- [Map URL](#)

Source

CMap::update() in *frontends/php/api/classes/CMap.php*.

Media

This class is designed to work with media.

Object references:

- [Media](#)

Available methods:

- [usermedia.get](#) - retrieving media

Methods to configure media via the user API:

- [user.addmedia](#) - creating media
- [user.updatemedia](#) - updating media
- [user.deletemedia](#) - deleting media

> Media object

The following objects are directly related to the usermedia API.

Media

Note:

Media are created, updated and deleted via the the [user API](#).

The media object defines how a media type should be used for a user. It has the following properties.

Property	Type	Description
mediaid active (required)	string integer	<i>(readonly)</i> ID of the media. Whether the media is enabled. Possible values: 0 - enabled; 1 - disabled.
mediatypeid (required)	string	ID of the media type used by the media.
period (required)	string	Time when the notifications can be sent as a time period .
sendto (required)	string	Address, user name or other identifier of the recipient.
severity (required)	integer	Trigger severities to send notifications about. Severities are stored in binary form with each bit representing the corresponding severity. For example, 12 equals 1100 in binary and means, that notifications will be sent from triggers with severities warning and average. Refer to the trigger object page for a list of supported trigger severities.
userid (required)	string	ID of the user that uses the media.

usermedia.get

Description

`integer/array usermedia.get(object parameters)`

The method allows to retrieve media according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
mediaids	string/array	Return only media with the given IDs.
usrgrpids	string/array	Return only media that are used by users in the given user groups.
userids	string/array	Return only media that are used by the given users.
mediatypeids	string/array	Return only media that use the given media types.
sortfield	string/array	Sort the result by the given properties. Possible values are: mediaid , userid and mediatypeid .
countOutput	flag	These parameters being common for all get methods are described in detail on the Generic Zabbix API information page .
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	

Parameter	Type	Description
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving media by user

Retrieve all media for the given user.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usermedia.get",
  "params": {
    "output": "extend",
    "userids": "1"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "mediaid": "8",
      "userid": "1",
      "mediatypeid": "3",
      "sendto": "+3711231233",
      "active": "0",
      "severity": "48",
      "period": "1-5,09:00-18:00"
    },
    {
      "mediaid": "9",
      "userid": "1",
      "mediatypeid": "1",
      "sendto": "john@company.com",
      "active": "0",
      "severity": "63",
      "period": "1-7,00:00-24:00"
    }
  ],
  "id": 1
}
```

Source

CUserMedia::get() in *frontends/php/api/classes/CUserMedia.php*.

Media type

This class is designed to work with media types.

Object references:

- **Media type**

Available methods:

- **mediatype.create** - creating new media types
- **mediatype.delete** - deleting media types
- **mediatype.get** - retrieving media types
- **mediatype.update** - updating media types

> Media type object

The following objects are directly related to the mediatype API.

Media type

The media type object has the following properties.

Property	Type	Description
mediatypeid	string	(<i>readonly</i>) ID of the media type.
description (required)	string	Name of the media type.
type (required)	integer	Transport used by the media type. Possible values: 0 - email; 1 - script; 2 - SMS; 3 - Jabber; 100 - Ez Texting.
exec_path	string	For script media types <code>exec_path</code> contains the name of the executed script. For Ez Texting <code>exec_path</code> contains the message text limit. Possible text limit values: 0 - USA (160 characters); 1 - Canada (136 characters).
gsm_modem	string	Required for script and Ez Texting media types. Serial device name of the GSM modem.
passwd	string	Required for SMS media types. Authentication password.
smtp_email	string	Required for Jabber and Ez Texting media types. Email address from which notifications will be sent.
smtp_helo	string	Required for email media types. SMTP HELO.
smtp_server	string	Required for email media types. SMTP server.
status	integer	Required for email media types. Whether the media type is enabled. Possible values: 0 - (<i>default</i>) enabled; 1 - disabled.
username	string	Username or Jabber identifier. Required for Jabber and Ez Texting media types.

mediatype.create

Description

object mediatype.create(object/array mediaTypes)

This method allows to create new media types.

Parameters

(object/array) Media types to create.

The method accepts media types with the **standard media type properties**.

Return values

(object) Returns an object containing the IDs of the created media types under the `mediatypeids` property. The order of the returned IDs matches the order of the passed media types.

Examples

Creating a media type

Create a new e-mail media type.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "mediatype.create",
  "params": {
    "description": "E-mail",
    "type": 0,
    "smtp_server": "rootmail@company.com",
    "smtp_helo": "company.com",
    "smtp_email": "zabbix@company.com"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "mediatypeids": [
      "7"
    ]
  },
  "id": 1
}
```

Source

CMediaType::create() in *frontends/php/api/classes/CMediaType.php*.

mediatype.delete

Description

object mediatype.delete(array mediaTypeIds)

This method allows to delete media types.

Parameters

(array) IDs of the media types to delete.

Return values

(object) Returns an object containing the IDs of the deleted media types under the `mediatypeids` property.

Examples

Deleting multiple media types

Delete two media types.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "mediatype.delete",
  "params": [
    "3",
    "5"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "mediatypeids": [
      "3",
      "5"
    ]
  },
  "id": 1
}
```

Source

CMediaType::delete() in *frontends/php/api/classes/CMediaType.php*.

mediatype.get

Description

integer/array mediatype.get(object parameters)

The method allows to retrieve media types according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
mediatypeids	string/array	Return only media types with the given IDs.
mediaids	string/array	Return only media types used by the given media.
userids	string/array	Return only media types used by the given users.
selectUsers	query	Return the users that use the media type in the users property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: mediatypeid. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	

Parameter	Type	Description
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving media types

Retrieve all configured media types.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "mediatype.get",
  "params": {
    "output": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "mediatypeid": "1",
      "type": "0",
      "description": "Email",
      "smtp_server": "mail.company.com",
      "smtp_helo": "company.com",
      "smtp_email": "zabbix@company.com",
      "exec_path": "",
      "gsm_modem": "",
      "username": "",
      "passwd": "",
      "status": "0"
    },
    {
      "mediatypeid": "2",
      "type": "3",
      "description": "Jabber",
      "smtp_server": "",
      "smtp_helo": "",
      "smtp_email": "",
      "exec_path": "",
      "gsm_modem": "",
      "username": "jabber@company.com",
      "passwd": "zabbix",
      "status": "0"
    }
  ],
}
```

```

    {
        "mediatypeid": "3",
        "type": "2",
        "description": "SMS",
        "smtp_server": "",
        "smtp_helo": "",
        "smtp_email": "",
        "exec_path": "",
        "gsm_modem": "/dev/ttyS0",
        "username": "",
        "passwd": "",
        "status": "0"
    }
],
    "id": 1
}

```

See also

- [User](#)

Source

CMediaType::get() in *frontends/php/api/classes/CMediaType.php*.

mediatype.update

Description

object mediatype.update(object/array mediaTypes)

This method allows to update existing media types.

Parameters

(object/array) **Media type properties** to be updated.

The mediatypeid property must be defined for each media type, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated media types under the mediatypeids property.

Examples

Enabling a media type

Enable a media type, that is, set its status to 0.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "mediatype.update",
    "params": {
        "mediatypeid": "6",
        "status": 0
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "mediatypeids": [
            "6"
        ]
    }
}

```

```

    ],
    "id": 1
}

```

Source

CMediaType::update() in *frontends/php/api/classes/CMediaType.php*.

Proxy

This class is designed to work with proxies.

Object references:

- [Proxy](#)
- [Proxy interface](#)

Available methods:

- [proxy.create](#) - create new proxies
- [proxy.delete](#) - delete proxies
- [proxy.get](#) - retrieve proxies
- [proxy.isreadable](#) - check if a proxy is readable
- [proxy.iswritable](#) - check if a proxy is writable
- [proxy.update](#) - update proxies

> Proxy object

The following objects are directly related to the proxy API.

Proxy

The proxy object has the following properties.

Property	Type	Description
proxyid	string	<i>(readonly)</i> ID of the proxy.
host (required)	string	Name of the proxy.
status (required)	integer	Type of proxy. Possible values: 5 - active proxy; 6 - passive proxy.
lastaccess	timestamp	<i>(readonly)</i> Time when the proxy last connected to the server.

Proxy interface

The proxy interface object defines the interface used to connect to a passive proxy. It has the following properties.

Property	Type	Description
interfaceid	string	<i>(readonly)</i> ID of the interface.
dns (required)	string	DNS name to connect to.
ip (required)	string	Can be empty if connections are made via IP address. IP address to connect to.
port (required)	string	Can be empty if connections are made via DNS names. Port number to connect to.

Property	Type	Description
useip (required)	integer	Whether the connection should be made via IP address. Possible values are: 0 - connect using DNS name; 1 - connect using IP address.
hostid	string	(<i>readonly</i>) ID of the proxy the interface belongs to.

proxy.create

Description

object proxy.create(object/array proxies)

This method allows to create new proxies.

Parameters

(object/array) Proxies to create.

Additionally to the **standard proxy properties**, the method accepts the following parameters.

Parameter	Type	Description
hosts	array	Hosts to be monitored by the proxy. If a host is already monitored by a different proxy, it will be reassigned to the current proxy.
interfaces	array	The hosts must have the <code>hostid</code> property defined. Host interface to be created for the passive proxy passed as an array. Required for passive proxies.

Return values

(object) Returns an object containing the IDs of the created proxies under the `proxyids` property. The order of the returned IDs matches the order of the passed proxies.

Examples

Create an active proxy

Create an action proxy "Active proxy" and assign a host to be monitored by it.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.create",
  "params": {
    "host": "Active proxy",
    "status": "5",
    "hosts": [
      {
        "hostid": "10279"
      }
    ]
  },
  "auth": "ab9638041ec6922cb14b07982b268f47",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "proxyids": [
      "10280"
    ]
  },
  "id": 1
}
```

Create a passive proxy

Create a passive proxy "Passive proxy" and assign two hosts to be monitored by it.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.create",
  "params": {
    "host": "Passive proxy",
    "status": "6",
    "interfaces": [
      {
        "ip": "127.0.0.1",
        "dns": "",
        "useip": "1",
        "port": "10051"
      }
    ],
    "hosts": [
      {
        "hostid": "10192"
      },
      {
        "hostid": "10139"
      }
    ]
  },
  "auth": "ab9638041ec6922cb14b07982b268f47",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "proxyids": [
      "10284"
    ]
  },
  "id": 1
}
```

See also

- [Host](#)
- [Proxy interface](#)

Source

CProxy::create() in *frontends/php/api/classes/CProxy.php*.

proxy.delete

Description

`object proxy.delete(array proxies)`

This method allows to delete proxies.

Parameters

(array) Proxies to delete.

The proxies must have the `proxyid` property defined.

Return values

(object) Returns an object containing the IDs of the deleted proxies under the `proxyids` property.

Examples

Delete multiple proxies

Delete two proxies.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.delete",
  "params": [
    {"proxyid": "10286"},
    {"proxyid": "10285"}
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "proxyids": [
      "10286",
      "10285"
    ]
  },
  "id": 1
}
```

Source

`CProxy::delete()` in *frontends/php/api/classes/CProxy.php*.

proxy.get

Description

`integer/array proxy.get(object parameters)`

The method allows to retrieve proxies according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
<code>proxyids</code>	string/array	Return only proxies with the given IDs.
<code>selectHosts</code>	query	Return hosts monitored by the proxy in the <code>hosts</code> property.

Does not support `shorten`.

Parameter	Type	Description
selectInterfaces	query	Return the proxy interface used by a passive proxy as an array in the <code>interfaces</code> property.
sortfield	string/array	Does not support <code>shorten</code> . Sort the result by the given properties.
countOutput	flag	Possible values are: <code>hostid</code> , <code>host</code> and <code>status</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve all proxies

Retrieve all configured proxies and their interfaces.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.get",
  "params": {
    "output": "extend",
    "selectInterfaces": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "interfaces": [],
      "host": "Active proxy",
      "status": "5",
      "lastaccess": "0",
      "proxyid": "30091"
    },
    {
      "interfaces": [
        {
          "interfaceid": "30109",
```



```

        "hostid": "30092",
        "useip": "1",
        "ip": "127.0.0.1",
        "dns": "",
        "port": "10051"
    }
],
    "host": "Passive proxy",
    "status": "6",
    "lastaccess": "0",
    "proxyid": "30092"
}
],
    "id": 1
}

```

See also

- [Host](#)
- [Proxy interface](#)

Source

CProxy::get() in *frontends/php/api/classes/CProxy.php*.

proxy.isreadable

Description

`boolean proxy.isreadable(array proxyIds)`

This method checks if the given proxies are available for reading.

Parameters

(array) IDs of the proxies to check.

Return values

(boolean) Returns true if the given proxies are available for reading.

Examples

Check multiple proxies

Check if the two proxies are readable.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "proxy.isreadable",
    "params": [
        "30091",
        "30092"
    ],
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}

```

See also

- [proxy.iswritable](#)

Source

CProxy::isReadable() in *frontends/php/api/classes/CProxy.php*.

proxy.iswritable

Description

`boolean proxy.iswritable(array proxyIds)`

This method checks if the given proxies are available for writing.

Parameters

(array) IDs of the proxies to check.

Return values

(boolean) Returns true if the given proxies are available for writing.

Examples

Check multiple proxies

Check if the two proxies are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.iswritable",
  "params": [
    "30091",
    "30092"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [proxy.isreadable](#)

Source

CProxy::isWritable() in *frontends/php/api/classes/CProxy.php*.

proxy.update

Description

`object proxy.update(object/array proxies)`

This method allows to update existing proxies.

Parameters

(object/array) Proxy properties to be updated.

The proxyid property must be defined for each proxy, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard proxy properties](#), the method accepts the following parameters.

Parameter	Type	Description
hosts	array	Hosts to be monitored by the proxy. If a host is already monitored by a different proxy, it will be reassigned to the current proxy.
interfaces	array	The hosts must have the <code>hostid</code> property defined. Host interface to replace the existing interface for the passive proxy passed as an array.
		Required for passive proxies.

Return values

(object) Returns an object containing the IDs of the updated proxies under the `proxyids` property.

Examples

Change hosts monitored by a proxy

Update the proxy to monitor the two given hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.update",
  "params": {
    "proxyid": "10293",
    "hosts": [
      "10294",
      "10295"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "proxyids": [
      "10293"
    ]
  },
  "id": 1
}
```

Change proxy status

Change the proxy to an active proxy and rename it to "Active proxy".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "proxy.update",
  "params": {
    "proxyid": "10293",
    "host": "Active proxy",
    "status": "5"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "proxyids": [
      "10293"
    ]
  },
  "id": 1
}
```

See also

- [Host](#)
- [Proxy interface](#)

Source

CProxy::update() in *frontends/php/api/classes/CProxy.php*.

Screen

This class is designed to work with screen.

Object references:

- [Screen](#)

Available methods:

- [screen.create](#) - creating new screen
- [screen.delete](#) - deleting screens
- [screen.exists](#) - checking if a screen exists
- [screen.get](#) - retrieving screens
- [screen.update](#) - updating screens

> Screen object

The following objects are directly related to the `screen` API.

Screen

The screen object has the following properties.

Property	Type	Description
screenid	string	<i>(readonly)</i> ID of the screen.
name (required)	string	Name of the screen.
hsize	integer	Width of the screen. Default: 1
vsize	integer	Height of the screen. Default: 1

screen.create

Description

object screen.create(object/array screens)

This method allows to create new screens.

Parameters

(object/array) Screens to create.

Additionally to the [standard screen properties](#), the method accepts the following parameters.

Parameter	Type	Description
screenitems	array	Screen items to be created for the screen.

Return values

(object) Returns an object containing the IDs of the created screens under the `screenids` property. The order of the returned IDs matches the order of the passed screens.

Examples

Creating a screen

Create a screen named "Graphs" with 2 rows and 3 columns and add a graph to the upper-left cell.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screen.create",
  "params": {
    "name": "Graphs",
    "hsize": 3,
    "vsize": 2,
    "screenitems": [
      {
        "resourcetype": 0,
        "resourceid": "612",
        "rowspan": 0,
        "colspan": 0,
        "x": 0,
        "y": 0
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "screenids": [
      "26"
    ]
  },
  "id": 1
}
```

See also

- [Screen item](#)

Source

CScreen::create() in *frontends/php/api/classes/CScreen.php*.

screen.delete

Description

object screen.delete(array screenIds)

This method allows to delete screens.

Parameters

(array) IDs of the screens to delete.

Return values

(object) Returns an object containing the IDs of the deleted screens under the `screenids` property.

Examples

Deleting multiple screens

Delete two screens.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screen.delete",
  "params": [
    "25",
    "26"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "screenids": [
      "25",
      "26"
    ]
  },
  "id": 1
}
```

Source

`CScreen::delete()` in *frontends/php/api/classes/CScreen.php*.

screen.exists

Description

`boolean screen.exists(object filter)`

This method checks if at least one screen that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
name	string/array	Names of the screens.
node	string	Name of the node the screens must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the nodes the screens must belong to.
screenid	string/array	IDs of the screens.

Return values

(boolean) Returns `true` if at least one screen that matches the given filter criteria exists.

Examples

Checking a screen by name

Check if a screen named "Graphs" already exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screen.exists",
  "params": {
    "name": "Graphs"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

CScreen::exists() in *frontends/php/api/classes/CScreen.php*.

screen.get

Description

integer/array screen.get(object parameters)

The method allows to retrieve screens according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
screenids	string/array	Return only screens with the given IDs.
screenitemids	string/array	Return only screen that contain the given screen items.
selectScreenItems	query	Return the screen items that are used in the screen.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>screenid</code> and <code>name</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving a screen by ID

Retrieve all data about screen "26" and its screen items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screen.get",
  "params": {
    "output": "extend",
    "selectScreenItems": "extend",
    "screenids": "26"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "screenitems": [
        {
          "screenitemid": "67",
          "screenid": "26",
          "resourcetype": "0",
          "resourceid": "612",
          "width": "320",
          "height": "200",
          "x": "0",
          "y": "0",
          "colspan": "0",
          "rowspan": "0",
          "elements": "25",
          "valign": "0",
          "halign": "0",
          "style": "0",
          "url": "",
          "dynamic": "0",
          "sort_triggers": "0"
        }
      ],
      "screenid": "26",
      "name": "CPU Graphs",
      "hsize": "3",
      "vsize": "2",
      "templateid": "0"
    }
  ],
  "id": 1
}
```

See also

- [Screen item](#)

Source

`CScreen::get()` in *frontends/php/api/classes/CScreen.php*.

screen.update

Description

`object screen.update(object/array screens)`

This method allows to update existing screens.

Parameters

(object/array) Screen properties to be updated.

The `screenid` property must be defined for each screen, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard screen properties](#), the method accepts the following parameters.

Parameter	Type	Description
screenitems	array	Screen items to replace existing screen items.
		Screen items are updated by coordinates, so each screen item must have the <code>x</code> and <code>y</code> properties defined.

Return values

(object) Returns an object containing the IDs of the updated screens under the `screenids` property.

Examples

Renaming a screen

Rename a screen to "CPU Graphs".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screen.update",
  "params": {
    "screenid": "26",
    "name": "CPU Graphs"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "screenids": [
      "26"
    ]
  },
  "id": 1
}
```

See also

- [Screen item](#)
- [screenitem.create](#)
- [screenitem.update](#)
- [screenitem.updatebyposition](#)

Source

CScreen::update() in *frontends/php/api/classes/CScreen.php*.

Screen item

This class is designed to work with screen items.

Object references:

- [Screen item](#)

Available methods:

- [screenitem.create](#) - creating new screen items
- [screenitem.delete](#) - deleting screen items
- [screenitem.get](#) - retrieving screen items
- [screenitem.isreadable](#) - checking if screen items are readable
- [screenitem.iswritable](#) - checking if screen items are writable
- [screenitem.update](#) - updating screen items
- [screenitem.updatebyposition](#) - updating screen items in a specific screen cell

> Screen item object

The following objects are directly related to the `screenitem` API.

Screen item

The screen item object defines an element displayed on a screen. It has the following properties.

Property	Type	Description
<code>screenitemid</code>	string	<i>(readonly)</i> ID of the screen item.
<code>colspan</code> (required)	integer	Number of columns the screen item will span across.
<code>resourcetype</code> (required)	integer	Type of screen item. Possible values: 0 - graph; 1 - simple graph; 2 - map; 3 - plain text; 4 - hosts info; 5 - triggers info; 6 - server info; 7 - clock; 8 - screen; 9 - triggers overview 10 - data overview; 11 - URL; 12 - history of actions; 13 - history of events; 14 - status of host group triggers; 15 - system status; 16 - status of host triggers.
<code>rowspan</code> (required)	integer	Number or rows the screen item will span across.
<code>screenid</code> (required)	string	ID of the screen that the item belongs to.
<code>dynamic</code>	integer	Whether the screen item is dynamic. Possible values: 0 - <i>(default)</i> not dynamic; 1 - dynamic.
<code>elements</code>	integer	Number of lines to display on the screen item. Default: 25.

Property	Type	Description
halign	integer	Specifies how the screen item must be aligned horizontally in the cell. Possible values: 0 - (<i>default</i>) center; 1 - left; 2 - right.
height	integer	Height of the screen item in pixels.
resourceid	string	Default: 200. ID of the object displayed on the screen item. Depending on the type of a screen item, the <code>resourceid</code> property can reference different objects.
sort_triggers	integer	Required for data overview, graph, map, plain text, screen, simple graph and trigger overview screen items. Unused by local and server time clocks, history of actions, history of events, hosts info, server info, system status and URL screen items. Order in which actions or triggers must be sorted. Possible values for history of actions screen elements: 3 - time, ascending; 4 - time, descending; 5 - type, ascending; 6 - type, descending; 7 - status, ascending; 8 - status, descending; 9 - retries left, ascending; 10 - retries left, descending; 11 - recipient, ascending; 12 - recipient, descending.
style	integer	Possible values for status of host group triggers and status of host triggers screen items: 0 - (<i>default</i>) last change, descending; 1 - severity, descending; 2 - host, ascending. Screen item display option. Possible values for data overview and triggers overview screen items: 0 - (<i>default</i>) display hosts on the left side; 1 - display hosts on the top. Possible values for hosts info and triggers info screen elements: 0 - (<i>default</i>) horizontal layout; 1 - vertical layout. Possible values for clock screen items: 0 - (<i>default</i>) local time; 1 - server time; 2 - host time. Possible values for plain text screen items: 0 - (<i>default</i>) display values as plain text; 1 - display values as HTML.
url	string	URL of the webpage to be displayed in the screen item. Used by URL screen items.

Property	Type	Description
valign	integer	Specifies how the screen item must be aligned vertically in the cell. Possible values: 0 - <i>(default)</i> middle; 1 - top; 2 - bottom.
width	integer	Width of the screen item in pixels.
x	integer	Default: 320. X-coordinates of the screen item on the screen, from left to right.
y	integer	Default: 0. Y-coordinates of the screen item on the screen, from top to bottom.
		Default: 0.

screenitem.create

Description

`object screenitem.create(object/array screenItems)`

This method allows to create new screen items.

Parameters

(object/array) Screen items to create.

The method accepts screen items with the **standard screen item properties**.

Return values

(object) Returns an object containing the IDs of the created screen items under the `screenitemids` property. The order of the returned IDs matches the order of the passed screen items.

Examples

Creating a screen item

Create a screen item displaying a graph in the left-upper cell of the screen.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screenitem.create",
  "params": {
    "screenid": 16,
    "resourcetype": 0,
    "resourceid": 612,
    "rowspan": 0,
    "colspan": 0,
    "x": 0,
    "y": 0
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
```

```

        "screenitemids": [
            "65"
        ]
    },
    "id": 1
}

```

See also

- [screen.update](#)

Source

CScreenItem::create() in *frontends/php/api/classes/CScreenItem.php*.

screenitem.delete

Description

object screenitem.delete(array screenItemIds)

This method allows to delete screen items.

Parameters

(array) IDs of the screen items to delete.

Return values

(object) Returns an object containing the IDs of the deleted screen items under the `screenitemids` property.

Examples

Deleting multiple screen items

Delete two screen items.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "screenitem.delete",
    "params": [
        "65",
        "63"
    ],
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "screenitemids": [
            "65",
            "63"
        ]
    },
    "id": 1
}

```

See also

- [screen.update](#)

Source

CScreenItem::delete() in *frontends/php/api/classes/CScreenItem.php*.

screenitem.get

Description

integer/array screenitem.get(object parameters)

The method allows to retrieve screen items according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
screenitemids	string/array	Return only screen items with the given IDs.
screenids	string/array	Return only screen items that belong to the given screen.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: screenitemid and screenid. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving screen items from screen

Retrieve all screen items from the given screen.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screenitem.get",
  "params": {
    "output": "extend",
    "screenids": "3"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
```

```

{
  "screenitemid": "20",
  "screenid": "3",
  "resourcetype": "0",
  "resourceid": "433",
  "width": "500",
  "height": "120",
  "x": "0",
  "y": "0",
  "colspan": "1",
  "rowspan": "1",
  "elements": "0",
  "valign": "1",
  "halign": "0",
  "style": "0",
  "url": "",
  "dynamic": "0",
  "sort_triggers": "0"
},
{
  "screenitemid": "21",
  "screenid": "3",
  "resourcetype": "0",
  "resourceid": "387",
  "width": "500",
  "height": "100",
  "x": "0",
  "y": "1",
  "colspan": "1",
  "rowspan": "1",
  "elements": "0",
  "valign": "1",
  "halign": "0",
  "style": "0",
  "url": "",
  "dynamic": "0",
  "sort_triggers": "0"
},
{
  "screenitemid": "22",
  "screenid": "3",
  "resourcetype": "1",
  "resourceid": "10013",
  "width": "500",
  "height": "148",
  "x": "1",
  "y": "0",
  "colspan": "1",
  "rowspan": "1",
  "elements": "0",
  "valign": "1",
  "halign": "0",
  "style": "0",
  "url": "",
  "dynamic": "0",
  "sort_triggers": "0"
},
{
  "screenitemid": "23",
  "screenid": "3",
  "resourcetype": "1",
  "resourceid": "22181",

```

```

        "width": "500",
        "height": "184",
        "x": "1",
        "y": "1",
        "colspan": "1",
        "rowspan": "1",
        "elements": "0",
        "valign": "1",
        "halign": "0",
        "style": "0",
        "url": "",
        "dynamic": "0",
        "sort_triggers": "0"
    }
],
    "id": 1
}

```

Source

CScreenItem::get() in *frontends/php/api/classes/CScreenItem.php*.

screenitem.isreadable

Description

boolean screenitem.isreadable(array screenItemIds)

This method checks if the given screen items are available for reading.

Parameters

(array) IDs of the screen items to check.

Return values

(boolean) Returns true if the given screen items are available for reading.

Examples

Check multiple screen items

Check if the two screen items are readable.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "screenitem.isreadable",
    "params": [
        "20",
        "21"
    ],
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}

```

See also

- [screenitem.iswritable](#)

Source

CScreenItem::isReadable() in *frontends/php/api/classes/CScreenItem.php*.

screenitem.iswritable

Description

`boolean screenitem.iswritable(array screenItemIds)`

This method checks if the given screen items are available for writing.

Parameters

(array) IDs of the screen items to check.

Return values

(boolean) Returns true if the given screen items are available for writing.

Examples

Check multiple screen items

Check if the two screen items are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screenitem.iswritable",
  "params": [
    "20",
    "21"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [screenitem.isreadable](#)

Source

CScreenItem::isWritable() in *frontends/php/api/classes/CScreenItem.php*.

screenitem.update

Description

`object screenitem.update(object/array screenItems)`

This method allows to update existing screen items.

Parameters

(object/array) **Screen item properties** to be updated.

The screenitemid property must be defined for each screen item, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated screen items under the screenitemids property.

Examples

Setting the size of the screen item

Set the width of the screen item to 500px and height to 300px.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screenitem.update",
  "params": {
    "screenitemid": "20",
    "width": 500,
    "height": 300
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "screenitemids": [
      "20"
    ]
  },
  "id": 1
}
```

See also

- [screenitem.updatebyposition](#)

Source

CScreenItem::update() in *frontends/php/api/classes/CScreenItem.php*.

screenitem.updatebyposition

Description

object screenitem.updatebyposition(array screenItems)

This method allows to update screen items in the given screen cells. If a cell is empty, a new screen item will be created.

Parameters

(array) **Screen item properties** to be updated.

The x, y and screenid properties must be defined for each screen item, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated and created screen items under the screenitemids property.

Examples

Changing a screen items resource ID

Change the resource ID for the screen element located in the upper-left cell of the screen.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "screenitem.updatebyposition",
  "params": [
    {
      "screenid": "16",
```

```

        "x": 0,
        "y": 0,
        "resourceid": "644"
    }
],
"auth": "038e1d7b1735c6a5436ee9eae095879e",
"id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "screenitemids": [
            "66"
        ]
    },
    "id": 1
}

```

See also

- [screenitem.update](#)

Source

CScreenItem::update() in *frontends/php/api/classes/CScreenItem.php*.

Script

This class is designed to work with scripts.

Object references:

- [Script](#)

Available methods:

- [script.create](#) - create new scripts
- [script.delete](#) - delete scripts
- [script.execute](#) - run scripts
- [script.get](#) - retrieve scripts
- [script.getscriptsbyhosts](#) - retrieve scripts for hosts
- [script.update](#) - update scripts

> Script object

The following objects are directly related to the script API.

Script

The script object has the following properties.

Property	Type	Description
scriptid	string	(<i>readonly</i>) ID of the script.
command (required)	string	Command to run.
name (required)	string	Name of the script.
confirmation	string	Confirmation pop up text. The pop up will appear when trying to run the script from the Zabbix frontend.
description	string	Description of the script.

Property	Type	Description
execute_on	integer	Where to run the script.
groupid	string	Possible values: 0 - run on Zabbix agent; 1 - <i>(default)</i> run on Zabbix server. ID of the host group that the script can be run on. If set to 0, the script will be available on all host groups.
host_access	integer	Default: 0. Host permissions needed to run the script.
type	integer	Possible values: 2 - <i>(default)</i> read; 3 - write. Script type.
usrgrpid	string	Possible values: 0 - <i>(default)</i> script; 1 - IPMI. ID of the user group that will be allowed to run the script. If set to 0, the script will be available for all user groups.
		Default: 0.

script.create

Description

object script.create(object/array scripts)

This method allows to create new scripts.

Parameters

(object/array) Scripts to create.

The method accepts scripts with the **standard script properties**.

Return values

(object) Returns an object containing the IDs of the created scripts under the `scriptids` property. The order of the returned IDs matches the order of the passed scripts.

Examples

Create a script

Create a script that will reboot a server. The script will require write access to the host and will display a configuration message before running in the frontend.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "script.create",
  "params": {
    "name": "Reboot server",
    "command": "reboot server 1",
    "host_access": 3,
    "confirmation": "Are you sure you would like to reboot the server?"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "scriptids": [
      "3"
    ]
  },
  "id": 1
}
```

Source

CScript::create() in *frontends/php/api/classes/CScript.php*.

script.delete

Description

object script.delete(array scriptIds)

This method allows to delete scripts.

Parameters

(array) IDs of the scripts to delete.

Return values

(object) Returns an object containing the IDs of the deleted scripts under the `scriptids` property.

Examples

Delete multiple scripts

Delete two scripts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "script.delete",
  "params": [
    "3",
    "4"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "scriptids": [
      "3",
      "4"
    ]
  },
  "id": 1
}
```

Source

CScript::delete() in *frontends/php/api/classes/CScript.php*.

script.execute

Description

```
object script.execute(object parameters)
```

This method allows to run a script on a host.

Parameters

(object) Parameters containing the ID of the script to run and the ID of the host.

Parameter	Type	Description
hostid (required)	string	ID of the host to run the script on.
scriptid (required)	string	ID of the script to run.

Return values

(object) Returns the result of script execution.

Property	Type	Description
response	string	Whether the script was run successfully.
value	string	Possible values: success or failed. Script output.

Examples

Run a script

Run a "ping" script on a host.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "script.execute",
  "params": {
    "scriptid": "1",
    "hostid": "30079"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "response": "success",
    "value": "PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.\n64 bytes from 127.0.0.1: icmp_req=1 tt"
  },
  "id": 1
}
```

Source

CScript::execute() in *frontends/php/api/classes/CScript.php*.

script.get

Description

```
integer/array script.get(object parameters)
```

The method allows to retrieve scripts according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	string/array	Return only scripts that can be run on the given host groups.
hostids	string/array	Return only scripts that can be run on the given hosts.
scriptids	string/array	Return only scripts with the given IDs.
usrgrpids	string/array	Return only scripts that can be run by users in the given user groups.
selectGroups	query	Return host groups that the script can be run on in the groups property.
selectHosts	query	Does not support shorten or arrays of properties. Return hosts that the script can be run on in the hosts property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: scriptid and name. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve all scripts

Retrieve all configured scripts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "script.get",
  "params": {
    "output": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```

{
  "jsonrpc": "2.0",
  "result": [
    {
      "scriptid": "1",
      "name": "Ping",
      "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
      "host_access": "2",
      "usrgrpuid": "0",
      "groupid": "0",
      "description": "",
      "confirmation": "",
      "type": "0",
      "execute_on": "1"
    },
    {
      "scriptid": "2",
      "name": "Traceroute",
      "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
      "host_access": "2",
      "usrgrpuid": "0",
      "groupid": "0",
      "description": "",
      "confirmation": "",
      "type": "0",
      "execute_on": "1"
    },
    {
      "scriptid": "3",
      "name": "Detect operating system",
      "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
      "host_access": "2",
      "usrgrpuid": "7",
      "groupid": "0",
      "description": "",
      "confirmation": "",
      "type": "0",
      "execute_on": "1"
    }
  ],
  "id": 1
}

```

See also

- [Host](#)
- [Host group](#)

Source

CScript::get() in *frontends/php/api/classes/CScript.php*.

script.getscriptsbyhosts

Description

object script.getscriptsbyhosts(array hostIds)

This method allows to retrieve scripts available on the given hosts.

Parameters

(string/array) IDs of hosts to return scripts for.

Return values

(object) Returns an object with host IDs as properties and arrays of available scripts as values.

Examples

Retrieve scripts by host IDs

Retrieve all scripts available on hosts "30079" and "30073".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "script.getscriptsbyhosts",
  "params": [
    "30079",
    "30073"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "30079": [
      {
        "scriptid": "3",
        "name": "Detect operating system",
        "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpuid": "7",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "hostid": "10001"
      },
      {
        "scriptid": "1",
        "name": "Ping",
        "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpuid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "hostid": "10001"
      },
      {
        "scriptid": "2",
        "name": "Traceroute",
        "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpuid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "hostid": "10001"
      }
    ],
    "30073": [

```

```

    {
        "scriptid": "3",
        "name": "Detect operating system",
        "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpid": "7",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "hostid": "10001"
    },
    {
        "scriptid": "1",
        "name": "Ping",
        "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "hostid": "10001"
    },
    {
        "scriptid": "2",
        "name": "Traceroute",
        "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "hostid": "10001"
    }
]
},
"id": 1
}

```

Source

CScript::getScriptsByHosts() in *frontends/php/api/classes/CScript.php*.

script.update

Description

object script.update(object/array scripts)

This method allows to update existing scripts.

Parameters

(object/array) **Script properties** to be updated.

The scriptid property must be defined for each script, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated scripts under the scriptids property.

Examples

Change script command

Change the command of the script to `"/bin/ping -c 10 {HOST.CONN} 2>&1"`.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "script.update",
  "params": {
    "scriptid": "1",
    "command": "/bin/ping -c 10 {HOST.CONN} 2>&1"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "scriptids": [
      "1"
    ]
  },
  "id": 1
}
```

Source

CScript::update() in *frontends/php/api/classes/CScript.php*.

Template

This class is designed to work with templates.

Object references:

- [Template](#)

Available methods:

- [template.create](#) - creating new templates
- [template.delete](#) - deleting templates
- [template.exists](#) - checking if a template exists
- [template.get](#) - retrieving templates
- [template.isreadable](#) - checking if templates are readable
- [template.iswritable](#) - checking if templates are writable
- [template.massadd](#) - adding related objects to templates
- [template.massremove](#) - removing related objects from templates
- [template.massupdate](#) - replacing or removing related objects from templates
- [template.update](#) - updating templates

> Template object

The following objects are directly related to the `template` API.

Template

The template object has the following properties.

Property	Type	Description
<code>templateid</code>	string	(<i>readonly</i>) ID of the template.

Property	Type	Description
host (required)	string	Technical name of the template.
name	string	Visible name of the host. Default: host property value.

template.create

Description

object template.create(object/array templates)

This method allows to create new templates.

Parameters

(object/array) Templates to create.

Additionally to the [standard template properties](#), the method accepts the following parameters.

Parameter	Type	Description
groups (required)	object/array	Host groups to add the template to. The host groups must have the <code>groupid</code> property defined.
templates	object/array	Templates to be linked to the template. The templates must have the <code>templateid</code> property defined.
macros	object/array	User macros to be created for the template.
hosts	object/array	Hosts to link the template to.

Return values

(object) Returns an object containing the IDs of the created templates under the `templateids` property. The order of the returned IDs matches the order of the passed templates.

Examples

Creating a template

Create a template and link it to two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.create",
  "params": {
    "host": "Linux template",
    "groups": {
      "groupid": 1
    },
    "hosts": [
      {
        "hostid": "10084"
      },
      {
        "hostid": "10090"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10086"
    ]
  },
  "id": 1
}
```

Source

CTemplate::create() in *frontends/php/api/classes/CTemplate.php*.

template.delete

Description

object template.delete(array templateIds)

This method allows to delete templates.

Parameters

(array) IDs of the templates to delete.

Return values

(object) Returns an object containing the IDs of the deleted templates under the `templateids` property.

Examples

Deleting multiple templates

Delete two templates.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.delete",
  "params": [
    "13",
    "32"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "13",
      "32"
    ]
  },
  "id": 1
}
```

Source

CTemplate::delete() in *frontends/php/api/classes/CTemplate.php*.

template.exists

Description

`boolean template.exists(object filter)`

This method checks if at least one template that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
host	string/array	Technical names of the templates.
name	string/array	Visible names of the templates.
node	string	Name of the node the templates must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. IDs of the node the hosts must belong to.
templateid	string/array	Template IDs.

Return values

(boolean) Returns `true` if at least one template that matches the given filter criteria exists.

Examples

Check template on a node

Check if a template with the technical name "Linux template" exists on the node with ID 1.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.exists",
  "params": {
    "host": "Linux template",
    "nodeids": [
      "1"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [template.isreadable](#)
- [template.iswritable](#)

Source

CTemplate::exists() in *frontends/php/api/classes/CTemplate.php*.

template.get

Description

`integer/array template.get(object parameters)`

The method allows to retrieve templates according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
templateids	string/array	Return only templates with the given template IDs.
groupids	string/array	Return only templates that belong to the given host groups.
parentTemplateids	string/array	Return only templates that are children of the given templates.
hostids	string/array	Return only templates that are linked to the given hosts.
graphids	string/array	Return only templates that contain the given graphs.
itemids	string/array	Return only templates that contain the given items.
triggerids	string/array	Return only templates that contain the given triggers.
with_items	flag	Return only templates that have items.
with_triggers	flag	Return only templates that have triggers.
with_graphs	flag	Return only templates that have graphs.
selectGroups	query	Return the host groups that the template belongs to in the <code>groups</code> property.
selectHosts	query	Does not support <code>shorten</code> or arrays of properties. Return the hosts that are linked to the template in the <code>hosts</code> property.
selectTemplates	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return the child templates in the <code>templates</code> property.
selectParentTemplates	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return the parent templates in the <code>parentTemplates</code> property.
selectItems	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return items from the template in the <code>items</code> property.
selectDiscoveries	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return low-level discoveries from the template in the <code>discoveries</code> property.
selectTriggers	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return triggers from the template in the <code>triggers</code> property.
selectGraphs	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return graphs from the template in the <code>graphs</code> property.
selectApplications	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return applications from the template in the <code>applications</code> property.
selectMacros	query	Supports <code>count</code> ; does not support <code>shorten</code> . Return the macros from the template in the <code>macros</code> property.
		Does not support <code>shorten</code> or arrays of properties.

Parameter	Type	Description
selectScreens	query	Return screens from the template in the screens property.
limitSelects	integer	Supports count; does not support shorten. Limits the number of records returned by subselects. Applies to the following subselects: selectTemplates - results will be sorted by name; selectHosts - sorted by host; selectParentTemplates - sorted by host; selectItems - sorted by name; selectDiscoveries - sorted by name; selectTriggers - sorted by description; selectGraphs - sorted by name; selectApplications - sorted by name; selectScreens - sorted by name.
sortfield	string/array	Sort the result by the given properties.
countOutput	flag	Possible values are: hostid, host, name, status. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving templates by name

Retrieve all data about two templates named "Template OS Linux" and "Template OS Windows".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.get",
  "params": {
    "output": "extend",
    "filter": {
      "host": [
        "Template OS Linux",
        "Template OS Windows"
      ]
    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
```



```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "hostid": "10001",
      "proxy_hostid": "0",
      "host": "Template OS Linux",
      "status": "3",
      "disable_until": "0",
      "error": "",
      "available": "0",
      "errors_from": "0",
      "lastaccess": "0",
      "ipmi_authtype": "0",
      "ipmi_privilege": "2",
      "ipmi_username": "",
      "ipmi_password": "",
      "ipmi_disable_until": "0",
      "ipmi_available": "0",
      "snmp_disable_until": "0",
      "snmp_available": "0",
      "maintenanceid": "0",
      "maintenance_status": "0",
      "maintenance_type": "0",
      "maintenance_from": "0",
      "ipmi_errors_from": "0",
      "snmp_errors_from": "0",
      "ipmi_error": "",
      "snmp_error": "",
      "jmx_disable_until": "0",
      "jmx_available": "0",
      "jmx_errors_from": "0",
      "jmx_error": "",
      "name": "Template OS Linux",
      "templateid": "10001"
    },
    {
      "hostid": "10081",
      "proxy_hostid": "0",
      "host": "Template OS Windows",
      "status": "3",
      "disable_until": "0",
      "error": "",
      "available": "0",
      "errors_from": "0",
      "lastaccess": "0",
      "ipmi_authtype": "0",
      "ipmi_privilege": "2",
      "ipmi_username": "",
      "ipmi_password": "",
      "ipmi_disable_until": "0",
      "ipmi_available": "0",
      "snmp_disable_until": "0",
      "snmp_available": "0",
      "maintenanceid": "0",
      "maintenance_status": "0",
      "maintenance_type": "0",
      "maintenance_from": "0",

```

```

        "ipmi_errors_from": "0",
        "snmp_errors_from": "0",
        "ipmi_error": "",
        "snmp_error": "",
        "jmx_disable_until": "0",
        "jmx_available": "0",
        "jmx_errors_from": "0",
        "jmx_error": "",
        "name": "Template OS Windows",
        "templateid": "10081"
    }
],
    "id": 1
}

```

See also

- [template.getobjects](#)
- [Host group](#)
- [Template](#)
- [User macro](#)
- [Host interface](#)

Source

CTemplate::get() in *frontends/php/api/classes/CTemplate.php*.

template.getobjects

Description

array `template.getobjects(object filter)`

This method allows to retrieve templates that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard template properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
node	string	Name of the node the templates must belong to.
nodeids	string/array	This will override the <code>nodeids</code> parameter. ID of the node the templates must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving templates by name

Retrieve all data about two templates named "Template OS Linux" and "Template OS Windows".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "template.getobjects",
    "params": {
        "host": [
            "Template OS Linux",
            "Template OS Windows"
        ]
    }
},

```

```
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "host": "Template OS Linux",
      "name": "Template OS Linux",
      "templateid": "10001"
    },
    {
      "host": "Template OS Windows",
      "name": "Template OS Windows",
      "templateid": "10081"
    }
  ],
  "id": 1
}
```

See also

- [template.get](#)

Source

CTemplate::getObject() in *frontends/php/api/classes/CTemplate.php*.

template.isreadable

Description

`boolean template.isreadable(array templateIds)`

This method checks if the given templates are available for reading.

Parameters

(array) IDs of the templates to check.

Return values

(boolean) Returns true if the given templates are available for reading.

Examples

Check multiple templates

Check if the two templates are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.isreadable",
  "params": [
    "10001",
    "10081"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
}
```

```
    "id": 1
}
```

See also

- [template.exists](#)
- [template.iswritable](#)

Source

CTemplate::isReadable() in *frontends/php/api/classes/CTemplate.php*.

template.iswritable

Description

boolean `template.iswritable(array templateIds)`

This method checks if the given templates are available for writing.

Parameters

(array) IDs of the templates to check.

Return values

(boolean) Returns true if the given templates are available for writing.

Examples

Check multiple templates

Check if the two templates are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.iswritable",
  "params": [
    "10001",
    "10081"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [template.isreadable](#)
- [template.exists](#)

Source

CTemplate::isWritable() in *frontends/php/api/classes/CTemplate.php*.

template.massadd

Description

object `template.massadd(object parameters)`

This method allows to simultaneously add multiple related objects to the given templates.

Parameters

(object) Parameters containing the IDs of the templates to update and the objects to add to the templates.

The method accepts the following parameters.

Parameter	Type	Description
templates (required)	object/array	Templates to be updated. The templates must have the <code>templateid</code> property defined.
groups	object/array	Host groups to add the given templates to. The host groups must have the <code>groupid</code> property defined.
hosts	object/array	Hosts to link the given templates to. The hosts must have the <code>templateid</code> property defined.
macros	object/array	User macros to be created for the given templates.
templates_link	object/array	Templates to link to the given templates. The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Adding templates to a group

Add two templates to the host group "2".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.massadd",
  "params": {
    "templates": [
      {
        "templateid": "10085"
      },
      {
        "templateid": "10086"
      }
    ],
    "groups": [
      {
        "groupid": "2"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10085",
      "10086"
    ]
  }
}
```

```

    },
    "id": 1
}

```

Linking a template to hosts

Link template "10073" to two hosts.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "template.massadd",
  "params": {
    "templates": [
      {
        "templateid": "10073"
      }
    ],
    "hosts": [
      {
        "hostid": "10106"
      },
      {
        "hostid": "10104"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10073"
    ]
  },
  "id": 1
}

```

See also

- [template.update](#)
- [Host](#)
- [Host group](#)
- [User macro](#)

Source

CTemplate::massAdd() in *frontends/php/api/classes/CTemplate.php*.

template.massremove

Description

object template.massremove(object parameters)

This method allows to remove related objects from multiple templates.

Parameters

(object) Parameters containing the IDs of the templates to update and the objects that should be removed.

Parameter	Type	Description
templateids (required)	string/array	IDs of the templates to be updated.
groupids	string/array	Host groups to remove the given templates from.
hostids	string/array	Hosts or templates to unlink the given templates from.
macros	string/array	User macros to delete from the given templates.
templateids_clear	string/array	Templates to unlink and clear from the given templates.
templateids_link	string/array	Templates to unlink from the given templates.

Return values

(object) Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Removing templates from a group

Remove two templates from group "2".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.massremove",
  "params": {
    "templateids": [
      "10085",
      "10086"
    ],
    "groupids": "2"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10085",
      "10086"
    ]
  },
  "id": 1
}
```

Unlinking templates from a host

Unlink template "10085" from two hosts.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.massremove",
  "params": {
    "templateids": "10085",
    "hostids": [
      "10106",
      "10104"
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10085"
    ]
  },
  "id": 1
}
```

See also

- [template.update](#)
- [User macro](#)

Source

CTemplate::massRemove() in *frontends/php/api/classes/CTemplate.php*.

template.massupdate

Description

`object template.massupdate(object parameters)`

This method allows to simultaneously replace or remove related objects and update properties on multiple templates.

Parameters

(object) Parameters containing the IDs of the templates to update and the properties that should be updated.

Additionally to the [standard template properties](#), the method accepts the following parameters.

Parameter	Type	Description
templates (required)	object/array	Templates to be updated. The templates must have the <code>templateid</code> property defined.
groups	object/array	Host groups to replace the current host groups the templates belong to. The host groups must have the <code>groupid</code> property defined.
hosts	object/array	Hosts and templates to replace the ones the templates are currently linked to. Both hosts and templates must use the <code>hostid</code> property to pass an ID.
macros	object/array	User macros to replace the current user macros on the given templates.
templates_clear	object/array	Templates to unlink and clear from the given templates.
templates_link	object/array	The templates must have the <code>templateid</code> property defined. Templates to replace the currently linked templates. The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Replacing host groups

Unlink and clear template "10091" from the given templates.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.massupdate",
  "params": {
    "templates": [
      {
        "templateid": "10085"
      },
      {
        "templateid": "10086"
      }
    ],
    "templates_clear": [
      {
        "templateid": "10091"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10085",
      "10086"
    ]
  },
  "id": 1
}
```

See also

- [template.update](#)
- [template.massadd](#)
- [Host group](#)
- [User macro](#)

Source

CTemplate::massUpdate() in *frontends/php/api/classes/CTemplate.php*.

template.update

Description

object `template.update(object/array templates)`

This method allows to update existing templates.

Parameters

(object/array) Template properties to be updated.

The `templateid` property must be defined for each template, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Additionally to the **standard template properties**, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to replace the current host groups the templates belong to.
hosts	object/array	The host groups must have the <code>groupid</code> property defined. Hosts and templates to replace the ones the templates are currently linked to.
macros	object/array	Both hosts and templates must use the <code>hostid</code> property to pass an ID. User macros to replace the current user macros on the given templates.
templates	object/array	Templates to replace the currently linked templates. Templates that are not passed are only unlinked.
templates_clear	object/array	The templates must have the <code>templateid</code> property defined. Templates to unlink and clear from the given templates.
		The templates must have the <code>templateid</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Renaming a template

Rename the template to "Template OS Linux".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "template.update",
  "params": {
    "templateid": "10086",
    "name": "Template OS Linux"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "templateids": [
      "10086"
    ]
  },
  "id": 1
}
```

Source

CTemplate::update() in *frontends/php/api/classes/CTemplate.php*.

Template screen

This class is designed to work with template screens.

Object references:

- [Template screen](#)

Available methods:

- [templatescreen.copy](#) - copy template screens
- [templatescreen.create](#) - create new template screens
- [templatescreen.delete](#) - delete template screens
- [templatescreen.exists](#) - check if a template screen exists
- [templatescreen.get](#) - retrieve template screens
- [templatescreen.isreadable](#) - check if template screens are readable
- [templatescreen.iswritable](#) - check if template screens are writable
- [templatescreen.update](#) - update template screens

> Template screen object

The following objects are directly related to the `templatescreen` API.

Template screen

The template screen object has the following properties.

Property	Type	Description
screenid	string	(readonly) ID of the template screen.
name (required)	string	Name of the template screen.
templateid (required)	string	ID of the template that the screen belongs to.
hsize	integer	Width of the template screen.
vsize	integer	Default: 1
		Height of the template screen.
		Default: 1

templatescreen.copy

Description

```
object templatescreen.copy(object parameters)
```

This method allows to copy template screens to the given templates.

Note:

Added in 2.0.4.

Parameters

(object) Parameters defining the template screens to copy and the target templates.

Parameter	Type	Description
screenids (required)	string/array	IDs of template screens to copy.
templateids (required)	string/array	IDs of templates to copy the screens to.

Return values

(boolean) Returns true if the copying was successful.

Examples

Copy a template screen

Copy template screen "25" to template "30085".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.copy",
  "params": {
    "screenIds": "25",
    "templateIds": "30085"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

Source

CTemplateScreen::copy() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.create

Description

object templatescreen.create(object/array templateScreens)

This method allows to create new template screens.

Parameters

(object/array) Template screens to create.

Additionally to the **standard template screen properties**, the method accepts the following parameters.

Parameter	Type	Description
screenitems	array	Template screen items to create on the screen.

Return values

(object) Returns an object containing the IDs of the created template screens under the `screenids` property. The order of the returned IDs matches the order of the passed template screens.

Examples

Create a template screen

Create a template screen named "Graphs" with 2 rows and 3 columns and add a graph to the upper-left cell.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.create",
  "params": {
    "name": "Graphs",
    "templateid": "10047",
  },
}
```

```

        "hsize": 3,
        "vsize": 2,
        "screenitems": [
            {
                "resourcetype": 0,
                "resourceid": "410",
                "rowspan": 0,
                "colspan": 0,
                "x": 0,
                "y": 0
            }
        ]
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "screenids": [
            "45"
        ]
    },
    "id": 1
}

```

See also

- [Template screen item](#)

Source

CTemplateScreen::create() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.delete

Description

object templatescreen.delete(array templateScreenIds)

This method allows to delete template screens.

Parameters

(array) IDs of the template screens to delete.

Return values

(object) Returns an object containing the IDs of the deleted template screens under the `screenids` property.

Examples

Delete multiple template screens

Delete two template screens.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "templatescreen.delete",
    "params": [
        "45",
        "46"
    ],
    "auth": "3a57200802b24cda67c4e4010b50c065",
}

```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "screenids": [
      "45",
      "46"
    ]
  },
  "id": 1
}
```

Source

CTemplateScreen::delete() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.exists

Description

boolean templatescreen.exists(object filter)

This method checks if at least one template screen that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
name	string/array	Names of the screens.
node	string	Name of the node the template screens must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the template screens must belong to.
screenid	string/array	IDs of the screens.
templateids	string/array	IDs of the templates that the screens belong to

Return values

(boolean) Returns true if at least one template screen that matches the given filter criteria exists.

Examples

Check screen by name

Check if screen "Zabbix server health" exists on template "10047".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.exists",
  "params": {
    "name": "Zabbix server health",
    "templateid": "10047"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [templatescreen.isreadable](#)
- [templatescreen.iswritable](#)

Source

CTemplateScreen::exists() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.get

Description

integer/array templatescreen.get(object parameters)

The method allows to retrieve template screens according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
hostids	string/array	Return only template screens that belong to the given hosts.
screenids	string/array	Return only template screens with the given IDs.
screenitemids	string/array	Return only template screens that contain the given screen items.
templateids	string/array	Return only template screens that belong to the given templates.
noInheritance	flag	Do not return inherited template screens.
selectScreenItems	query	Return the screen items that are used in the template screen in the screenitems property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>screenid</code> and <code>name</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;

- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieve screens from template

Retrieve all screens from template "10001" and all of the screen items.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.get",
  "params": {
    "output": "extend",
    "selectScreenItems": "extend",
    "templateids": "10001"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "screenid": "3",
      "name": "System performance",
      "hsize": "2",
      "vsize": "2",
      "templateid": "10001",
      "screenitems": [
        {
          "screenitemid": "20",
          "screenid": "3",
          "resourcetype": "0",
          "resourceid": "433",
          "width": "500",
          "height": "120",
          "x": "0",
          "y": "0",
          "colspan": "1",
          "rowspan": "1",
          "elements": "0",
          "valign": "1",
          "halign": "0",
          "style": "0",
          "url": ""
        },
        {
          "screenitemid": "21",
          "screenid": "3",
          "resourcetype": "0",
          "resourceid": "387",
          "width": "500",
          "height": "100",
          "x": "0",
          "y": "1",
          "colspan": "1",
          "rowspan": "1",
          "elements": "0",
          "valign": "1",
          "halign": "0",
          "style": "0",
          "url": ""
        }
      ]
    }
  ]
}
```



```

        "url": ""
    },
    {
        "screenitemid": "22",
        "screenid": "3",
        "resourcetype": "1",
        "resourceid": "10013",
        "width": "500",
        "height": "148",
        "x": "1",
        "y": "0",
        "colspan": "1",
        "rowspan": "1",
        "elements": "0",
        "valign": "1",
        "halign": "0",
        "style": "0",
        "url": ""
    },
    {
        "screenitemid": "23",
        "screenid": "3",
        "resourcetype": "1",
        "resourceid": "22181",
        "width": "500",
        "height": "184",
        "x": "1",
        "y": "1",
        "colspan": "1",
        "rowspan": "1",
        "elements": "0",
        "valign": "1",
        "halign": "0",
        "style": "0",
        "url": ""
    }
]
},
    "id": 1
}

```

See also

- [Template screen item](#)

Source

CTemplateScreen::get() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.isreadable

Description

boolean templatescreen.isreadable(array templateScreenIds)

This method checks if the given template screens are available for reading.

Note:

Added in 2.0.4.

Parameters

(array) IDs of the template screens to check.

Return values

(boolean) Returns true if the given template screens are available for reading.

Examples

Check multiple template screens

Check if the two template screens are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.isreadable",
  "params": [
    "3",
    "5"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [templatescreen.exists](#)
- [templatescreen.iswritable](#)

Source

CTemplateScreen::isReadable() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.iswritable

Description

boolean templatescreen.iswritable(array templateScreenIds)

This method checks if the given template screens are available for writing.

Note:

Added in 2.0.4.

Parameters

(array) IDs of the template screens to check.

Return values

(boolean) Returns true if the given template screens are available for writing.

Examples

Check multiple template screens

Check if the two template screens are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.iswritable",
  "params": [
    "3",
    "5"
  ],
}
```

```
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [templatescreen.isreadable](#)
- [templatescreen.exists](#)

Source

CTemplateScreen::isWritable() in *frontends/php/api/classes/CTemplateScreen.php*.

templatescreen.update

Description

`object templatescreen.update(object/array templateScreens)`

This method allows to update existing template screens.

Parameters

(object/array) Template screen properties to be updated.

The `screenid` property must be defined for each template screen, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard template screen properties](#), the method accepts the following parameters.

Parameter	Type	Description
screenitems	array	Screen items to replace existing screen items. Screen items are updated by coordinates, so each screen item must have the <code>x</code> and <code>y</code> properties defined.

Return values

(object) Returns an object containing the IDs of the updated template screens under the `screenids` property.

Examples

Rename a template screen

Rename the template screen to "Performance graphs".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreen.update",
  "params": {
    "screenid": "3",
    "name": "Performance graphs"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "screenids": [
      "3"
    ]
  },
  "id": 1
}
```

Source

CTemplateScreen::update() in *frontends/php/api/classes/CTemplateScreen.php*.

Template screen item

This class is designed to work with template screen items.

Object references:

- [Template screen item](#)

Available methods:

- [templatescreenitem.get](#) - retrieve template screen items

> Template screen item object

The following objects are directly related to the `templatescreenitem` API.

Template screen item

The template screen item object defines an element displayed on a template screen. It has the following properties.

Property	Type	Description
<code>screenitemid</code>	string	<i>(readonly)</i> ID of the template screen item.
<code>colspan</code> (required)	integer	Number of columns the template screen item will span across.
<code>resourceid</code> (required)	string	ID of the object from the parent template displayed on the template screen item. Depending on the type of screen item, the <code>resourceid</code> property can reference different objects. Unused by clock and URL template screen items.
<code>resourcetype</code> (required)	integer	<i>Note: the <code>resourceid</code> property always references an object used in the parent template object, even if the screen item itself is inherited on a host or template.</i> Type of template screen item. Possible values: 0 - graph; 1 - simple graph; 3 - plain text; 7 - clock; 11 - URL.
<code>rowspan</code> (required)	integer	Number or rows the template screen item will span across.
<code>screenid</code> (required)	string	ID of the template screen that the item belongs to.
<code>elements</code>	integer	Number of lines to display on the template screen item. Default: 25.

Property	Type	Description
halign	integer	Specifies how the template screen item must be aligned horizontally in the cell. Possible values: 0 - <i>(default)</i> center; 1 - left; 2 - right.
height	integer	Height of the template screen item in pixels.
style	integer	Default: 200. Template screen item display option. Possible values for clock screen items: 0 - <i>(default)</i> local time; 1 - server time; 2 - host time. Possible values for plain text screen items: 0 - <i>(default)</i> display values as plain text; 1 - display values as HTML.
url	string	URL of the webpage to be displayed in the template screen item. Used by URL template screen items.
valign	integer	Specifies how the template screen item must be aligned vertically in the cell. Possible values: 0 - <i>(default)</i> middle; 1 - top; 2 - bottom.
width	integer	Width of the template screen item in pixels.
x	integer	Default: 320. X-coordinates of the template screen item on the screen, from left to right.
y	integer	Default: 0. Y-coordinates of the template screen item on the screen, from top to bottom.
		Default: 0.

templatescreenitem.get

Description

integer/array templatescreenitem.get(object parameters)

The method allows to retrieve template screen items according to the given parameters.

Note:

Added in 2.0.3.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
screenids	string/array	Return only template screen items that belong to the given template screens.

Parameter	Type	Description
screenitemids	string/array	Return only template screen items with the given IDs. Returns an additional <code>real_resourceid</code> property for each template screen item, that belongs to a screen from the given hosts or templates. The <code>real_resourceid</code> property contains the ID of object displayed on the screen.
hostids	string/array	
sortfield	string/array	
countOutput	flag	Possible values are: <code>screenitemid</code> and <code>screenid</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve template screen items for screen

Return all template screen items from template screen "15".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "templatescreenitem.get",
  "params": {
    "output": "extend",
    "screenids": "15"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "screenitemid": "42",
      "screenid": "15",
      "resourcetype": "0",
      "resourceid": "454",
      "width": "500",
      "height": "200",
      "x": "0",
      "y": "0",

```

```

        "colspan": "1",
        "rowspan": "1",
        "elements": "0",
        "valign": "1",
        "halign": "0",
        "style": "0",
        "url": ""
    },
    {
        "screenitemid": "43",
        "screenid": "15",
        "resourcetype": "0",
        "resourceid": "455",
        "width": "500",
        "height": "270",
        "x": "1",
        "y": "0",
        "colspan": "1",
        "rowspan": "1",
        "elements": "0",
        "valign": "1",
        "halign": "0",
        "style": "0",
        "url": ""
    }
],
    "id": 1
}

```

Source

CTemplateScreenItem::get() in *frontends/php/api/classes/CTemplateScreenItem.php*.

Trigger

This class is designed to work with triggers.

Object references:

- [Trigger](#)

Available methods:

- [trigger.adddependencies](#) - adding new trigger dependencies
- [trigger.create](#) - creating new triggers
- [trigger.delete](#) - deleting triggers
- [trigger.deletedependencies](#) - deleting trigger dependencies
- [trigger.exists](#) - checking if a trigger exists
- [trigger.get](#) - retrieving triggers
- [trigger.getobjects](#) - retrieving triggers by filters
- [trigger.isreadable](#) - checking if triggers are readable
- [trigger.iswritable](#) - checking if triggers are writable
- [trigger.update](#) - updating triggers

> Trigger object

The following objects are directly related to the `trigger` API.

Trigger

The trigger object has the following properties.

Property	Type	Description
triggerid	string	(<i>readonly</i>) ID of the trigger.
description (required)	string	Name of the trigger.
expression (required)	string	Reduced trigger expression.
comments	string	Additional comments to the trigger.
error	string	(<i>readonly</i>) Error text if there have been any problems when updating the state of the trigger.
flags	integer	(<i>readonly</i>) Origin of the trigger. Possible values are: 0 - (<i>default</i>) a plain trigger; 4 - a discovered trigger.
lastchange	timestamp	(<i>readonly</i>) Time when the trigger last changed its state.
priority	integer	Severity of the trigger. Possible values are: 0 - (<i>default</i>) not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.
status	integer	Whether the trigger is enabled or disabled. Possible values are: 0 - (<i>default</i>) enabled; 1 - disabled.
templateid	string	(<i>readonly</i>) ID of the parent template trigger.
type	integer	Whether the trigger can generate multiple problem events. Possible values are: 0 - (<i>default</i>) do not generate multiple events; 1 - generate multiple events.
url	string	URL associated with the trigger.
value	integer	(<i>readonly</i>) Whether the trigger is in OK or problem state. Possible values are: 0 - (<i>default</i>) OK; 1 - problem.
value_flags	integer	(<i>readonly</i>) Whether the trigger is currently in an unknown state. Possible values are: 0 - (<i>default</i>) trigger state is up to date; 1 - current trigger state is unknown.

trigger.adddependencies

Description

object trigger.adddependencies(object/array triggerDependencies)

This method allows to create new trigger dependencies.

Parameters

(object/array) Trigger dependencies to create.

Each trigger dependency has the following parameters:

Parameter	Type	Description
triggerid (required)	string	ID of the dependent trigger.
dependsOnTriggerid (required)	string	ID of the trigger that the trigger depends on.

Return values

(object) Returns an object containing the IDs of the dependent triggers under the `triggerids` property.

Examples

Add a trigger dependency

Make trigger "14092" dependent on trigger "13565."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.adddependencies",
  "params": {
    "triggerid": "14092",
    "dependsOnTriggerid": "13565"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "14092"
    ]
  },
  "id": 1
}
```

See also

- [trigger.update](#)

Source

CTrigger::addDependencies() in *frontends/php/api/classes/CTrigger.php*.

trigger.create

Description

object trigger.create(object/array triggers)

This method allows to create new triggers.

Parameters

(object/array) Triggers to create.

Additionally to the [standard trigger properties](#) the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Triggers that the trigger is dependent on. The triggers must have the <code>triggerid</code> property defined.

Attention:

The trigger expression has to be given in its expanded form.

Return values

(object) Returns an object containing the IDs of the created triggers under the `triggerids` property. The order of the returned IDs matches the order of the passed triggers.

Examples**Creating a trigger**

Create a trigger with a single trigger dependency.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.create",
  "params": {
    "description": "Processor load is too high on {HOST.NAME}",
    "expression": "{Linux server:system.cpu.load[percpu,avg1].last(0)}>5",
    "dependencies": [
      {
        "triggerid": "14062"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "14102"
    ]
  },
  "id": 1
}
```

Source

CTrigger::create() in *frontends/php/api/classes/CTrigger.php*.

trigger.delete**Description**

object trigger.delete(array triggerIds)

This method allows to delete triggers.

Parameters

(array) IDs of the triggers to delete.

Return values

(object) Returns an object containing the IDs of the deleted triggers under the `triggerids` property.

Examples**Delete multiple triggers**

Delete two triggers.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.delete",
  "params": [
    "12002",
    "12003"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "12002",
      "12003"
    ]
  },
  "id": 1
}
```

Source

CTrigger::delete() in *frontends/php/api/classes/CTrigger.php*.

trigger.deletedependencies

Description

object trigger.deletedependencies(string/array triggers)

This method allows to delete all trigger dependencies from the given triggers.

Parameters

(string/array) Triggers to delete the trigger dependencies from.

Return values

(object) Returns an object containing the IDs of the affected triggers under the `triggerids` property.

Examples

Deleting dependencies from multiple triggers

Delete all dependencies from two triggers.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.deleteDependencies",
  "params": [
    {
      "triggerid": "14544"
    },
    {
      "triggerid": "14545"
    }
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "14544",
      "14545"
    ]
  },
  "id": 1
}
```

See also

- [trigger.update](#)

Source

CTrigger::deleteDependencies() in *frontends/php/api/classes/CTrigger.php*.

trigger.exists

Description

`boolean trigger.exists(object filter)`

This method checks if at least one trigger that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
expression (required)	string	Exploded trigger expression.
host (required)	string/array	Technical names of the hosts the triggers must belong to.
hostid (required)	string/array	IDs of the hosts the triggers must belong to.
description node	string/array	Names of the triggers.
	string	Name of the node the triggers must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the triggers must belong to.

Note:

Only one of the three parameters is required: host, hostid or expression.

Return values

(boolean) Returns true if at least one trigger that matches the given filter criteria exists.

Examples

Check a trigger by expression

Check if a trigger with the given expression exists.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.exists",
  "params": {
    "expression": "{Linux server:vfs.file.cksum[/etc/passwd].diff(0)}>0"
  },
}
```

```

    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}

```

See also

- [trigger.isreadable](#)
- [trigger.iswritable](#)

Source

CTrigger::exists() in *frontends/php/api/classes/CTrigger.php*.

trigger.get

Description

integer/array trigger.get(object parameters)

The method allows to retrieve triggers according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
triggerids	string/array	Return only triggers with the given IDs.
groupids	string/array	Return only triggers that belong to hosts from the given host groups.
templateids	string/array	Return only triggers that belong to the given templates.
hostids	string/array	Return only triggers that belong to the given hosts.
itemids	string/array	Return only triggers that contain the given items.
applicationids	string/array	Return only triggers that contain items from the given applications.
functions	string/array	Return only triggers that use the given functions. Refer to the supported trigger functions page for a list of supported functions.
group	string	Return only triggers that belong to hosts from the host group with the given name.
host	string	Return only triggers that belong to host with the given name.
inherited	boolean	If set to <code>true</code> return only triggers inherited from a template.
templated	boolean	If set to <code>true</code> return only triggers that belong to templates.
monitored	flag	Return only enabled triggers that belong to monitored hosts and contain only enabled items.
active	flag	Return only enabled triggers that belong to monitored hosts.
maintenance	boolean	If set to <code>true</code> return only enabled triggers that belong to hosts in maintenance.
withUnacknowledgedEvents	flag	Return only triggers that have unacknowledged events.
withAcknowledgedEvents	flag	Return only triggers with all events acknowledged.

Parameter	Type	Description
withLastEventUnacknowledged	flag	Return only triggers with the last event unacknowledged.
skipDependent	flag	Skip triggers in a problem state that are dependent on other triggers. Note that the other triggers are ignored if disabled, have disabled items or disabled item hosts.
lastChangeSince	timestamp	Return only triggers that have changed their state after the given time.
lastChangeTill	timestamp	Return only triggers that have changed their state before the given time.
only_true	flag	Return only triggers that have recently been in a problem state.
min_severity	integer	Return only triggers with severity greater or equal than the given severity.
expandData	flag	Return additional data about the first host in the trigger expression.
expandDescription	flag	Adds the following properties to each trigger: hostname - (<i>string</i>) visible name of the host; host - (<i>string</i>) technical name of the host; hostid - (<i>string</i>) ID of the host. Expand macros in the name of the trigger.
expandExpression	flag	Expand functions and macros in the trigger expression.
selectGroups	query	Return the host groups that the trigger belongs to in the groups property.
selectHosts	query	Does not support shorten or arrays of properties. Return the hosts that the trigger belongs to in the hosts property.
selectItems	query	Does not support shorten. Return items contained by the trigger in the items property.
selectFunctions	query	Does not support shorten. Return functions used in the trigger in the functions property.
		The function objects represents the functions used in the trigger expression and has the following properties: functionid - (<i>string</i>) ID of the function; itemid - (<i>string</i>) ID of the item used in the function; function - (<i>string</i>) name of the function; parameter - (<i>string</i>) parameter passed to the function.
selectDependencies	query	Does not support shorten or arrays of properties. Return triggers that the trigger depends on in the dependencies property.
selectDiscoveryRule	query	Does not support shorten or arrays of properties. Return the low-level discovery rule that created the trigger.
selectLastEvent	query	Does not support shorten. (added in 2.0.6) Return the last significant trigger event in the lastEvent property.
		Does not support shorten or arrays of properties.

Parameter	Type	Description
filter	object	Return only those results that exactly match the given filter. Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Supports additional filters: host - technical name of the host that the trigger belongs to; hostid - ID of the host that the trigger belongs to. Limits the number of records returned by subselects.
limitSelects	integer	Applies to the following subselects: selectHosts - results will be sorted by host. Sort the result by the given properties.
sortfield	string/array	Possible values are: triggerid, description, status, priority, lastchange and hostname. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
countOutput	flag	
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving data by trigger ID

Retrieve all data and the functions used in trigger "14062."

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.get",
  "params": {
    "triggerids": "14062",
    "output": "extend",
    "selectFunctions": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
```

```

"result": [
  {
    "functions": [
      {
        "functionid": "13513",
        "itemid": "24350",
        "function": "diff",
        "parameter": "0"
      }
    ],
    "triggerid": "14062",
    "expression": "{13513}>0",
    "description": "/etc/passwd has been changed on {HOST.NAME}",
    "url": "",
    "status": "0",
    "value": "0",
    "priority": "2",
    "lastchange": "0",
    "comments": "",
    "error": "",
    "templateid": "10016",
    "type": "0",
    "value_flags": "0",
    "flags": "0"
  }
],
"id": 1
}

```

Retrieving triggers in problem state

Retrieve the ID, name and severity of all triggers in problem state and sort them by severity in descending order.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "trigger.get",
  "params": {
    "output": [
      "triggerid",
      "description",
      "priority"
    ],
    "filter": {
      "value": 1
    },
    "sortfield": "priority",
    "sortorder": "DESC"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": [
    {
      "triggerid": "13907",
      "description": "Zabbix self-monitoring processes < 100% busy",
      "priority": "4"
    },
    {

```



```

        "triggerid": "13824",
        "description": "Zabbix discoverer processes more than 75% busy",
        "priority": "3"
    }
],
    "id": 1
}

```

See also

- [trigger.getobjects](#)
- [Discovery rule](#)
- [Item](#)
- [Host](#)
- [Host group](#)

Source

CTrigger::get() in *frontends/php/api/classes/CTrigger.php*.

trigger.getobjects

Description

array trigger.getobjects(object filter)

This method allows to retrieve triggers that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard trigger properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
host	string/array	Technical name of the host that the trigger belongs to.
hostid	string/array	ID of the host that the trigger belongs to.
node	string	Name of the node the host group must belong to.
nodeids	string/array	This will override the nodeids parameter. ID of the node the host group must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving triggers by name

Retrieve triggers with the name `"/etc/passwd has been changed on {HOST.NAME}"` from two hosts.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "trigger.getobjects",
    "params": {
        "description": "/etc/passwd has been changed on {HOST.NAME}",
        "hostid": [
            "30069",
            "30049"
        ]
    },
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "triggerid": "13938",
      "expression": "{13385}>0",
      "description": "/etc/passwd has been changed on {HOST.NAME}",
      "url": "",
      "status": "0",
      "value": "0",
      "priority": "2",
      "lastchange": "0",
      "comments": "",
      "error": "Agent is unavailable.",
      "templateid": "10016",
      "type": "0",
      "value_flags": "1",
      "flags": "0"
    },
    {
      "triggerid": "14062",
      "expression": "{13513}>0",
      "description": "/etc/passwd has been changed on {HOST.NAME}",
      "url": "",
      "status": "0",
      "value": "0",
      "priority": "2",
      "lastchange": "0",
      "comments": "",
      "error": "",
      "templateid": "10016",
      "type": "0",
      "value_flags": "0",
      "flags": "0"
    }
  ],
  "id": 1
}
```

See also

- [trigger.get](#)

Source

CTrigger::getObject() in *frontends/php/api/classes/CTrigger.php*.

trigger.isreadable

Description

boolean trigger.isreadable(array triggerIds)

This method checks if the given triggers are available for reading.

Parameters

(array) IDs of the triggers to check.

Return values

(boolean) Returns true if the given triggers are available for reading.

Examples

Check multiple triggers

Check if the two triggers are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.isreadable",
  "params": [
    "13938",
    "14062"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [trigger.exists](#)
- [trigger.iswritable](#)

Source

CTrigger::isReadable() in *frontends/php/api/classes/CTrigger.php*.

trigger.iswritable

Description

`boolean trigger.iswritable(array triggerIds)`

This method checks if the given triggers are available for writing.

Parameters

(array) IDs of the triggers to check.

Return values

(boolean) Returns true if the given triggers are available for writing.

Examples

Check multiple triggers

Check if the two triggers are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.iswritable",
  "params": [
    "13938",
    "14062"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
}
```

```
}  
    "id": 1  
}
```

See also

- `trigger.isreadable`
- `trigger.exists`

Source

`CTrigger::isWritable()` in *frontends/php/api/classes/CTrigger.php*.

trigger.update

Description

`object trigger.update(object/array triggers)`

This method allows to update existing triggers.

Parameters

(object/array) Trigger properties to be updated.

The `triggerid` property must be defined for each trigger, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the **standard trigger properties** the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Triggers that the trigger is dependent on. The triggers must have the <code>triggerid</code> property defined.

Attention:

The trigger expression has to be given in its expanded form.

Return values

(object) Returns an object containing the IDs of the updated triggers under the `triggerids` property.

Examples

Enabling a trigger

Enable a trigger, that is, set its status to 0.

Request:

```
{  
    "jsonrpc": "2.0",  
    "method": "trigger.update",  
    "params": {  
        "triggerid": "13938",  
        "status": 0  
    },  
    "auth": "038e1d7b1735c6a5436ee9eae095879e",  
    "id": 1  
}
```

Response:

```
{  
    "jsonrpc": "2.0",  
    "result": {  
        "triggerids": [  
            "13938"  
        ]  
    }  
}
```

```

    ],
    },
    "id": 1
}

```

See also

- [trigger.adddependencies](#)
- [trigger.deletedependencies](#)

Source

CTrigger::update() in *frontends/php/api/classes/CTrigger.php*.

Trigger prototype

This class is designed to work with trigger prototypes.

Object references:

- [Trigger prototype](#)

Available methods:

- [triggerprototype.create](#) - creating new trigger prototypes
- [triggerprototype.delete](#) - deleting trigger prototypes
- [triggerprototype.get](#) - retrieving trigger prototypes
- [triggerprototype.update](#) - updating trigger prototypes

> Trigger prototype object

The following objects are directly related to the triggerprototype API.

Trigger

The trigger prototype object has the following properties.

Property	Type	Description
triggerid	string	(<i>readonly</i>) ID of the trigger prototype.
description (required)	string	Name of the trigger prototype.
expression (required)	string	Reduced trigger expression.
comments	string	Additional comments to the trigger prototype.
priority	integer	Severity of the trigger prototype. Possible values: 0 - (<i>default</i>) not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.
status	integer	Whether the trigger prototype is enabled or disabled. Possible values: 0 - (<i>default</i>) enabled; 1 - disabled.
templateid	string	(<i>readonly</i>) ID of the parent template trigger prototype.

Property	Type	Description
type	integer	Whether the trigger prototype can generate multiple problem events. Possible values: 0 - (<i>default</i>) do not generate multiple events; 1 - generate multiple events.
url	string	URL associated with the trigger prototype.

triggerprototype.create

Description

object triggerprototype.create(object/array triggerPrototypes)

This method allows to create new trigger prototypes.

Parameters

(object/array) Trigger prototypes to create.

The method accepts trigger prototypes with the **standard trigger prototype properties**.

Attention:

The trigger expression has to be given in its expanded form and must contain at least one item prototype.

Return values

(object) Returns an object containing the IDs of the created trigger prototypes under the `triggerids` property. The order of the returned IDs matches the order of the passed trigger prototypes.

Examples

Creating a trigger prototype

Create a trigger prototype to detect when a file system has less than 20% free disk space.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "triggerprototype.create",
  "params": {
    "description": "Free disk space is less than 20% on volume {#FSNAME}",
    "expression": "{Zabbix server:vfs.fs.size[{#FSNAME},pfree].last(0)}<20"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "15331"
    ]
  },
  "id": 1
}
```

Source

CTriggerPrototype::create() in *frontends/php/api/classes/CTriggerPrototype.php*.

triggerprototype.delete

Description

object triggerprototype.delete(array triggerPrototypeIds)

This method allows to delete trigger prototypes.

Parameters

(array) IDs of the trigger prototypes to delete.

Return values

(object) Returns an object containing the IDs of the deleted trigger prototypes under the triggerids property.

Examples

Deleting multiple trigger prototypes

Delete two trigger prototypes.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "triggerprototype.delete",
  "params": [
    "12002",
    "12003"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "12002",
      "12003"
    ]
  },
  "id": 1
}
```

Source

CTriggerPrototype::delete() in *frontends/php/api/classes/CTriggerPrototype.php*.

triggerprototype.get

Description

integer/array triggerprototype.get(object parameters)

The method allows to retrieve trigger prototypes according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
active	flag	Return only enabled trigger prototypes that belong to monitored hosts.
applicationids	string/array	Return only trigger prototypes that contain items from the given applications.

Parameter	Type	Description
discoveryids	string/array	Return only trigger prototypes that belong to the given LLD rules.
functions	string/array	Return only triggers that use the given functions. Refer to the supported trigger functions page for a list of supported functions.
group	string	Return only trigger prototypes that belong to hosts from the host groups with the given name.
groupids	string/array	Return only trigger prototypes that belong to hosts from the given host groups.
host	string	Return only trigger prototypes that belong to hosts with the given name.
hostids	string/array	Return only trigger prototypes that belong to the given hosts.
inherited	boolean	If set to <code>true</code> return only trigger prototypes inherited from a template.
maintenance	boolean	If set to <code>true</code> return only enabled trigger prototypes that belong to hosts in maintenance.
min_severity	integer	Return only trigger prototypes with severity greater or equal than the given severity.
monitored	flag	Return only enabled trigger prototypes that belong to monitored hosts and contain only enabled items.
templated	boolean	If set to <code>true</code> return only trigger prototypes that belong to templates.
templateids	string/array	Return only trigger prototypes that belong to the given templates.
triggerids	string/array	Return only trigger prototypes with the given IDs.
expandData	flag	Return additional data about the first host in the trigger expression.
expandExpression	flag	Adds the following properties to each trigger prototype: <code>hostname</code> - <i>(string)</i> visible name of the host; <code>host</code> - <i>(string)</i> technical name of the host; <code>hostid</code> - <i>(string)</i> ID of the host. <i>(added in 2.0.3)</i> Expand functions and macros in the trigger expression.
selectDiscoveryRule	query	Return the LLD rule that the trigger prototype belongs to.
selectFunctions	query	Does not support <code>shorten</code> . Return functions used in the trigger prototype in the <code>functions</code> property.
selectGroups	query	The function objects represents the functions used in the trigger expression and has the following properties: <code>functionid</code> - <i>(string)</i> ID of the function; <code>itemid</code> - <i>(string)</i> ID of the item used in the function; <code>function</code> - <i>(string)</i> name of the function; <code>parameter</code> - <i>(string)</i> parameter passed to the function.
selectHosts	query	Does not support <code>shorten</code> or arrays of properties. Return the hosts that the trigger prototype belongs to in the <code>hosts</code> property.
		Does not support <code>shorten</code> .

Parameter	Type	Description
selectItems	query	Return items and item prototypes used the trigger prototype in the <code>items</code> property.
filter	object	Does not support <code>shorten</code>. Return only those results that exactly match the given filter. Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Supports additional filters: <code>host</code> - technical name of the host that the trigger prototype belongs to; <code>hostid</code> - ID of the host that the trigger prototype belongs to.
limitSelects	integer	Limits the number of records returned by subselects.
sortfield	string/array	Applies to the following subselects: <code>selectHosts</code> - results will be sorted by <code>host</code>. Sort the result by the given properties. Possible values are: <code>triggerid</code>, <code>description</code>, <code>status</code> and <code>priority</code>.
countOutput	flag	These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve trigger prototypes from an LLD rule

Retrieve all trigger prototypes and their functions from an LLD rule.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "triggerprototype.get",
  "params": {
    "output": "extend",
    "selectFunctions": "extend",
    "discoveryids": "22450"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "functions": [
        {
          "functionid": "12598",
          "itemid": "22454",
          "function": "last",
          "parameter": "0"
        }
      ],
      "triggerid": "13272",
      "expression": "{12598}<20",
      "description": "Free inodes is less than 20% on volume {#FSNAME}",
      "url": "",
      "status": "0",
      "priority": "2",
      "comments": "",
      "templateid": "0",
      "type": "0",
      "flags": "2"
    },
    {
      "functions": [
        {
          "functionid": "13500",
          "itemid": "22686",
          "function": "last",
          "parameter": "0"
        }
      ],
      "triggerid": "13266",
      "expression": "{13500}<201",
      "description": "Free disk space is less than 20% on volume {#FSNAME}",
      "url": "",
      "status": "0",
      "priority": "2",
      "comments": "",
      "templateid": "0",
      "type": "0",
      "flags": "2"
    }
  ],
  "id": 1
}
```

See also

- [Discovery rule](#)
- [Item](#)
- [Host](#)
- [Host group](#)

Source

CTriggerPrototype::get() in *frontends/php/api/classes/CTriggerPrototype.php*.

triggerprototype.update

Description

`object triggerprototype.update(object/array triggerPrototypes)`

This method allows to update existing trigger prototypes.

Parameters

(object/array) **Trigger prototype properties** to be updated.

The `triggerid` property must be defined for each trigger prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Attention:

The trigger expression has to be given in its expanded form and must contain at least one item prototype.

Return values

(object) Returns an object containing the IDs of the updated trigger prototypes under the `triggerids` property.

Examples

Enabling a trigger prototype

Enable a trigger prototype, that is, set its status to 0.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "triggerprototype.update",
  "params": {
    "triggerid": "13938",
    "status": 0
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "triggerids": [
      "13938"
    ]
  },
  "id": 1
}
```

Source

`CTriggerPrototype::update()` in *frontends/php/api/classes/CTriggerPrototype.php*.

User

This class is designed to work with users.

Object references:

- **User**

Available methods:

- **user.addmedia** - adding media to users
- **user.create** - creating new users
- **user.delete** - deleting users
- **user.deletemedia** - deleting media from users

- **user.get** - retrieving users
- **user.isreadable** - checking if users are readable
- **user.iswritable** - checking if users are writable
- **user.login** - logging in to the API
- **user.logout** - logging out of the API
- **user.update** - updating users
- **user.updatemedia** - updating user media
- **user.updateprofile** - updating the currently logged in user

> User object

The following objects are directly related to the `user` API.

User

The user object has the following properties.

Property	Type	Description
<code>userid</code>	string	(<i>readonly</i>) ID of the user.
<code>alias</code> (<i>required</i>)	string	User alias.
<code>attempt_clock</code>	timestamp	(<i>readonly</i>) Time of the last unsuccessful login attempt.
<code>attempt_failed</code>	integer	(<i>readonly</i>) Recent failed login attempt count.
<code>attempt_ip</code>	string	(<i>readonly</i>) IP address from where the last unsuccessful login attempt came from.
<code>autologin</code>	integer	Whether to enable auto-login. Possible values: 0 - (<i>default</i>) auto-login disabled; 1 - auto-login enabled.
<code>autologout</code>	integer	User session life time in seconds. If set to 0, the session will never expire. Default: 900.
<code>lang</code>	string	Language code of the user's language.
<code>name</code>	string	Default: <code>en_GB</code> . Name of the user.
<code>refresh</code>	integer	Automatic refresh period in seconds.
<code>rows_per_page</code>	integer	Default: 30. Amount of object rows to show per page.
<code>surname</code>	string	Default: 50. Surname of the user.
<code>theme</code>	string	User's theme. Possible values: <code>default</code> - (<i>default</i>) system default; <code>classic</code> - Classic; <code>originalblue</code> - Original blue; <code>darkblue</code> - Black & Blue; <code>darkorange</code> - Dark orange.
<code>type</code>	integer	Type of the user. Possible values: 1 - (<i>default</i>) Zabbix user; 2 - Zabbix admin; 3 - Zabbix super admin.
<code>url</code>	string	URL of the page to redirect the user to after logging in.

user.addmedia

Description

`object user.addmedia(object parameters)`

This method allows to add new media to multiple users.

Parameters

(object) Parameters defining the media to create and the users to add them to.

Parameter	Type	Description
medias (required)	object/array	Media to create for the given users.
users (required)	object/array	The media <code>userid</code> property must not be defined. Users to add the media to.
		The users must have the <code>userid</code> property defined.

Return values

(object) Returns an object containing the IDs of the created media under the `mediaids` property.

Examples

Adding a media to multiple users

Create a common e-mail media for two users. The media must send notifications about all alerts at any time.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.addmedia",
  "params": {
    "users": [
      {
        "userid": "1"
      },
      {
        "userid": "2"
      }
    ],
    "medias": {
      "mediatypeid": "1",
      "sendto": "support@company.com",
      "active": 0,
      "severity": 63,
      "period": "1-7,00:00-24:00"
    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "mediaids": [
      "12",
      "13"
    ]
  },
  "id": 1
}
```

See also

- [user.update](#)
- [user.updatemedia](#)
- [Media](#)
- [User](#)

Source

CUser::addMedia() in *frontends/php/api/classes/CUser.php*.

user.authenticate

Warning:

This method is a deprecated alias of [user.login](#).

user.create

Description

`object user.create(object/array users)`

This method allows to create new users.

Parameters

(object/array) Users to create.

Additionally to the [standard user properties](#), the method accepts the following parameters.

Parameter	Type	Description
passwd (required)	string	User's password.
usrgrps (required)	array	User groups to add the user to. The user groups must have the <code>usrgrpid</code> property defined.
<code>user_medias</code>	array	Media to create for the user. The media <code>userid</code> property must not be defined.

Return values

(object) Returns an object containing the IDs of the created users under the `userids` property. The order of the returned IDs matches the order of the passed users.

Examples

Creating a user

Create a new user, add him to a user group and create a new media for him.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.create",
  "params": {
    "alias": "John",
    "passwd": "Doe123",
    "usrgrps": [
      {
        "usrgrpid": "7"
      }
    ],
    "user_medias": [
      {
        "mediatypeid": "1",
```

```

        "sendto": "support@company.com",
        "active": 0,
        "severity": 63,
        "period": "1-7,00:00-24:00"
    }
]
},
"auth": "038e1d7b1735c6a5436ee9eae095879e",
"id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "12"
        ]
    },
    "id": 1
}

```

See also

- [Media](#)
- [User group](#)

Source

CUser::create() in *frontends/php/api/classes/CUser.php*.

user.delete

Description

object user.delete(array users)

This method allows to delete users.

Parameters

(array) Users to delete.

The users must have the `userid` property defined.

Return values

(object) Returns an object containing the IDs of the deleted users under the `userids` property.

Examples

Deleting multiple users

Delete two users.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "user.delete",
    "params": [
        {"userid": "1"},
        {"userid": "5"}
    ],
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "userids": [
      "1",
      "5"
    ]
  },
  "id": 1
}
```

Source

CUser::delete() in *frontends/php/api/classes/CUser.php*.

user.deletemedia

Description

object user.deletemedia(string/array mediaIds)

This method allows to delete media.

Parameters

(string/array) IDs of the media to delete.

Return values

(object) Returns an object containing the IDs of the deleted media under the `mediaids` property.

Examples

Deleting multiple media

Delete two media.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.deletemedia",
  "params": [
    "11",
    "13"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "mediaids": [
      "11",
      "13"
    ]
  },
  "id": 1
}
```

See also

- [user.update](#)
- [user.updatemedia](#)

Source

CUser::deleteMedia() in *frontends/php/api/classes/CUser.php*.

user.get

Description

integer/array user.get(object parameters)

The method allows to retrieve users according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
mediaids	string/array	Return only users that use the given media.
mediatypesids	string/array	Return only users that use the given media types.
userids	string/array	Return only users with the given IDs.
usrgrpsids	string/array	Return only users that belong to the given user groups.
getAccess	flag	Adds additional information about user permissions. Adds the following properties for each user: gui_access - (<i>integer</i>) user's frontend authentication method. Refer to the gui_access property of the user group object for a list of possible values. debug_mode - (<i>integer</i>) indicates whether debug is enabled for the user. Possible values: 0 - debug disabled, 1 - debug enabled. users_status - (<i>integer</i>) indicates whether the user is disabled. Possible values: 0 - user enabled, 1 - user disabled.
selectMedias	query	Return media used by the user in the medias property.
selectMediatypes	query	Does not support shorten or arrays of properties. (<i>added in 2.0.1</i>) Return media types used by the user in the mediatypes property.
selectUsrgrps	query	Does not support shorten or arrays of properties. Return user groups that the user belongs to in the usrgrps property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: userid and alias. These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving users

Retrieve all of the configured users.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.get",
  "params": {
    "output": "extend"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "userid": "1",
      "alias": "Admin",
      "name": "Zabbix",
      "surname": "Administrator",
      "url": "",
      "autologin": "1",
      "autologout": "0",
      "lang": "ru_RU",
      "refresh": "0",
      "type": "3",
      "theme": "default",
      "attempt_failed": "0",
      "attempt_ip": "",
      "attempt_clock": "0",
      "rows_per_page": "50"
    },
    {
      "userid": "2",
      "alias": "guest",
      "name": "Default2",
      "surname": "User",
      "url": "",
      "autologin": "0",
      "autologout": "900",
      "lang": "en_GB",
      "refresh": "30",
      "type": "1",
      "theme": "default",
      "attempt_failed": "0",
      "attempt_ip": "",
      "attempt_clock": "0",
      "rows_per_page": "50"
    }
  ],
  "id": 1
}
```

See also

- [Media](#)
- [Media type](#)
- [User group](#)

Source

CUser::get() in *frontends/php/api/classes/CUser.php*.

user.isreadable

Description

`boolean user.isreadable(array userIds)`

This method checks if the given users are available for reading.

Parameters

(array) IDs of the users to check.

Return values

(boolean) Returns true if the given users are available for reading.

Examples

Check multiple users

Check if the two users are readable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.isreadable",
  "params": [
    "4",
    "6"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [user.iswritable](#)

Source

CUser::isReadable() in *frontends/php/api/classes/CUser.php*.

user.iswritable

Description

`boolean user.iswritable(array userIds)`

This method checks if the given users are available for writing.

Parameters

(array) IDs of the users to check.

Return values

(boolean) Returns true if the given users are available for writing.

Examples

Check multiple users

Check if the two users are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.iswritable",
  "params": [
    "4",
    "6"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [user.isreadable](#)

Source

CUser::isWritable() in *frontends/php/api/classes/CUser.php*.

user.login

Description

string/object user.login(object parameters)

This method allows to log in to the API and generate an authentication token.

Parameters

Attention:
This method is available to unauthenticated users and should be called without the `auth` parameter in the JSON-RPC request. Starting from Zabbix 2.4 the method will return an error if the `auth` parameter is given.

(object) Parameters containing the user name and password.

The method accepts the following parameters.

Parameter	Type	Description
password (required)	string	User password. Unused for HTTP authentication.
user (required)	string	User name.
userData	flag	Return information about the authenticated user.

Attention:
When using HTTP authentication, the user name in the API request must match the one used in the `Authorization` header. The password will not be validated and can be omitted.

Return values

(string/object) If the `userData` parameter is used, returns an object containing information about the authenticated user.

Additionally to the **standard user properties**, the following information is returned:

Property	Type	Description
<code>debug_mode</code>	boolean	Whether debug mode is enabled for the user.
<code>gui_access</code>	integer	User's authentication method to the frontend.
<code>node</code>	object	Refer to the <code>gui_access</code> property of the user group object for a list of possible values. Local node of the user.
<code>sessionid</code>	string	The object has the following properties: <code>name</code> - (string) Name of the node; <code>nodeid</code> - (string) ID of the node. Authentication token, which must be used in the following API requests.
<code>userip</code>	string	IP address of the user.

Note:

If a user has been successfully authenticated after one or more failed attempts, the method will return the current values for the `attempt_clock`, `attempt_failed` and `attempt_ip` properties and then reset them.

If the `userData` parameter is not used, the method returns an authentication token.

Note:

The generated authentication token should be remembered and used in the `auth` parameter of the following JSON-RPC requests. It is also required when using HTTP authentication.

Examples

Authenticating a user

Authenticate a user.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.login",
  "params": {
    "user": "Admin",
    "password": "zabbix"
  },
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": "0424bd59b807674191e7d77572075f33",
  "id": 1
}
```

Requesting authenticated user's information

Authenticate and return additional information about the user.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.login",
  "params": {
    "user": "Admin",
```

```

        "password": "zabbix",
        "userData": true
    },
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "userid": "1",
        "alias": "Admin",
        "name": "Zabbix",
        "surname": "Administrator",
        "url": "",
        "autologin": "1",
        "autologout": "0",
        "lang": "ru_RU",
        "refresh": "0",
        "type": "3",
        "theme": "default",
        "attempt_failed": "0",
        "attempt_ip": "127.0.0.1",
        "attempt_clock": "1355919038",
        "rows_per_page": "50",
        "debug_mode": true,
        "userip": "127.0.0.1",
        "node": {
            "name": "- unknown -",
            "nodeid": null
        },
        "sessionid": "5b56eee8be445e98f0bd42b435736e42",
        "gui_access": "0"
    },
    "id": 1
}

```

See also

- [user.logout](#)

Source

CUser::login() in *frontends/php/api/classes/CUser.php*.

user.logout

Description

string/object `user.logout(array)`

This method allows to log out of the API and invalidates the current authentication token.

Parameters

(array) The method accepts an empty array.

Return values

(boolean) Returns true if the user has been logged out successfully.

Examples

Logging out

Log out from the API.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.logout",
  "params": [],
  "id": 1,
  "auth": "16a46baf181ef9602e1687f3110abf8a"
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [user.login](#)

Source

CUser::login() in *frontends/php/api/classes/CUser.php*.

user.update

Description

`object user.update(object/array users)`

This method allows to update existing users.

Parameters

(object/array) User properties to be updated.

The `userid` property must be defined for each user, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard user properties](#), the method accepts the following parameters.

Parameter	Type	Description
passwd	string	User's password.
usrgrps	array	User groups to replace existing user groups.
		The user groups must have the <code>usrgrp_id</code> property defined.

Return values

(object) Returns an object containing the IDs of the updated users under the `userids` property.

Examples

Renaming a user

Rename a user to John Doe.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.update",
  "params": {
    "userid": "1",
    "name": "John",
    "surname": "Doe"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "userids": [
      "1"
    ]
  },
  "id": 1
}
```

See also

- [user.updateprofile](#)

Source

CUser::update() in *frontends/php/api/classes/CUser.php*.

user.updatemedia

Description

object user.updatemedia(object parameters)

This method allows to update media for multiple users.

Parameters

(object) Parameters defining the media and users to be updated.

Parameter	Type	Description
medias (required)	object/array	Media to replace existing media. If a media has the mediaid property defined it will be updated, otherwise a new media will be created.
users (required)	object/array	Users to update. The users must have the userid property defined.

Return values

(object) Returns an object containing the IDs of the updated users under the userids property.

Examples

Replacing media for multiple users

Replace all media used by the two users with a common e-mail media. The media must send notifications about all alerts at any time.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.updatemedia",
  "params": {
    "users": [
      {
        "userid": "1"
      },
      {
        "userid": "2"
      }
    ]
  },
}
```



```

        "medias": {
            "mediatypeid": "1",
            "sendto": "support@company.com",
            "active": 0,
            "severity": 63,
            "period": "1-7,00:00-24:00"
        }
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "1",
            "2"
        ]
    },
    "id": 1
}

```

See also

- [user.addmedia](#)
- [user.deletemedia](#)
- [user.updatemedia](#)
- [Media](#)
- [User](#)

Source

CUser::updateMedia() in *frontends/php/api/classes/CUser.php*.

user.updateprofile

Description

`object user.updateprofile(object parameters)`

This method allows to update the currently logged in user.

Parameters

(object/array) User properties to be updated.

The `userid` property must not be defined. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard user properties](#), the method accepts the following parameters.

Parameter	Type	Description
passwd	string	User's password.
usrgrps	array	User groups to replace existing user groups.
		The user groups must have the <code>usrgrpid</code> property defined.

Return values

(object) Returns an object containing the ID of the updated user under the `userids` property.

Examples

Renaming the current user

Rename the current user to John Doe.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "user.updateprofile",
  "params": {
    "name": "John",
    "lastname": "Doe"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "userids": [
      "1"
    ]
  },
  "id": 1
}
```

See also

- [user.update](#)

Source

CUser::update() in *frontends/php/api/classes/CUser.php*.

User group

This class is designed to work with user groups.

Object references:

- [User group](#)

Available methods:

- [usergroup.create](#) - creating new user groups
- [usergroup.delete](#) - deleting user groups
- [usergroup.exists](#) - checking if a user group exists
- [usergroup.get](#) - retrieving user groups
- [usergroup.getobjects](#) - retrieving user groups by filters
- [usergroup.isreadable](#) - checking if user groups are readable
- [usergroup.iswritable](#) - checking if user groups are writable
- [usergroup.massadd](#) - adding permissions and users to user groups
- [usergroup.massupdate](#) - simultaneously updating multiple user groups
- [usergroup.update](#) - updating user groups

> User group object

The following objects are directly related to the usergroup API.

User group

The user group object has the following properties.

Property	Type	Description
usrgrpid	string	(<i>readonly</i>) ID of the user group.

Property	Type	Description
name (required)	string	Name of the user group.
debug_mode	integer	Whether debug mode is enabled or disabled.
gui_access	integer	Possible values are: 0 - <i>(default)</i> disabled; 1 - enabled. Frontend authentication method of the users in the group.
users_status	integer	Possible values: 0 - <i>(default)</i> use the system default authentication method; 1 - use internal authentication; 2 - disable access to the frontend. Whether the user group is enabled or disabled.
		Possible values are: 0 - <i>(default)</i> enabled; 1 - disabled.

Permission

The permission object has the following properties.

Property	Type	Description
id (required)	string	ID of the host group to add permission to.
permission (required)	integer	Access level to the host group. Possible values: 0 - access denied; 2 - read-only access; 3 - read-write access.

usergroup.create

Description

`object usergroup.create(object/array userGroups)`

This method allows to create new user groups.

Parameters

(object/array) User groups to create.

Additionally to the **standard user group properties**, the method accepts the following parameters.

Parameter	Type	Description
rights	object/array	Permissions to assign to the group
userids	string/array	IDs of users to add to the user group.

Return values

(object) Returns an object containing the IDs of the created user groups under the `usrgrpids` property. The order of the returned IDs matches the order of the passed user groups.

Examples

Creating a user group

Create a user group, which denies access to host group "2", and add a user to it.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.create",
  "params": {
    "name": "Operation managers",
    "rights": {
      "permission": 0,
      "id": "2"
    },
    "userids": "12"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "usrgrpids": [
      "20"
    ]
  },
  "id": 1
}
```

See also

- [Permission](#)

Source

CUserGroup::create() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.delete

Description

object usergroup.delete(array userGroupIds)

This method allows to delete user groups.

Parameters

(array) IDs of the user groups to delete.

Return values

(object) Returns an object containing the IDs of the deleted user groups under the `usrgrpids` property.

Examples

Deleting multiple user groups

Delete two user groups.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.delete",
  "params": [
    "20",
    "21"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "usrgrpids": [
      "20",
      "21"
    ]
  },
  "id": 1
}
```

Source

CUserGroup::delete() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.exists

Description

boolean usergroup.exists(object filter)

This method checks if at least one user group that matches the given filter criteria exists.

Parameters

(object) Criteria to search by.

The following parameters are supported as search criteria.

Parameter	Type	Description
name	string/array	Names of the user groups.
node	string	Name of the node the user groups must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the user groups must belong to.

Return values

(boolean) Returns true if at least one user group that matches the given filter criteria exists.

Examples

Checking if a user group exists

Check if user group "Zabbix administrators"

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.exists",
  "params": {
    "name": "Zabbix administrators"
  },
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
}
```

```

    "id": 1
}

```

See also

- [usergroup.isreadable](#)
- [usergroup.iswritable](#)

Source

CUserGroup::exists() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.get

Description

`integer/array usergroup.get(object parameters)`

The method allows to retrieve user groups according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
status	integer	Return only user groups with the given status. Refer to the user group page for a list of supported statuses.
userids	string/array	Return only user groups that contain the given users.
usrgrpids	string/array	Return only user groups with the given IDs.
with_gui_access	integer	Return only user groups with the given frontend authentication method. Refer to the user group page for a list of supported methods.
selectUsers	query	Return the users from the user group in the users property. Does not support <code>shorten</code> or arrays of properties.
limitSelects	integer	Limits the number of records returned by subselects.
sortfield	string/array	Sort the result by the given properties. Possible values are: <code>usrgrp_id</code> , <code>name</code> .
countOutput	flag	These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;

- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving enabled user groups

Retrieve all enabled user groups.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.get",
  "params": {
    "output": "extend",
    "status": 0
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "usrgrpid": "7",
      "name": "Zabbix administrators",
      "gui_access": "0",
      "users_status": "0",
      "debug_mode": "1"
    },
    {
      "usrgrpid": "8",
      "name": "Guests",
      "gui_access": "0",
      "users_status": "0",
      "debug_mode": "0"
    },
    {
      "usrgrpid": "11",
      "name": "Enabled debug mode",
      "gui_access": "0",
      "users_status": "0",
      "debug_mode": "1"
    },
    {
      "usrgrpid": "12",
      "name": "No access to the frontend",
      "gui_access": "2",
      "users_status": "0",
      "debug_mode": "0"
    },
    {
      "usrgrpid": "14",
      "name": "Read only",
      "gui_access": "0",
      "users_status": "0",
      "debug_mode": "0"
    },
    {
      "usrgrpid": "18",
      "name": "Deny",
      "gui_access": "0",
      "users_status": "0",

```

```

        "debug_mode": "0"
    }
],
    "id": 1
}

```

See also

- [usergroup.getobjects](#)
- [User](#)

Source

CUserGroup::get() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.getobjects

Description

array usergroup.getobjects(object filter)

This method allows to retrieve user groups that match the given filter criteria.

Parameters

(object) Criteria to search by.

Additionally to the standard [standard user group properties](#) the following parameters are supported as search criteria.

Parameter	Type	Description
name	string	Name of the user group.
node	string	Name of the node the user groups must belong to.
nodeids	string/array	This will override the nodeids parameter. IDs of the nodes the user groups must belong to.

Return values

(array) Returns an array of objects with all properties.

Examples

Retrieving a user group by name

Retrieve all data about the user group "Zabbix administrators".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "usergroup.getobjects",
    "params": {
        "name": "Zabbix administrators"
    },
    "auth": "3a57200802b24cda67c4e4010b50c065",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": [
        {
            "usrgrp_id": "7",
            "name": "Zabbix administrators",
            "gui_access": "0",
            "users_status": "0",
            "debug_mode": "1"
        }
    ]
}

```



```

    }
  ],
  "id": 1
}

```

See also

- [usergroup.get](#)

Source

CUserGroup::getObject() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.isreadable

Description

boolean usergroup.isreadable(array userGroupIds)

This method checks if the given user groups are available for reading.

Parameters

(array) IDs of the user groups to check.

Return values

(boolean) Returns true if the given user groups are available for reading.

Examples

Check multiple user groups

Check if the two user groups are readable.

Request:

```

{
  "jsonrpc": "2.0",
  "method": "usergroup.isreadable",
  "params": [
    "21",
    "22"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}

```

Response:

```

{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}

```

See also

- [usergroup.exists](#)
- [usergroup.iswritable](#)

Source

CUserGroup::isReadable() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.iswritable

Description

boolean usergroup.iswritable(array userGroupIds)

This method checks if the given user groups are available for writing.

Parameters

(array) IDs of the user groups to check.

Return values

(boolean) Returns true if the given user groups are available for writing.

Examples

Check multiple user groups

Check if the two user groups are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.iswritable",
  "params": [
    "21",
    "22"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [usergroup.isreadable](#)
- [usergroup.exists](#)

Source

CUserGroup::isWritable() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.massadd

Description

object usergroup.massadd(object parameters)

This method allows to simultaneously add permissions and users to multiple user groups.

Parameters

(object) Parameters containing the IDs of the user groups to update and the permissions and users to add.

The method accepts the following parameters.

Parameter	Type	Description
usrgrpids (required)	string/array	IDs of user groups to update.
rights	object/array	Permissions to assign to the user groups.
usersids	string/array	IDs of the users to add to the user groups.

Return values

(object) Returns an object containing the IDs of the updated user groups under the `usrgrpids` property.

Examples

Denying access to host group

Deny two user groups access to host group "2".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.massadd",
  "params": {
    "usrgrpids": [
      "17",
      "19"
    ],
    "rights": {
      "permission": 0,
      "id": "2"
    }
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "usrgrpids": [
      "17",
      "19"
    ]
  },
  "id": 1
}
```

See also

- [Permission](#)
- [usergroup.massupdate](#)
- [usergroup.update](#)

Source

CUserGroup::massAdd() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.massupdate

Description

object usergroup.massupdate(object parameters)

This method allows to simultaneously update properties, users or permissions for multiple user groups.

Parameters

(object) Parameters containing the IDs of the user groups to update and the properties that should be updated.

Additionally to the [standard user group properties](#), the method accepts the following parameters.

Parameter	Type	Description
usrgrpids (required)	string/array	IDs of user groups to update.
rights	string/array	Permissions to replace the current permissions assigned to the user group.
usersids	object/array	IDs of the users to replace the users in the group.

Return values

(object) Returns an object containing the IDs of the updated user groups under the `usrgrpids` property.

Examples

Changing permissions for a user group

Update the permissions for two user groups to only allow read-write access to two host groups.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.massupdate",
  "params": {
    "usrgrpids": [
      "17",
      "19"
    ],
    "rights": [
      {
        "permission": 3,
        "id": "2"
      },
      {
        "permission": 3,
        "id": "3"
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "usrgrpids": [
      "17",
      "19"
    ]
  },
  "id": 1
}
```

See also

- [Permission](#)
- [usergroup.massadd](#)
- [usergroup.update](#)

Source

CUserGroup::massUpdate() in *frontends/php/api/classes/CUserGroup.php*.

usergroup.update

Description

object usergroup.update(object/array userGroups)

This method allows to update existing user groups.

Parameters

(object/array) User group properties to be updated.

The `usrgrpId` property must be defined for each user group, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the [standard user group properties](#), the method accepts the following parameters.

Parameter	Type	Description
rights	object/array	Permissions to replace the current permissions assigned to the user group.
userids	string/array	IDs of the users to replace the users in the group.

Return values

(object) Returns an object containing the IDs of the updated user groups under the `usrgrpids` property.

Examples

Disabling a user group

Disable a user group.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usergroup.update",
  "params": {
    "usrgrp_id": "17",
    "users_status": "1"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "usrgrpids": [
      "17"
    ]
  },
  "id": 1
}
```

See also

- [Permission](#)
- [usergroup.massadd](#)
- [usergroup.massupdate](#)

Source

`CUserGroup::update()` in *frontends/php/api/classes/CUserGroup.php*.

User macro

This class is designed to work with host and global macros.

Object references:

- [Global macro](#)
- [Host macro](#)

Available methods:

- [usermacro.create](#) - creating new host macros
- [usermacro.createglobal](#) - creating new global macros
- [usermacro.delete](#) - deleting host macros
- [usermacro.deleteglobal](#) - deleting global macros
- [usermacro.get](#) - retrieving host and global macros
- [usermacro.update](#) - updating host macros

- **usermacro.updateglobal** - updating global macros

> User macro object

The following objects are directly related to the usermacro API.

Global macro

The global macro object has the following properties.

Property	Type	Description
globalmacroid	string	(readonly) ID of the global macro.
macro (required)	string	Macro string.
value (required)	string	Value of the macro.

Host macro

The host macro object defines a macro available on a host or template. It has the following properties.

Property	Type	Description
hostmacroid	string	(readonly) ID of the host macro.
hostid (required)	string	ID of the host that the macro belongs to.
macro (required)	string	Macro string.
value (required)	string	Value of the macro.

usermacro.create

Description

`object usermacro.create(object/array hostMacros)`

This method allows to create new host macros.

Parameters

(object/array) Host macros to create.

The method accepts host macros with the **standard host macro properties**.

Return values

(object) Returns an object containing the IDs of the created host macros under the `hostmacroids` property. The order of the returned IDs matches the order of the passed host macros.

Examples

Creating a host macro

Creat a host macro "{\$SNMP_COMMUNITY}" with the value "public" on host "10198".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usermacro.create",
  "params": {
    "hostid": "10198",
    "macro": "{$SNMP_COMMUNITY}",
    "value": "public"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
}
```

```
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostmacroids": [
      "11"
    ]
  },
  "id": 1
}
```

Source

CUserMacro::create() in *frontends/php/api/classes/CUserMacro.php*.

usermacro.createglobal

Description

object usermacro.createglobal(object/array globalMacros)

This method allows to create new global macros.

Parameters

(object/array) Global macros to create.

The method accepts global macros with the **standard global macro properties**.

Return values

(object) Returns an object containing the IDs of the created global macros under the `globalmacroids` property. The order of the returned IDs matches the order of the passed global macros.

Examples

Creating a global macro

Create a global macro "{ \$SNMP_COMMUNITY }" with value "public".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usermacro.createglobal",
  "params": {
    "macro": "{ $SNMP_COMMUNITY }",
    "value": "public"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "globalmacroids": [
      "6"
    ]
  },
  "id": 1
}
```

Source

CUserMacro::createGlobal() in *frontends/php/api/classes/CUserMacro.php*.

usermacro.delete

Description

object usermacro.delete(array hostMacroIds)

This method allows to delete host macros.

Parameters

(array) IDs of the host macros to delete.

Return values

(object) Returns an object containing the IDs of the deleted host macros under the `hostmacroids` property.

Examples

Deleting multiple host macros

Delete two host macros.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usermacro.delete",
  "params": [
    "32",
    "11"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "hostmacroids": [
      "32",
      "11"
    ]
  },
  "id": 1
}
```

Source

CUserMacro::delete() in *frontends/php/api/classes/CUserMacro.php*.

usermacro.deleteglobal

Description

object usermacro.deleteglobal(array globalMacroIds)

This method allows to delete global macros.

Parameters

(string/array) IDs of the global macros to delete.

Return values

(object) Returns an object containing the IDs of the deleted global macros under the `globalmacroids` property.

Examples

Deleting multiple global macros

Delete two global macros.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usermacro.deleteglobal",
  "params": [
    "32",
    "11"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "globalmacroids": [
      "32",
      "11"
    ]
  },
  "id": 1
}
```

Source

CUserMacro::deleteGlobal() in *frontends/php/api/classes/CUserMacro.php*.

usermacro.get

Description

integer/array usermacro.get(object parameters)

The method allows to retrieve host and global macros according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
globalmacro	flag	Return global macros instead of host macros.
globalmacroids	string/array	Return only global macros with the given IDs.
groupids	string/array	Return only host macros that belong to hosts or templates from the given host groups.
hostids	string/array	Return only host macros that belong to the given hosts.
hostmacroids	string/array	Return only host macros with the given IDs.
templateids	string/array	Return only host macros that belong to the given templates.
selectGroups	query	Return host groups that the host macro belongs to in the groups property. Does not support shorten or arrays of properties. Used only when retrieving host macros.
selectHosts	query	Return hosts that the host macro belongs to in the hosts property. Does not support shorten or arrays of properties. Used only when retrieving host macros.

Parameter	Type	Description
selectTemplates	query	Return templates that the host macro belongs to in the <code>templates</code> property.
sortfield	string/array	Does not support <code>shorten</code> or arrays of properties. Used only when retrieving host macros. Sort the result by the given properties.
countOutput	flag	Possible value: <code>macro</code> . These parameters being common for all <code>get</code> methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving host macros for a host

Retrieve all host macros defined for host "10198".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "usermacro.get",
  "params": {
    "output": "extend",
    "hostids": "10198"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "hostmacroid": "9",
      "hostid": "10198",
      "macro": "{$INTERFACE}",
      "value": "eth0"
    },
    {
      "hostmacroid": "11",
      "hostid": "10198",
      "macro": "{$SNMP_COMMUNITY}",

```

```

        "value": "public"
    }
],
    "id": 1
}

```

Retrieving global macros

Retrieve all global macros.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "usermacro.get",
    "params": {
        "output": "extend",
        "globalmacro": true
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": [
        {
            "globalmacroid": "6",
            "macro": "{$SNMP_COMMUNITY}",
            "value": "public"
        }
    ],
    "id": 1
}

```

Source

CUserMacro::get() in *frontends/php/api/classes/CUserMacro.php*.

usermacro.update

Description

object usermacro.update(object/array hostMacros)

This method allows to update existing host macros.

Parameters

(object/array) **Host macro properties** to be updated.

The `hostmacroid` property must be defined for each host macro, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated host macros under the `hostmacroids` property.

Examples

Changing the value of a host macro

Change the value of a host macro to "public".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "usermacro.update",

```

```

    "params": {
        "hostmacroid": "1",
        "value": "public"
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "hostmacroids": [
            "1"
        ]
    },
    "id": 1
}

```

Source

CUserMacro::update() in *frontends/php/api/classes/CUserMacro.php*.

usermacro.updateglobal

Description

object usermacro.updateglobal(object/array globalMacros)

This method allows to update existing global macros.

Parameters

(object/array) **Global macro properties** to be updated.

The globalmacroid property must be defined for each global macro, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

(object) Returns an object containing the IDs of the updated global macros under the globalmacroids property.

Examples

Changing the value of a global macro

Change the value of a global macro to "public".

Request:

```

{
    "jsonrpc": "2.0",
    "method": "usermacro.updateglobal",
    "params": {
        "globalmacroid": "1",
        "value": "public"
    },
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": {
        "globalmacroids": [
            "1"
        ]
    }
}

```

```

    },
    "id": 1
}

```

Source

CUserMacro::updateGlobal() in *frontends/php/api/classes/CUserMacro.php*.

Web scenario

This class is designed to work with web scenarios.

Object references:

- [Web scenario](#)
- [Scenario step](#)

Available methods:

- [webcheck.create](#) - creating new web scenarios
- [webcheck.delete](#) - deleting web scenarios
- [webcheck.get](#) - retrieving web scenarios
- [webcheck.isreadable](#) - checking if web scenarios are readable
- [webcheck.iswritable](#) - checking if web scenarios are writable
- [webcheck.update](#) - updating web scenarios

> Web scenario object

The following objects are directly related to the webcheck API.

Web scenario

The web scenario object has the following properties.

Property	Type	Description
httpstid	string	<i>(readonly)</i> ID of the web scenario.
applicationid (required)	string	ID of the application that the web scenario belongs to.
name (required)	string	Name of the web scenario.
agent	string	User agent string that will be used by the web scenario.
authentication	integer	Authentication method that will be used by the web scenario. Possible values: 0 - <i>(default)</i> none; 1 - basic HTTP authentication; 2 - NTLM authentication.
delay	integer	Execution interval of the web scenario in seconds. Default: 60.
http_password	string	Password used for authentication.
http_user	string	Required for web scenarios with basic HTTP or NTLM authentication. User name used for authentication.
macros	string	Required for web scenarios with basic HTTP or NTLM authentication. Web scenario variables.
nextcheck	timestamp	<i>(readonly)</i> Time of the next web scenario execution.

Property	Type	Description
status	integer	Whether the web scenario is enabled. Possible values are: 0 - <i>(default)</i> enabled; 1 - disabled.

Scenario step

The scenario step object defines a specific web scenario check. It has the following properties.

Property	Type	Description
webstepid	string	<i>(readonly)</i> ID of the scenario step.
name (required)	string	Name of the scenario step.
no (required)	integer	Sequence number of the step in a web scenario.
url (required)	string	URL to be checked.
httptestid	string	<i>(readonly)</i> ID of the web scenario that the step belongs to.
posts	string	HTTP POST variables as a string.
required	string	Text that must be present in the response.
status_codes	string	Ranges of required HTTP status codes separated by commas.
timeout	integer	Request timeout in seconds. Default: 15.

webcheck.create

Description

object webcheck.create(object/array webScenarios)

This method allows to create new web scenarios.

Note:

Creating a web scenario will automatically create a set of **web monitoring items**.

Parameters

(object/array) Web scenarios to create.

Additionally to the **standard web scenario properties**, the method accepts the following parameters.

Parameter	Type	Description
hostid (required)	string	ID of the host that the web scenario belongs to.
steps (required)	array	Web scenario steps.

Return values

(object) Returns an object containing the IDs of the created web scenarios under the `httptestids` property. The order of the returned IDs matches the order of the passed web scenarios.

Examples

Creating a web scenario

Create a web scenario to monitor the company home page. The scenario will have two steps, to check the home page and the "About" page and make sure they return the HTTP status code 200.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "webcheck.create",
  "params": {
    "name": "Homepage check",
    "applicationid": "654",
    "hostid": "30079",
    "steps": [
      {
        "name": "Homepage",
        "url": "http://mycompany.com",
        "status_codes": 200,
        "no": 1
      },
      {
        "name": "Homepage / About",
        "url": "http://mycompany.com/about",
        "status_codes": 200,
        "no": 2
      }
    ]
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "httptestids": [
      "5"
    ]
  },
  "id": 1
}
```

See also

- [Scenario step](#)

Source

CWebCheck::create() in *frontends/php/api/classes/CWebCheck.php*.

webcheck.delete

Description

object webcheck.delete(array webScenarioIds)

This method allows to delete web scenarios.

Parameters

(array) IDs of the web scenarios to delete.

Return values

(object) Returns an object containing the IDs of the deleted web scenarios under the `httptestids` property.

Examples

Deleting multiple web scenarios

Delete two web scenarios.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "webcheck.delete",
  "params": [
    "2",
    "3"
  ],
  "auth": "3a57200802b24cda67c4e4010b50c065",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "httpstestids": [
      "2",
      "3"
    ]
  },
  "id": 1
}
```

Source

CWebCheck::delete() in *frontends/php/api/classes/CWebCheck.php*.

webcheck.get

Description

integer/array webcheck.get(object parameters)

The method allows to retrieve web scenarios according to the given parameters.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
applicationids	string/array	Return only web scenarios that belong to the given applications.
hostids	string/array	Return only web scenarios that belong to the given hosts.
httpstestids	string/array	Return only web scenarios with the given IDs.
selectHosts	query	Return the host that the web scenario belongs to as an array in the <code>hosts</code> property.
selectSteps	query	Does not support shorten or arrays of properties. Return an object containing the web scenario steps in the <code>steps</code> property.
sortfield	string/array	Does not support shorten or arrays of properties. Sort the result by the given properties.
countOutput	flag	Possible values are: <code>httpstestid</code> and <code>name</code> . These parameters being common for all get methods are described in detail on the Generic Zabbix API information page.
editable	boolean	
excludeSearch	flag	
filter	object	

Parameter	Type	Description
limit	integer	
nodeids	string/array	
output	query	
preservekeys	flag	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	flag	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving a web scenario

Retrieve all data about web scenario "5".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "webcheck.get",
  "params": {
    "output": "extend",
    "selectSteps": "extend",
    "httpstestids": "5"
  },
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "steps": [
        {
          "httpstestid": "5",
          "name": "Homepage",
          "no": "1",
          "url": "http://mycompany.com",
          "timeout": "30",
          "posts": "",
          "required": "",
          "status_codes": "200",
          "webstepid": "5"
        },
        {
          "httpstestid": "5",
          "name": "Homepage / About",
          "no": "2",
          "url": "http://mycompany.com/about",
          "timeout": "30",
          "posts": "",
          "required": "",
          "status_codes": "200",
          "webstepid": "6"
        }
      ]
    }
  ]
}
```

```

        }
    ],
    "httptestid": "5",
    "name": "Homepage check",
    "applicationid": "654",
    "nextcheck": "0",
    "delay": "60",
    "status": "0",
    "macros": "",
    "agent": "",
    "authentication": "0",
    "http_user": "",
    "http_password": ""
    }
],
"id": 1
}

```

See also

- [Host](#)
- [Scenario step](#)

Source

CWebCheck::get() in *frontends/php/api/classes/CWebCheck.php*.

webcheck.isreadable

Description

`boolean webcheck.isreadable(array webScenarioIds)`

This method checks if the given web scenarios are available for reading.

Parameters

(array) IDs of the web scenarios to check.

Return values

(boolean) Returns true if the given web scenarios are available for reading.

Examples

Check multiple web scenarios

Check if the two web scenarios are readable.

Request:

```

{
    "jsonrpc": "2.0",
    "method": "webcheck.isreadable",
    "params": [
        "3",
        "5"
    ],
    "auth": "038e1d7b1735c6a5436ee9eae095879e",
    "id": 1
}

```

Response:

```

{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}

```

See also

- [webcheck.iswritable](#)

Source

CWebCheck::isReadable() in *frontends/php/api/classes/CWebCheck.php*.

webcheck.iswritable

Description

`boolean webcheck.iswritable(array webScenarioIds)`

This method checks if the given web scenarios are available for writing.

Parameters

(array) IDs of the web scenarios to check.

Return values

(boolean) Returns true if the given web scenarios are available for writing.

Examples

Check multiple web scenarios

Check if the two web scenarios are writable.

Request:

```
{
  "jsonrpc": "2.0",
  "method": "webcheck.iswritable",
  "params": [
    "3",
    "5"
  ],
  "auth": "038e1d7b1735c6a5436ee9eae095879e",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": true,
  "id": 1
}
```

See also

- [webcheck.isreadable](#)

Source

CWebCheck::isWritable() in *frontends/php/api/classes/CWebCheck.php*.

webcheck.update

Description

`object webcheck.update(object/array webScenarios)`

This method allows to update existing web scenarios.

Parameters

(object/array) Web scenario properties to be updated.

The `httptestid` property must be defined for each web scenario, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the **standard web scenario properties**, the method accepts the following parameters.

Parameter	Type	Description
steps	array	Scenario steps to replace existing steps.

Return values

(object) Returns an object containing the IDs of the updated web scenarios under the `httptestid` property.

Examples

Enabling a web scenario

Enable a web scenario, that is, set its status to "0".

Request:

```
{
  "jsonrpc": "2.0",
  "method": "webcheck.update",
  "params": {
    "httptestid": "5",
    "status": 0
  },
  "auth": "700ca65537074ec963db7efabda78259",
  "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "httptestids": [
      "5"
    ]
  },
  "id": 1
}
```

See also

- **Scenario step**

Source

CWebCheck::update() in *frontends/php/api/classes/CWebCheck.php*.

Generic Zabbix API information

Note:
Setting *Disable GUI access* for a **user group** will also disable API access.

Attention:
When performing an API request, a `Content-Type: application/json` header must be included. Otherwise, the request will fail with a 412 Precondition Failed status code.

Supported data types

The Zabbix API supports the following data types:

Type	Description
bool	A boolean value, accepts either <code>true</code> or <code>false</code> .
flag	The value is considered to be <code>true</code> if it is passed and not equal to <code>null</code> and <code>false</code> otherwise.

Type	Description
integer	A whole number.
float	A floating point number.
string	A text string.
timestamp	A Unix timestamp
array	An ordered sequence of values, that is, a plain array.
object	An associative array.
query	A value which defines, what data should be returned.
	Can be defined as an array of property names to return only specific properties, or as one of the predefined values: shorten - returns just the primary key; refer - returns the primary key and the foreign keys used to link related objects; extend - returns all object properties; count - returns the number of retrieved records; supported only by certain subselects.

Common "get" method parameters

The following parameters are supported by all get methods:

Parameter	Type	Description
countOutput	flag	Return the number of records in the result instead of the actual data.
editable	boolean	If set to true return only objects that the user has write permissions to.
excludeSearch	flag	Default: false. Return results that do not match the criteria given in the search parameter.
filter	object	Return only those results that exactly match the given filter.
limit	integer	Accepts an array, where the keys are property names, and the values are either a single value or an array of values to match against. Limit the number of records returned.
nodeids	string/array	Returns objects that belong to the given nodes.
output	query	Object properties to be returned.
preservekeys	flag	Default: refer. Use IDs as keys in the resulting array.
search	object	Return results that match the given wildcard search. Accepts an array, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a LIKE "%...%" search.
searchByAny	boolean	Works only with text fields. If set to true return results that match any of the criteria given in the filter or search parameter instead of all of them.
searchWildcardsEnabled	boolean	Default: false. If set to true enables the use of "*" as a wildcard character in the search parameter. Default: false.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties. Refer to a specific API get method description for a list of properties that can be used for sorting. Macros are not expanded before sorting.
sortorder	string/array	Order of sorting. If an array is passed, each value will be matched to the corresponding property given in the sortfield parameter. Possible values are: ASC - ascending; DESC - descending.
startSearch	flag	The search parameter will compare the beginning of fields, that is, perform a LIKE "...%" search instead.

Zabbix API changes from 1.8 to 2.0

Warning:

List of changes for all ***.get** methods is complete.
List of changes for other methods is still not complete.

Generic changes

- API access property has been removed. No reasons given.
- For all ***.get** methods a new *searchByAny* parameter has been added.

Action

Actions API changed significantly:

- Fields were removed from "operations" array: object, objectid, shortdata, longdata, default_msg.
- shortdata, longdata, default_msg were moved to array opmessage. shortdata and longdata were renamed to subject and message.
- For different operation types added different arrays of parameters:
 - * for "Send message" operation added arrays opmessage, opmessage_usr, opmessage_grp
 - * for "Execute remote command" operation added arrays opcommand_hst, opcommand_grp
 - * for "Add host group" and "Remove host group" added array opgroup
 - * for "Link to template" and "Unlink from template" added array optemplate
- * Added new parameters for //action.get// method: usrgrpids, scriptids
- * Parameter names for method //action.get// have been changed from → to: select_conditions → selectConditions

Alert

Parameter names for method *alert.get* have been changed from → to: select_mediatypes → selectMediatypes, select_users → selectUsers, select_hosts → selectHosts.

Apiinfo

API version has been increased. Method *apiinfo.version* now returns 1.4 as a result.

Application

Parameter names for method *application.get* have been changed from → to: select_hosts → selectHosts, select_items → selectItems.

Dcheck

Added new parameters for *dcheck.get* method: dserviceids, selectHosts.

Dhost

No specific changes for *dhost.get* method.

Discoveryrule

New API class for low level discovery rules has been added.

Drule

No specific changes for *drule.get* method.

Added create/update methods.

Dservice

No specific changes for *dservice.get* method.

Event

Methods *event.create* and *event.delete* not supported anymore and have been removed.

Parameter names for method *event.get* have been changed from → to: *select_hosts* → *selectHosts*, *select_items* → *selectItems*, *select_triggers* → *selectTriggers*

Parameter *hide_unknown* replaced by *showUnknown*

Added new parameters for extending output: *select_alerts*, *select_acknowledges*, *groupCount*

Added new parameters for optional filtering: *filter*, *search*, *searchByAny*, *startSearch*, *excludeSearch*, *searchWildcardsEnabled*.

Graph

Added new parameters for *graph.get* method: *discoveryids*, *selectDiscoveryRule*.

Parameter names for method *graph.get* have been changed from → to: *select_groups* → *selectGroups*, *select_templates* → *selectTemplates*, *select_hosts* → *selectHosts*, *select_items* → *selectItems*, *select_graph_items* → *selectGraphItems*.

Graphitem

Parameter name for method *graphitem.get* has been changed from → to: *select_graphs* → *selectGraphs*.

Graphprototype

New API class for graph prototypes has been added.

History

No specific changes for *history.get* method.

Host

Each host can now have multiple host interfaces instead of just one.

Added new parameters for *host.get* method: *interfaceids*, *applicationids*, *httptestids*, *with_simple_graph_items*, *with_applications*, *withInventory*, *selectDiscoveries*, *selectScreens*, *selectInterfaces*

Parameter names for method *host.get* have been changed from → to: *select_groups* → *selectGroups*, *select_items* → *selectItems*, *select_triggers* → *selectTriggers*, *select_graphs* → *selectGraphs*, *select_dhosts* → *selectDHosts*, *select_dservices* → *selectDServices*, *select_applications* → *selectApplications*, *select_macros* → *selectMacros*, *select_profile* → *selectInventory*.

host.get parameter 'selectInventory' now supports not only boolean values, but also an array of fields that needs to be fetched (similar to 'output' option).

host.create, *host.update* and *host.massUpdate* parameter 'profile' and 'extendedProfile' are replaced by 'inventory'.

host.create has new parameter 'inventory_mode' which specifies if host inventory is: -1 - disabled, 0 - manual or 1 - automatic, for the host being created. *host.massupdate* with macros now removes all existing macros from hosts and create new ones instead of trying to update macros based on hostmacroid. To update host macros *host.update* can be used.

Hostinterface

New API class for host interfaces has been added.

Hostgroup

Added new parameters for *hostgroup.get* method: *with_hosts_and_templates*, *with_simple_graph_items*, *with_applications*

Parameter names for method *hostgroup.get* have been changed from → to: *select_hosts* → *selectHosts*, *select_templates* → *selectTemplates*

Item

Added new parameters for *item.get* method: *interfaceids*, *discoveryids*, *selectInterfaces*, *selectDiscoveryRule*, *selectItemDiscovery*.

Parameter names for method *item.get* have been changed from → to: *select_hosts* → *selectHosts*, *select_triggers* → *selectTriggers*, *select_graphs* → *selectGraphs*, *select_applications* → *selectApplications*.

An item field "description" has been renamed to "name".

New item field "description" added. It can be used for multi-line item description.

New item field "inventory_link" added.

Itemprototype

New API class for item prototypes has been added.

Iconmap

New API class for icon maps has been added.

Image

No changes for *image.get* method.

Maintenance

Parameter names for method *maintenance.get* have been changed from → to: *select_groups* → *selectGroups*, *select_hosts* → *selectHosts*

Map

Map object gets 16 additional fields

- 3 to control map display during configuration:
 - *grid_size* - integer, optional, default 50
 - *grid_show* - integer, optional, default 1
 - *grid_align* - integer, optional, default 1
- 11 to control element labels
 - *label_format* - integer, optional, default 0
 - *label_type_hostgroup* - integer, optional, default 2
 - *label_type_host* - integer, optional, default 2
 - *label_type_trigger* - integer, optional, default 2
 - *label_type_map* - integer, optional, default 2
 - *label_type_image* - integer, optional, default 2
 - *label_string_hostgroup* - string, optional, default ""
 - *label_string_host* - string, optional, default ""
 - *label_string_trigger* - string, optional, default ""
 - *label_string_map* - string, optional, default ""
 - *label_string_image* - string, optional, default ""
- 1 to represent icon map linkage
 - *iconmapid* - id, optional, default NULL
- 1 to control if macros are expanded on map configuration screen
 - *expand_macros* - int, optional, default 0

Map element gets 6 additional fields

- 5 to control areas
 - *elementsubtype* integer, optional, default 0
 - *areatype* integer, optional, default 0
 - *width* integer, optional, default 200
 - *height* integer optional, default 200
 - *viewtype* integer, optional, default 0
- 1 to enable icon mapping for element
 - *use_iconmap* - integer, optional, default 1

Added new parameters for *map.get* method: *selectIconMap*, *expandUrls*.

Parameter names for method *map.get* have been changed from → to: *select_links* → *selectLinks*, *select_selements* → *selectSelements*.

map.addElements has been renamed to *map.createSelements*, *map.updateElements* to *map.updateSelements* and *map.deleteElements* to *map.deleteSelements*.

Mapelement

New API class for map elements has been added.

Mediatype

Added new mediatype field "status".

Parameter names for method *mediatype.get* have been changed from → to: *select_users* → *selectUsers*, *select_medias* → *selectMedias*.

Proxy

Added new parameters for *proxy.get* method: selectInterfaces.

Parameter names for method *proxy.get* have been changed from → to: select_hosts → selectHosts.

Screen

Added new parameters for *screen.get* method: groupCount.

Parameter names for method *screen.get* have been changed from → to: select_screenitems → selectScreenItems.

Screenitem

New API class for screen items has been added.

Script

Added new script fields: description, question, type, execute_on.

Added new parameters for *script.get* method: usrgroups.

Parameter names for method *script.get* have been changed from → to: select_groups → selectGroups, select_hosts → selectHosts.

Template

Added new parameters for *template.get* method: selectDiscoveries, selectScreens.

Parameter names for method *template.get* have been changed from → to: select_groups → selectGroups, select_hosts → selectHosts, select_templates → selectTemplates, selectParentTemplates → selectParentTemplates, select_items → selectItems, select_triggers → selectTriggers, select_graphs → selectGraphs, select_applications → selectApplications, select_macros → selectMacros.

Trigger

Added new value "hostname" as sortfield.

Added new parameters for *trigger.get* method: discoveryids, selectDiscoveryRule.

Parameter names for method *trigger.get* have been changed from → to: select_groups → selectGroups, select_hosts → selectHosts, select_items → selectItems, select_functions → selectFunctions, select_dependencies → selectDependencies.

Triggerprototype

New API class for trigger prototypes has been added.

User

Method *user.addMedia* if it's successful now returns array of created user media IDs. Previously it returned array of updated User IDs.

Method *user.deleteMedia* now should receive an array of user media IDs which should be deleted. Previously it should receive media objects which should be deleted.

Added new parameters for *user.get* method: selectMedias.

Parameter names for method *user.get* have been changed from → to: select_usrgroups → selectUsrgroups, select_mediatypes → selectMediatypes, get_access → getAccess.

Method *user.get* with the option "output": "extend" now will not show user password hashes, i.e. they are removed from output.

Usergroup

From UserGroup object removed field api_access.

Removed "with_api_access" parameter for *usergroup.get* method.

Parameter names for method *usergroup.get* have been changed from → to: select_users → selectUsers.

Usermacro

Method *usermacro.updateglobal* now uses macro IDs instead of macro names.

Method *usermacro.massadd* renamed to *usermacro.create* and will now accept an array of macros.

Method *usermacro.massupdate* renamed to *usermacro.update*. It will now accept an array of macros, and update the data via IDs instead of macro names.

Method *usermacro.deletehostmacro* renamed to *usermacro.delete*.

Removed the *usermacro.massRemove* method.

Parameter names for method *usermacro.get* have been changed from → to: select_groups → selectGroups, select_hosts →

selectHosts, select_templates → selectTemplates.

Usermedia

Added new class *Usermedia* with a method *get*.

Webcheck

New API class for web checks has been added.

Zabbix API changes in 2.0

2.0.17

service

Bug fixes:

[ZBX-10232](#) service.getSla: fixed SQL errors with invalid "year" parameter in IT services report

webcheck

Bug fixes:

[ZBX-10316](#) removed faulty web scenario step name validation

2.0.16

hostgroup

Bug fixes:

[ZBX-9738](#) hostgroup.delete: fixed deletion of related action operations when deleting a host group

item

Bug fixes:

[ZBX-8235](#) item.update: fixed losing initial values when updating templated items

2.0.15

hosts

Bug fixes:

[ZBX-9093](#) host.create: fixed "inventory" property causing SQL errors when using MySQL strict mode

maintenance

Bug fixes:

[ZBX-5656](#) fixed "timeperiods" validation when passing a single timeperiod object

2.0.14

triggers

Bug fixes:

[ZBX-6174](#) trigger.get: fixed "skipDependent" option not handling cases when triggers upon which other triggers depend are disabled (or have disabled items or disabled item hosts).

2.0.13

configuration

Bug fixes:

[ZBX-8151](#) configuration.import: fixed XXE vulnerability while importing XML with external entities

2.0.12

graphprototype

Bug fixes:

[ZBX-6151](#) fixed graph prototype validation so item prototypes are no longer allowed from multiple discovery rules

trigger

Bug fixes:

[ZBX-7674](#) trigger.delete: fixed trigger unlink from IT Services

triggerprototype

Bug fixes:

[ZBX-6151](#) fixed trigger prototype validation so item prototypes are no longer allowed from multiple discovery rules

2.0.11

user

Bug fixes:

[ZBX-7693](#) fixed admin users being able to edit media for other users

[ZBX-7703](#) user.login: fixed being able to switch users without proper credentials when using HTTP authentication

2.0.10

action

Bug fixes:

[ZBX-7101](#) action.get: fixed generated SQL queries when simultaneously selecting more than one condition and operation

2.0.9

General

Bug fixes:

[ZBX-7091](#) fixed multiple SQL injection vulnerabilities

service

Bug fixes:

[ZBX-6992](#) service.getsla: fixed calculation of SLA for PostgreSQL

2.0.8

No changes.

2.0.7

itemprototype

Bug fixes:

[ZBX-6644](#) itemprototype.get: fixed sorting by history and trends fields

map

Bug fixes:

[ZBX-6561](#) map.get: fixed incorrect API response (only for zabbix administrators)

2.0.6

services

Bug fixes:

[ZBX-6261](#) service.get: fixed method not working with Oracle

templates

Bug fixes:

[ZBX-6408](#) template.delete: fixed action conditions not being deleted and actions not being disabled when deleting a template

triggers

Changes:

[ZBX-6389](#) trigger.get: added option "selectLastEvent" to select latest trigger event

webcheck

Bug fixes:

[ZBX-6413](#) fixed wrongly implemented web scenario name validation, which raised exception even if name was not set

2.0.5

actions

Bug fixes:

[ZBX-6142](#) action.update: fixed not being able to remove the "current host" target from remote command operations

event

Bug fixes:

[ZBX-6005](#) event.get: fixed not returning network discovery or auto-registration events for admin users when passing only the "source" parameter

host

Bug fixes:

[ZBX-6214](#) hostgroup.massupdate: fixed being able to unlink all groups from a host

maintenance

Bug fixes:

[ZBX-5900](#) fixed timeperiod validation in maintenance periods

map

Bug fixes:

[ZBX-5945](#) map.get: fixed selectSelements used together with expandUrls not returning map URIs when their IDs match map element URL IDs; selements urls property will now be returned as an array instead of a hash

services

Bug fixes:

[ZBX-5532](#) service.get: added support of 1.8 service time reading

Changes:

[ZBX-6002](#) service.getsla: improved SLA calculation performance

template

Bug fixes:

[ZBX-3534](#) fixed not being able to link templates with read-only permissions

[ZBX-5806](#) template.massadd: fixed returning an empty hostids array

[ZBX-5806](#) template.massremove: fixed returning an empty hostids array

=== user ===

Bug fixes:

[ZBX-6097](#) user.login: fixed the ability to override LDAP configuration

2.0.4

General

Bug fixes:

[ZBX-5648](#) disallowed explicitly setting the "templateid" property for objects

application

Bug fixes:

[ZBX-5977](#) application.massadd: disallowed possibility to add discovery rule in applications

discoveryrule

Bug fixes:

[ZBX-5977](#) removed possibility to add an LLD rule to an application

[ZBX-5684](#) discoveryrule.get: added support for 'interfaceids' parameter

[ZBX-5751](#) discoveryrule.get: fixed parameter "selectGraph" not working

graph

Bug fixes:

[ZBX-5693](#) fixed possibility to create graph with the same name as existing graph prototype within template; improved error message

[ZBX-5740](#) graph.get: removed unsupported "discoveryids" parameter

graphprototype

Bug fixes:

[ZBX-5693](#) fixed possibility to create graph prototype with the same name as existing graph within template; improved error message

hostgroup

Bug fixes:

[ZBX-5860](#) hostgroup.massupdate: fixed multiple issues

[ZBX-5837](#) hostgroup.get: fixed filtering by proxy status in HostGroup API get method

image

Bug fixes:

[ZBX-5737](#) image.create: fixed image saving using Oracle

[ZBX-5704](#) image.create: fixed validation on file uploading

item

Bug fixes:

[ZBX-5684](#) item.get: added possibility to unset 'flags' parameter in CItem get()

[ZBX-5684](#) item.get: fixed inconsistency that CItem get() returned discovery rules and item prototypes by default

[ZBX-5740](#) item.get: removed unsupported "discoveryids" parameter

[ZBX-5742](#) item.get: fixed sorting of triggers returned by the selectTriggers parameter

[ZBX-4097](#) item.get: fixed the group, host and application parameters being case insensitive

itemprototype

Bug fixes:

[ZBX-5684](#) itemprototype.get: added support for 'graphids' parameter

[ZBX-5751](#) itemprototype.get: fixed parameter "selectGraph" not working. Fixed "limitSelects" for "selectTriggers" not working

[ZBX-5740](#) itemprototype.get: fixed parameter "selectTriggers" not working

[ZBX-5742](#) itemprototype.get: fixed sorting of triggers returned by the selectTriggers parameter

=== maintenance ===

Bug fixes:

[ZBX-5656](#) maintenance.create: fixed "timeperiods" validation when passing a single timeperiod object

screenitem

Bug fixes:

[ZBX-1946](#) fixed the colspan and rowspan properties not being validated

scripts

Bug fixes:

[ZBX-5744](#) script.getscriptsbyhost: fixed returning an unnecessary array of hosts for each script

=== templatescreen ===

Changes:

[ZBX-4833](#) templatescreen.copy: implemented

[ZBX-4833](#) templatescreen.isreadable: implemented

[ZBX-4833](#) templatescreen.iswritable: implemented

Bug fixes:

[ZBX-5648](#) templatescreen.update: disallowed updating the "templateid" property

trigger

Bug fixes:

[ZBX-5643](#) fixed sql error by adding permission check for triggers when adding nonexistent dependencies

[ZBX-5643](#) fixed being able to create circular dependencies

[ZBX-3920](#) added check for dependency duplicates

[ZBX-5740](#) trigger.get: removed unsupported "discoveryids" parameter

[ZBX-4097](#) trigger.get: fixed the group and host parameters being case insensitive

[ZBX-5111](#) trigger.update: fixed child trigger data reset after changes in parent trigger

=== triggerprototype ===

Bug fixes:

[ZBX-5698](#) triggerprototype.create: fixed bug when response doesn't return "triggerids"

[ZBX-5665](#) triggerprototype.delete: fixed bug when deleting a trigger prototype and not deleting created triggers

[ZBX-4097](#) triggerprototype.get: fixed the group and host parameters being case insensitive

=== user ===

Bug fixes:

[ZBX-3907](#) user.logout: fixed user.logout method

usermacro

Bug fixes:

[ZBX-4983](#) usermacro.replacemacros: fixed an error when attempting to replace existing macro with identical macro

2.0.3

alert

Changes:

[ZBX-5489](#) alert.create: removed

[ZBX-5489](#) alert.delete: removed

configuration

Bug fixes:

[ZBX-5307](#) fixed undefined index error when importing a trigger that depends on an unexisting trigger

[ZBX-5307](#) fixed errors when importing item prototypes without applications

event

Changes:

[ZBX-5489](#) event.create: removed

[ZBX-5489](#) event.delete: removed

[ZBX-5489](#) event.deleteByTriggerIds: removed

graph

Bug fixes:

[ZBX-5363](#) fixed checking that the Y axis MAX value is greater than Y axis MIN value when the MAX value is not fixed

graphprototype

Bug fixes:

[ZBX-5363](#) fixed check that the Y axis MAX value is greater than Y axis MIN value when the MAX value is not fixed

host

Bug fixes:

[ZBX-5250](#) host.massRemove: fixed macros parameter not working

item

Bug fixes:

[ZBX-4852](#) fixed delay_flex attribute validation

[ZBX-5510](#) item.get: fixed Oracle DB error when trying to retrieve the name of the item and using the group parameter

itemprototype

Bug fixes:

[ZBX-5434](#) fixed error when inheriting an item prototype with the same key as an item on the target host or template

maintenance

Changes:

[ZBX-5351](#) maintenance.get: implemented the selectTimeperiods parameter

map

Bug fixes:

[ZBX-5317](#) map.update: fixed duplicate entry error when updating map link triggers

template

Bug fixes:

[ZBX-5250](#) template.massRemove: fixed macros parameter not working

templatescreenitem

Changes:

[ZBXNEXT-1253](#) templatescreenitem.get: implemented

trigger

Changes:

[ZBX-2060](#) implemented the expandExpression parameter

Bug fixes:

[ZBX-5307](#) fixed being able to create dependencies from a templated trigger to an inherited trigger

[ZBX-5307](#) fixed inherited trigger dependencies not being validated on hosts

[ZBX-5308](#) fixed templated triggers not overriding dependencies on inherited triggers

triggerprototype

Changes:

[ZBX-2060](#) triggerprototype.get: implemented the expandExpression parameter

usergroup

Bug fixes:

[ZBX-5440](#) fixed not being able to remove the last user or the last permission from a user group

[ZBX-5440](#) usergroup.massUpdate: fixed checking that the usrgrpids parameter is passed

usergroup

Changes:

[ZBX-5292](#) the value attribute can now be set to an empty string

webcheck

Bug fixes:

[ZBX-4316](#) fixed checking that the status_codes attribute of a step object is valid

2.0.2

application

Bug fixes:

[ZBX-5252](#) fixed multiple level templated application inheritance

graphprototype

Bug fixes:

[ZBX-5349](#) graphprototype.create: fixed graph prototypes being created as plain graphs

host

Changes:

[ZBX-5208](#) host.update: the interfaces parameter will now also accept a single interface

Bug fixes:

[ZBX-5079](#) host.get: fixed selectTriggers and selectGraphs returning objects that the user has no permissions to

[ZBX-5060](#) host.update: fixed inventory_mode attribute not being updated without updating inventory

[ZBX-5208](#) host.update: fixed interface updating not working

[ZBX-5060](#) host.massUpdate: fixed inventory_mode attribute not being updated without updating inventory

hostgroup

Changes:

[ZBX-5267](#) hostgroup.delete: fixed method not accepting a single ID

[ZBX-4007](#) hostgroup.get: fixed the graphid attribute being returned for hostgroup objects

hostinterface

Bug fixes:

[ZBX-5208](#) hostinterface.massRemove: fixed checking that the default interface is not deleted

item

Bug fixes:

[ZBX-2055](#) fixed flexible interval validation

proxy

Bug fixes:

[ZBX-5311](#) proxy.delete: fixed checking that a proxy is not used for monitoring or network discovery

screenitem

Bug fixes:

[ZBX-4007](#) fixed using resources from child nodes

service

Bug fixes:

[ZBX-5065](#) fixed checking that services do not create circular dependencies

[ZBX-5065](#) service.delete: fixed not being able to delete services with soft linked dependencies

template

Bug fixes:

[ZBX-5079](#) template.get: fixed selectTriggers and selectGraphs returning objects that the user has no permissions to

triggerprototype

Bug fixes:

[ZBX-4963](#) fixed inherited trigger prototypes lacking certain data

usergroup

Bug fixes:

[ZBX-5213](#) usergroup.massUpdate: fixed update failing when updating only permissions or users

2.0.1

graph

Bug fixes:

[ZBX-5053](#) fixed checking that the Y axis MAX value is greater than Y axis MIN value

graphprototype

Bug fixes:

[ZBX-5053](#) fixed checking that the Y axis MAX value is greater than Y axis MIN value

[ZBX-4861](#) fixed graph item color not being validated

item

Changes:

[ZBX-5150](#) value_type attribute can accept value 3, numeric, for aggregate items

itemprototype

Changes:

[ZBX-5150](#) value_type attribute can accept value 3, numeric, for aggregate items

user

Changes:

[ZBX-5032](#) user.get: implemented the selectMediatypes parameter

usermedia

Changes:

[ZBX-5032](#) usermedia.create: removed

[ZBX-5032](#) usermedia.delete: removed

[ZBX-5032](#) usermedia.isReadable: removed

[ZBX-5032](#) usermedia.isWritable: removed

[ZBX-5032](#) usermedia.update: removed

1 Frequently asked questions / Troubleshooting

Frequently asked questions or FAQ.

1. Q: Can I flush/clear the queue (as depicted in *Administration* → *Queue*)?
A: No.
2. Q: How do I migrate from one database to another?
A: Dump data only (for MySQL, use flag -t or --no-create-info), create the new database using schema files from Zabbix and import the data.
3. Q: I would like to replace all spaces with underscores in my item keys because they worked in older versions but space is not a valid symbol for an item key in 1.8 (or any other reason to mass-modify item keys). How should I do it and what should I beware of?
A: You may use a database query to replace all occurrences of spaces in item keys with underscores:
update items set key_=replace(key_,' ','_');
Triggers will be able to use these items without any additional modifications, but you might have to change any item references in these locations:
 - * Notifications (actions)
 - * Map element and link labels
 - * Calculated item formulas
4. Q: My graphs have dots instead of lines or empty areas. Why so?
A: Data is missing. This can happen for a variety of reasons - performance problems on Zabbix database, Zabbix server, network, monitored devices...
5. Q: Zabbix daemons fail to start up with a message *Listener failed with error: socket() for [[:]:10050] failed with error 22: Invalid argument*.
A: This error arises at attempt to run Zabbix agent compiled on version 2.6.27 or above on a platform with a kernel 2.6.26 and lower. Note that static linking will not help in this case because it is the socket() system call that does not support SOCK_CLOEXEC flag on earlier kernels. [ZBX-3395](#)
6. Q: I try to set up a flexible user parameter (one that accepts parameters) with a command that uses a positional parameter like \$1, but it doesn't work (uses item parameter instead). How to solve this?
A: Use a double dollar sign like **\$\$1**

7. Q: All dropdowns have a scrollbar and look ugly in Opera 11. Why so?
A: It's a known bug in Opera 11.00 and 11.01; see [Zabbix issue tracker](#) for more information.
8. Q: What is the structure of IDs in the database for distributed monitoring?
A: **NNNSSDDDDDDDDDDDD**, where NNN - nodeid (to which node the ID belongs to), SSS - source nodeid (in which node was the ID created), DDDDDDDDDDD - the ID itself.
For instance, ID created on source node 5 for node 14 might look like 1400500000012345. Note that in this example the length of ID is less than 17 digits, because the target node number has less than three digits.
9. Q: How can I change graph background colour in a custom theme?
A: See graph_theme table in the database and [theming guide](#).

Installation troubleshooting

See the [installation-specific troubleshooting section](#).

See also

* [Troubleshooting page on zabbix.org](#)

2 Installation

1 Database creation scripts

Overview

A Zabbix database must be created during the installation of Zabbix server or proxy.

This section provides scripts for creating a Zabbix database. A separate schema script is provided for each supported database.

Note:

schema.sql, images.sql and data.sql files are located in the *database* subdirectory of Zabbix sources. If Zabbix was installed from distribution packages, refer to the distribution documentation.

Attention:

For a Zabbix proxy database, **only** schema.sql should be imported (no images.sql nor data.sql)

Attention:

The order of database file import has been changed - **images** must be imported before data.sql now.

Scripts

MySQL

We assume that a *username* user with *password* password exists and has permissions to create database objects.

```
shell> mysql -u<username> -p<password>
mysql> create database zabbix character set utf8 collate utf8_bin;
mysql> quit;
shell> mysql -u<username> -p<password> zabbix < database/mysql/schema.sql
# stop here if you are creating database for Zabbix proxy
shell> mysql -u<username> -p<password> zabbix < database/mysql/images.sql
shell> mysql -u<username> -p<password> zabbix < database/mysql/data.sql
```

PostgreSQL

We assume that a *username* user exists and has permissions to create database objects.

```
shell> psql -U <username>
psql> create database zabbix;
psql> \q
shell> cd database/postgresql
shell> psql -U <username> zabbix < schema.sql
# stop here if you are creating database for Zabbix proxy
shell> psql -U <username> zabbix < images.sql
shell> psql -U <username> zabbix < data.sql
```

Oracle

We assume that a *zabbix* database user with *password* password exists and has permissions to create database objects in ORCL service located on the *host* Oracle database server with a *user* shell user having write access to /tmp directory.

If you are creating a database for Zabbix server you need to have images somewhere on the Oracle host, for example in the /tmp/zabbix_images folder. Copy all images from misc/images/png_modern to /tmp/zabbix_images directory on the Oracle host:

```
shell> cd /path/to/zabbix-sources
shell> scp -r misc/images/png_modern user@host:/tmp/zabbix_images
```

Edit the database/oracle/images.sql file and set images_dir variable to the /tmp/zabbix_images path:

```
CREATE OR REPLACE DIRECTORY image_dir AS '/tmp/zabbix_images'
```

Now prepare the database:

```
shell> sqlplus zabbix/password@host/ORCL
sqlplus> @database/oracle/schema.sql
# stop here if you are creating database for Zabbix proxy
sqlplus> @database/oracle/images.sql
sqlplus> @database/oracle/data.sql
```

Note:

Zabbix requires a UTF-8 database character set. If database is not UTF-8, it can be converted by running: ALTER DATABASE NATIONAL CHARACTER SET UTF8;

IBM DB2

```
shell> db2 "create database zabbix using codeset utf-8 territory us pagesize 32768"
shell> cd database/ibm_db2
shell> db2batch -d zabbix -f schema.sql
# stop here if you are creating database for Zabbix proxy
shell> db2batch -d zabbix -f images.sql
shell> db2batch -d zabbix -f data.sql
```

Zabbix frontend uses OFFSET and LIMIT clauses in SQL queries. For this to work, IBM DB2 server must have DB2_COMPATIBILITY_VECTOR variable be set to 3. Run the following command before starting the database server:

```
shell> db2set DB2_COMPATIBILITY_VECTOR=3
```

SQLite

```
shell> cd database/sqlite3
shell> sqlite3 /var/lib/sqlite/zabbix.db < schema.sql
# stop here if you are creating database for Zabbix proxy
shell> sqlite3 /var/lib/sqlite/zabbix.db < images.sql
shell> sqlite3 /var/lib/sqlite/zabbix.db < data.sql
```

Note:

If using SQLite with Zabbix proxy, database will be automatically created if it does not exist.

Return to the [installation section](#).

2 Zabbix agent on Microsoft Windows

Configuring agent

Zabbix agent runs as a Windows service.

You can run a single instance of Zabbix agent or multiple instances of the agent on a Microsoft Windows host. A single instance can use the default configuration file C:\zabbix_agentd.conf or a configuration file specified in the command line. In case of multiple instances each agent instance must have its own configuration file (one of the instances can use the default configuration file).

An example configuration file is available in Zabbix source archive as conf/zabbix_agentd.win.conf.

See the [configuration file](#) options for details on configuring Zabbix Windows agent.

Hostname parameter

To perform **active checks** on a host Zabbix agent needs to have the hostname defined. Moreover, the hostname value set on the agent side should exactly match the "**Host name**" configured for the host in the frontend.

The hostname value on the agent side can be defined by either the **Hostname** or **Hostnameltem** parameter in the agent **configuration file** - or the default values are used if any of these parameters are not specified.

The default value for **Hostnameltem** parameter is the value returned by the "system.hostname" agent key and for Windows platform it returns the NetBIOS host name.

The default value for **Hostname** is the value returned by the **Hostnameltem** parameter. So, in effect, if both these parameters are unspecified the actual hostname will be the host NetBIOS name; Zabbix agent will use NetBIOS host name to retrieve the list of active checks from Zabbix server and send results to it.

Attention:

The **system.hostname** key always returns the NetBIOS host name which is limited to 15 symbols and in UPPERCASE only - regardless of the length and lowercase/uppercase characters in the real host name.

Starting from Zabbix agent 1.8.6 version for Windows the "system.hostname" key supports an optional parameter - *type* of the name. The default value of this parameter is "netbios" (for backward compatibility) and the other possible value is "host".

Attention:

The **system.hostname[host]** key always returns the full, real (case sensitive) Windows host name.

So, to simplify the configuration of zabbix_agentd.conf file and make it unified, two different approaches could be used.

1. leave **Hostname** or **Hostnameltem** parameters undefined and Zabbix agent will use NetBIOS host name as the hostname;
2. leave **Hostname** parameter undefined and define **Hostnameltem** like this:

Hostnameltem=system.hostname[host]

and Zabbix agent will use the full, real (case sensitive) Windows host name as the hostname.

Host name is also used as part of Windows service name which is used for installing, starting, stopping and uninstalling the Windows service. For example, if Zabbix agent configuration file specifies **Hostname=Windows_db_server**, then the agent will be installed as a Windows service "Zabbix Agent [Windows_db_server]". Therefore, to have a different Windows service name for each Zabbix agent instance, each instance must use a different host name.

Installing agent as Windows service

To install a single instance of Zabbix agent with the default configuration file c:\zabbix_agentd.conf:

```
zabbix_agentd.exe --install
```

Attention:

On a 64-bit system, a 64-bit Zabbix agent version is required for all checks related to running 64-bit processes to work correctly.

If you wish to use a configuration file other than c:\zabbix_agentd.conf, you should use the following command for service installation:

```
zabbix_agentd.exe --config <your_configuration_file> --install
```

A full path to the configuration file should be specified.

Multiple instances of Zabbix agent can be installed as services like this:

```
zabbix_agentd.exe --config <configuration_file_for_instance_1> --install --multiple-agents
zabbix_agentd.exe --config <configuration_file_for_instance_2> --install --multiple-agents
...
zabbix_agentd.exe --config <configuration_file_for_instance_N> --install --multiple-agents
```

The installed service should now be visible in Control Panel.

Starting agent

To start the agent service, you can use Control Panel or do it from command line.

To start a single instance of Zabbix agent with the default configuration file:

```
zabbix_agentd.exe --start
```

To start a single instance of Zabbix agent with another configuration file:

```
zabbix_agentd.exe --config <your_configuration_file> --start
```

To start one of multiple instances of Zabbix agent:

```
zabbix_agentd.exe --config <configuration_file_for_this_instance> --start --multiple-agents
```

Stopping agent

To stop the agent service, you can use Control Panel or do it from command line.

To stop a single instance of Zabbix agent started with the default configuration file:

```
zabbix_agentd.exe --stop
```

To stop a single instance of Zabbix agent started with another configuration file:

```
zabbix_agentd.exe --config <your_configuration_file> --stop
```

To stop one of multiple instances of Zabbix agent:

```
zabbix_agentd.exe --config <configuration_file_for_this_instance> --stop --multiple-agents
```

Uninstalling agent Windows service

To uninstall a single instance of Zabbix agent using the default configuration file:

```
zabbix_agentd.exe --uninstall
```

To uninstall a single instance of Zabbix agent using a non-default configuration file:

```
zabbix_agentd.exe --config <your_configuration_file> --uninstall
```

To uninstall multiple instances of Zabbix agent from Windows services:

```
zabbix_agentd.exe --config <configuration_file_for_instance_1> --uninstall --multiple-agents
zabbix_agentd.exe --config <configuration_file_for_instance_2> --uninstall --multiple-agents
...
zabbix_agentd.exe --config <configuration_file_for_instance_N> --uninstall --multiple-agents
```

3 Troubleshooting installation issues

This page covers potential issues that could arise during installation of Zabbix

Access control with Apache

Zabbix frontend includes *.htaccess* files that limit access to directories *api*, *conf* and *include*. Since Apache 2.4 [introduced a new access control syntax](#), running versions of Zabbix 2.0.10 and older might cause the following error to appear in the Apache error logs:

Invalid command 'Order', perhaps misspelled or defined by a module not included in the server configuration

Using the outdated syntax in the Apache configuration files might prevent Apache from starting at all. To solve this problem either upgrade to Zabbix 2.0.11 or enable the Apache *mod_access_compat* module.

See [Apache documentation](#) for details.

3 Daemon configuration

1 Zabbix server

The parameters supported in a Zabbix server configuration file:

Parameter	Mandatory	Range	Default	Description
AlertScriptsPath	no		/usr/local/share/zabbix/alertscripts	Alert scripts (depends on compile-time installation variable <i>datadir</i>).

Parameter	Mandatory	Range	Default	Description
CacheSize	no	128K-2G	8M	Size of configuration cache, in bytes. Shared memory size for storing host, item and trigger data. The upper limit used to be 1G before version 2.0.5.
CacheUpdateFrequency	no	1-3600	60	How often Zabbix will perform update of configuration cache, in seconds.
DBHost	no		localhost	Database host name. In case of MySQL localhost or empty string results in using a socket. In case of PostgreSQL only empty string results in attempt to use socket.
DBName	yes			Database name. For SQLite3 path to database file must be provided. DBUser and DBPassword are ignored.
DBPassword	no			Database password. Ignored for SQLite. Comment this line if no password is used.
DBPort	no	1024-65535	3306	Database port when not using local socket. Ignored for SQLite.
DBSchema	no			Schema name. Used for IBM DB2.
DBSocket	no		/tmp/mysql.sock	Path to MySQL socket.
DBUser	no			Database user. Ignored for SQLite.
DebugLevel	no	0-4	3	Specifies debug level: 0 - no debug 1 - critical information 2 - error information 3 - warnings 4 - for debugging (produces lots of information)
DisableHousekeeping	no	0-1	0	If set to 1, housekeeping is disabled.
ExternalScripts	no		/usr/local/share/zabbix/external_scripts	Location of external scripts (depends on compile-time installation variable <i>datadir</i>).
Fping6Location	no		/usr/sbin/fping6	Location of fping6. Make sure that fping6 binary has root ownership and SUID flag set. Make empty ("Fping6Location=") if your fping utility is capable to process IPv6 addresses.
FpingLocation	no		/usr/sbin/fping	Location of fping. Make sure that fping binary has root ownership and SUID flag set!

Parameter	Mandatory	Range	Default	Description
HistoryCacheSize	no	128K-2G	8M	Size of history cache, in bytes. Shared memory size for storing history data. The upper limit used to be 1G before version 2.0.5.
HistoryTextCacheSize	no	128K-2G	16M	Size of text history cache, in bytes. Shared memory size for storing character, text or log history data. The upper limit used to be 1G before version 2.0.5.
HousekeepingFrequency	no	1-24	1	How often Zabbix will perform housekeeping procedure (in hours). Housekeeping is removing outdated information from the database. <i>Note:</i> To prevent housekeeper from being overloaded (for example, when history and trend periods are greatly reduced), no more than 4 times HousekeepingFrequency hours of outdated information are deleted in one housekeeping cycle, for each item. Thus, if HousekeepingFrequency is 1, no more than 4 hours of outdated information (starting from the oldest entry) will be deleted per cycle.
Include	no			You may include individual files or all files in a directory in the configuration file. See special notes about limitations.
JavaGateway	no			IP address (or hostname) of Zabbix Java gateway. Only required if Java pollers are started. This option is supported in version 2.0.0 and higher.
JavaGatewayPort	no	1024-32767	10052	Port that Zabbix Java gateway listens on. This option is supported in version 2.0.0 and higher.
ListenIP	no		0.0.0.0	List of comma delimited IP addresses that the trapper should listen on. Trapper will listen on all network interfaces if this parameter is missing. Multiple IP addresses are supported in version 1.8.3 and higher.

Parameter	Mandatory	Range	Default	Description
ListenPort	no	1024-32767	10051	Listen port for trapper.
LogFileSize	no	0-1024	1	Maximum size of log file in MB. 0 - disable automatic log rotation.
LogFile	no			Name of log file.
LogSlowQueries	no	0-3600000	0	How long a database query may take before being logged (in milliseconds). 0 - don't log slow queries. This option becomes enabled starting with DebugLevel=3. Supported since Zabbix 1.8.2 version.
MaxHousekeeperDelete	no	0-1000000	500	No more than 'MaxHousekeeperDelete' rows (corresponding to [tablename], [field], [value]) will be deleted per one task in one housekeeping cycle. SQLite3 does not use this parameter, deletes all corresponding rows without a limit. If set to 0 then no limit is used at all. In this case you must know what you are doing! This parameter is supported since Zabbix 1.8.2 and applies only to deleting history and trends of already deleted items.
NodeID	no	0-999	0	Unique NodeID in distributed setup.
NodeNoEvents	no	0-1	0	0 - standalone server If set to '1' local events won't be sent to master node. This won't impact ability of this node to propagate events from its child nodes.
NodeNoHistory	no	0-1	0	If set to '1' local history won't be sent to master node. This won't impact ability of this node to propagate history from its child nodes.
PidFile	no		/tmp/zabbix_server.pid	Name of PID file.
ProxyConfigFrequency	no	1-604800	3600	How often Zabbix server sends configuration data to a Zabbix proxy in seconds. Used only for proxies in a passive mode. This option is supported in version 1.8.3 and higher.

Parameter	Mandatory	Range	Default	Description
ProxyDataFrequency	no	1-3600	1	How often Zabbix server requests history data from a Zabbix proxy in seconds. Used only for proxies in a passive mode. This option is supported in version 1.8.3 and higher.
SenderFrequency	no	5-3600	30	How often Zabbix will try to send unsent alerts (in seconds).
SNMPTrapperFile	no		/tmp/zabbix_traps.tmp	Temporary file used for passing data from SNMP trap daemon to the server. Must be the same as in zabbix_trap_receiver.pl or SNMPTT configuration file. This option is supported in version 2.0.0 and higher.
SourceIP	no			Source IP address for outgoing connections.
SSHKeyLocation	no			Location of public and private keys for SSH checks and actions
StartDBSyncers	no	1-100	4	Number of pre-forked instances of DB Syncers. The upper limit used to be 64 before version 1.8.5. This option is supported in version 1.8.3 and higher.
StartDiscoverers	no	0-250	1	Number of pre-forked instances of discoverers. The upper limit used to be 255 before version 1.8.5.
StartHTTPPollers	no	0-1000	1	Number of pre-forked instances of HTTP pollers. The upper limit used to be 255 before version 1.8.5.
StartIPMIPollers	no	0-1000	0	Number of pre-forked instances of IPMI pollers. The upper limit used to be 255 before version 1.8.5.
StartJavaPollers	no	0-1000	0	Number of pre-forked instances of Java pollers. This option is supported in version 2.0.0 and higher.
StartPingers	no	0-1000	1	Number of pre-forked instances of ICMP pingers. The upper limit used to be 255 before version 1.8.5.
StartPollersUnreachable	no	0-1000	1	Number of pre-forked instances of pollers for unreachable hosts (including IPMI). The upper limit used to be 255 before version 1.8.5. This option is missing in version 1.8.3.

Parameter	Mandatory	Range	Default	Description
StartPollers	no	0-1000	5	Number of pre-forked instances of pollers. The upper limit used to be 255 before version 1.8.5.
StartProxyPollers	no	0-250	1	Number of pre-forked instances of pollers for passive proxies. The upper limit used to be 255 before version 1.8.5. This option is supported in version 1.8.3 and higher.
StartSNMPTrapper	no	0-1	0	If set to 1, SNMP trapper process will be started. This option is supported in version 2.0.0 and higher.
StartTrappers	no	0-1000	5	Number of pre-forked instances of trappers. Trappers accept incoming connections from Zabbix sender, active agents, active proxies and child nodes. At least one trapper process must be running to display server availability in the frontend. The upper limit used to be 255 before version 1.8.5.
Timeout	no	1-30	3	Specifies how long we wait for agent, SNMP device or external check (in seconds).
TmpDir	no		/tmp	Temporary directory.
TrapperTimeout	no	1-300	300	Specifies how many seconds trapper may spend processing new data.
TrendCacheSize	no	128K-2G	4M	Size of trend cache, in bytes. Shared memory size for storing trends data. The upper limit used to be 1G before version 2.0.5.
UnavailableDelay	no	1-3600	60	How often host is checked for availability during the unavailability period, in seconds.
UnreachableDelay	no	1-3600	15	How often host is checked for availability during the unreachability period, in seconds.
UnreachablePeriod	no	1-3600	45	After how many seconds of unreachability treat a host as unavailable.

Note:

Zabbix supports configuration files only in UTF-8 encoding without BOM.

2 Zabbix proxy

The parameters supported in a Zabbix proxy configuration file:

Parameter	Mandatory	Range	Default	Description
CacheSize	no	128K-2G	8M	Size of configuration cache, in bytes.
ConfigFrequency	no	1-604800	3600	Shared memory size, for storing hosts and items data. The upper limit used to be 1G before version 2.0.5. How often proxy retrieves configuration data from Zabbix server in seconds. For a proxy in the passive mode this parameter will be ignored.
DBHost	no		localhost	Database host name. In case of MySQL localhost or empty string results in using a socket. In case of PostgreSQL only empty string results in attempt to use socket.
DBName	yes			Database name. For SQLite3 path to database file must be provided. DBUser and DBPassword are ignored.
DBPassword	no			Warning: Do not attempt to use the same database Zabbix server is using. Database password. Ignored for SQLite. Comment this line if no password is used.
DBSchema	no			Schema name. Used for IBM DB2.
DBSocket	no		3306	Path to MySQL socket. Database port when not using local socket. Ignored for SQLite.
DBUser				Database user. Ignored for SQLite.
DataSenderFrequency	no	1-3600	1	Proxy will send collected data to the server every N seconds.
DebugLevel	no	0-4	3	Specifies debug level: 0 - no debug 1 - critical information 2 - error information 3 - warnings 4 - for debugging (produces lots of information)
ExternalScripts	no		/usr/local/share/zabbix/externalscripts	Location of external scripts (depends on compile-time installation variable <i>datadir</i>).
Fping6Location	no		/usr/sbin/fping6	Location of fping6. Make sure that fping6 binary has root ownership and SUID flag set. Make empty ("Fping6Location=") if your fping utility is capable to process IPv6 addresses.

Parameter	Mandatory	Range	Default	Description
FpingLocation	no		/usr/sbin/fping	Location of fping. Make sure that fping binary has root ownership and SUID flag set!
HeartbeatFrequency	no	0-3600	60	Frequency of heartbeat messages in seconds. Used for monitoring availability of proxy on server side. 0 - heartbeat messages disabled. For a proxy in the passive mode this parameter will be ignored.
HistoryCacheSize	no	128K-2G	8M	Size of history cache, in bytes. Shared memory size for storing history data. The upper limit used to be 1G before version 2.0.5.
HistoryTextCacheSize	no	128K-2G	16M	Size of text history cache, in bytes. Shared memory size for storing character, text or log history data. The upper limit used to be 1G before version 2.0.5.
Hostname	no		Set by HostnameItem	Unique, case sensitive Proxy name. Make sure the proxy name is known to the server! Allowed characters: alphanumeric, '.', '-', '_' and '-'. Maximum length: 64
HostnameItem	no		system.hostname	Item used for setting Hostname if it is undefined (this will be run on the proxy similarly as on an agent). Does not support UserParameters, performance counters or aliases, but does support system.run[]. Ignored if Hostname is set. This option is supported in version 1.8.6 and higher.

Parameter	Mandatory	Range	Default	Description
HousekeepingFrequency	no	1-24	1	How often Zabbix will perform housekeeping procedure (in hours). Housekeeping is removing outdated information from the database. <i>Note:</i> To prevent housekeeper from being overloaded (for example, when configuration parameters ProxyLocalBuffer or ProxyOfflineBuffer are greatly reduced), no more than 4 times HousekeepingFrequency hours of outdated information are deleted in one housekeeping cycle. Thus, if HousekeepingFrequency is 1, no more than 4 hours of outdated information (starting from the oldest entry) will be deleted per cycle.
Include	no			You may include individual files or all files in a directory in the configuration file. See special notes about limitations.
JavaGateway	no			IP address (or hostname) of Zabbix Java gateway. Only required if Java pollers are started. This option is supported in version 2.0.0 and higher.
JavaGatewayPort	no	1024-32767	10052	Port that Zabbix Java gateway listens on. This option is supported in version 2.0.0 and higher.
ListenIP	no		0.0.0.0	List of comma delimited IP addresses that the trapper should listen on. Trapper will listen on all network interfaces if this parameter is missing. Multiple IP addresses are supported in version 1.8.3 and higher.
ListenPort	no	1024-32767	10051	Listen port for trapper.
LogFileSize	no	0-1024	1	Maximum size of log file in MB. 0 - disable automatic log rotation.
LogFile	no			Name of log file. If not set, syslog is used.

Parameter	Mandatory	Range	Default	Description
LogSlowQueries	no	0-3600000	0	How long a database query may take before being logged (in milliseconds). 0 - don't log slow queries. This option becomes enabled starting with DebugLevel=3. Supported since Zabbix 1.8.2 version.
PidFile	no		/tmp/zabbix_proxy.pid	Name of PID file.
ProxyLocalBuffer	no	0-720	0	Proxy will keep data locally for N hours, even if the data have already been synced with the server. This parameter may be used if local data will be used by third party applications.
ProxyMode	no	0-1	0	Proxy operating mode. 0 - proxy in the active mode 1 - proxy in the passive mode This option is supported in version 1.8.3 and higher.
ProxyOfflineBuffer	no	1-720	1	Proxy will keep data for N hours in case of no connectivity with Zabbix server. Older data will be lost.
ServerPort	no	1024-32767	10051	Port of Zabbix trapper on Zabbix server. For a proxy in the passive mode this parameter will be ignored.
Server	yes			IP address (or hostname) of Zabbix server. Active proxy will get configuration data from the server. For a proxy in the passive mode this parameter will be ignored.
SNMPTrapperFile	no		/tmp/zabbix_traps.tmp	Temporary file used for passing data from SNMP trap daemon to the proxy. Must be the same as in zabbix_trap_receiver.pl or SNMPTT configuration file. This option is supported in version 2.0.0 and higher.
SourceIP	no			Source IP address for outgoing connections.
SSHKeyLocation	no			Location of public and private keys for SSH checks and actions
StartDBSyncers	no	1-100	4	Number of pre-forked instances of DB Syncers. The upper limit used to be 64 before version 1.8.5. This option is supported in version 1.8.3 and higher.

Parameter	Mandatory	Range	Default	Description
StartDiscoverers	no	0-250	1	Number of pre-forked instances of discoverers. The upper limit used to be 255 before version 1.8.5.
StartHTTPPollers	no	0-1000	1	Number of pre-forked instances of HTTP pollers.
StartIPMIPollers	no	0-1000	0	Number of pre-forked instances of IPMI pollers. The upper limit used to be 255 before version 1.8.5.
StartJavaPollers	no	0-1000	0	Number of pre-forked instances of Java pollers. This option is supported in version 2.0.0 and higher.
StartPingers	no	0-1000	1	Number of pre-forked instances of ICMP pingers. The upper limit used to be 255 before version 1.8.5.
StartPollersUnreachable	no	0-1000	1	Number of pre-forked instances of pollers for unreachable hosts (including IPMI). The upper limit used to be 255 before version 1.8.5. This option is missing in version 1.8.3.
StartPollers	no	0-1000	5	Number of pre-forked instances of pollers. The upper limit used to be 255 before version 1.8.5.
StartSNMPTrapper	no	0-1	0	If set to 1, SNMP trapper process will be started. This option is supported in version 2.0.0 and higher.
StartTrappers	no	0-1000	5	Number of pre-forked instances of trappers. Trappers accept incoming connections from Zabbix sender and active agents. The upper limit used to be 255 before version 1.8.5.
Timeout	no	1-30	3	Specifies how long we wait for agent, SNMP device or external check (in seconds).
TmpDir	no		/tmp	Temporary directory.
TrapperTimeout	no	1-300	300	Specifies how many seconds trapper may spend processing new data.
UnavailableDelay	no	1-3600	60	How often host is checked for availability during the unavailability period, in seconds.
UnreachableDelay	no	1-3600	15	How often host is checked for availability during the unreachability period, in seconds.
UnreachablePeriod	no	1-3600	45	After how many seconds of unreachability treat a host as unavailable.

Note:

Zabbix supports configuration files only in UTF-8 encoding without BOM.

3 Zabbix agent (UNIX)

The parameters supported in a Zabbix agent configuration file (*zabbix_agentd.conf*):

Parameter	Mandatory	Range	Default	Description
Alias	no			Sets an alias for an item key. The alias can be used instead of the full item key in item configuration. For example: <i>Alias</i> =system.cpu.load[percpu,5];system.cp <i>Item key</i> : system.cpu.load[percpu,5] <i>Item name</i> : Processor load (\$2 min average per core) An alias can be especially useful to substitute a long and complex item key with a smaller and simpler one. Multiple aliases can be defined.
AllowRoot	no		0	Allow the agent to run as 'root'. If disabled and the agent is started by 'root', the agent will try to switch to user 'zabbix' instead. Has no effect if started under a regular user. 0 - do not allow 1 - allow
BufferSend	no	1-3600	5	Do not keep data longer than N seconds in buffer.
BufferSize	no	2-65535	100	Maximum number of values in a memory buffer. The agent will send all collected data to Zabbix server or proxy if the buffer is full.
DebugLevel	no	0-4	3	Specifies debug level: 0 - no debug 1 - critical information 2 - error information 3 - warnings 4 - for debugging (produces lots of information)
EnableRemoteCommands	no		0	Whether remote commands from Zabbix server are allowed. 0 - not allowed 1 - allowed
Hostname	no		Set by HostnameItem	Unique, case sensitive hostname. Required for active checks and must match hostname as configured on the server. Allowed characters: alphanumeric, '.', '-', '_' and '~'. Maximum length: 64

Parameter	Mandatory	Range	Default	Description
HostnameItem	no		system.hostname	Item used for setting Hostname if it is undefined. Does not support UserParameters, performance counters or aliases, but does support system.run[] regardless of EnableRemoteCommands value. Ignored if Hostname is set. This option is supported in version 1.8.6 and higher.
Include	no			You may include individual files or all files in a directory in the configuration file. See special notes about limitations.
ListenIP	no		0.0.0.0	List of comma delimited IP addresses that the agent should listen on. Multiple IP addresses are supported in version 1.8.3 and higher.
ListenPort	no	1024-32767	10050	Agent will listen on this port for connections from the server.
LogFileSize	no	0-1024	1	Maximum size of log file in MB. 0 - disable automatic log rotation.
LogFile	no			Name of log file.
LogRemoteCommands	no		0	If not set, syslog is used. Enable logging of executed shell commands as warnings. 0 - disabled 1 - enabled
MaxLinesPerSecond	no	1-1000	100	Maximum number of new lines the agent will send per second to Zabbix server or proxy when processing 'log' and 'eventlog' active checks. The provided value will be overridden by the parameter 'maxlines', provided in 'log' or 'eventlog' item key. <i>Note:</i> Zabbix will process 4 times more new lines than set in <i>MaxLinesPerSecond</i> to seek the required string in log items.
PidFile	no		/tmp/zabbix_agentd.pid	Name of PID file.
RefreshActiveChecks	no	60-3600	120	How often list of active checks is refreshed, in seconds. <i>Note:</i> that after failing to refresh active checks the next refresh will be attempted after 60 seconds.

Parameter	Mandatory	Range	Default	Description
Server	no			List of comma delimited IP addresses (or hostnames) of Zabbix servers. No spaces allowed. Incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then '127.0.0.1', '::127.0.0.1', '::ffff:127.0.0.1' are treated equally.
ServerActive	no			IP:port (or hostname:port) of Zabbix server or Zabbix proxy for active checks. Multiple comma-delimited addresses can be provided to use several independent Zabbix servers in parallel. No spaces allowed. If port is not specified, default port is used. IPv6 addresses must be enclosed in square brackets if port for that host is specified. If port is not specified, square brackets for IPv6 addresses are optional. If this parameter is not specified, active checks are disabled.
SourceIP	no			Source IP address for outgoing connections.
StartAgents	no	0-100	3	Number of pre-forked instances of zabbix_agentd that process passive checks. If set to 0, disables passive checks and the agent will not listen on any TCP port. The upper limit used to be 16 before version 1.8.5.
Timeout	no	1-30	3	Spend no more than Timeout seconds on processing
UnsafeUserParameters	no	0,1	0	Allow all characters to be passed in arguments to user-defined parameters. Supported since Zabbix 1.8.2.
UserParameter	no			User-defined parameter to monitor. There can be several user-defined parameters. Format: UserParameter=<key>,<shell command> Note that shell command must not return empty string or EOL only. Example: UserParameter=system.test,who wc -l

Note:

In Zabbix agent 2.0.0 version configuration parameters related to active and passive checks have been changed. See the ["See also"](#) section at the bottom of this page to read more details about these changes.

Note:

Zabbix supports configuration files only in UTF-8 encoding without [BOM](#).

See also

1. [Differences in the Zabbix agent configuration for active and passive checks starting from version 2.0.0](#)

4 Zabbix agent (Windows)

The parameters supported in a Zabbix agent (Windows) configuration file:

Parameter	Mandatory	Range	Default	Description
Alias	no			Sets an alias for an item key. The alias can be used instead of the full item key in item configuration. For example, to retrieve paging file usage percentage from the server, you may use the following item key: <pre>perf_counter[\Paging File(_Total)\% Usage]</pre> However, you may also define an alias by adding the following line to configuration file: <pre>Alias=pg_usage:perf_counter[\Paging File(_Total)\% Usage]</pre> and then use "pg_usage" instead of the full item key in item configuration. An alias can be especially useful to substitute a long and complex item key with a smaller and simpler one. Multiple aliases can be defined. Aliases cannot be reused in PerfCounter parameters.
BufferSend	no	1-3600	5	Do not keep data longer than N seconds in buffer.
BufferSize	no	2-65535	100	Maximum number of values in a memory buffer. The agent will send all collected data to Zabbix server or proxy if the buffer is full.
DebugLevel	no	0-4	3	Specifies debug level: 0 - no debug 1 - critical information 2 - error information 3 - warnings 4 - for debugging (produces lots of information)

Parameter	Mandatory	Range	Default	Description
EnableRemoteCommands	no		0	Whether remote commands from Zabbix server are allowed. 0 - not allowed 1 - allowed
Hostname	no		Set by HostnameItem	Unique, case sensitive hostname. Required for active checks and must match hostname as configured on the server. Allowed characters: alphanumeric, '.', '-', '_' and '-'. Maximum length: 64
HostnameItem	no		system.hostname	Item used for setting Hostname if it is undefined. Does not support UserParameters, performance counters or aliases, but does support system.run[] regardless of EnableRemoteCommands value. Ignored if Hostname is set. This option is supported in version 1.8.6 and higher. See also a more detailed description .
Include	no			You may include individual file in the configuration file.
ListenIP	no		0.0.0.0	List of comma-delimited IP addresses that the agent should listen on. Multiple IP addresses are supported since Zabbix 1.8.3.
ListenPort	no	1024-32767	10050	Agent will listen on this port for connections from the server.
LogFileSize	no	0-1024	1	Maximum size of log file in MB. 0 - disable automatic log rotation.
LogFile	no			Name of log file. If not set, Windows Event Log is used.
LogRemoteCommands	no		0	Enable logging of executed shell commands as warnings. 0 - disabled 1 - enabled

Parameter	Mandatory	Range	Default	Description
MaxLinesPerSecond	no	1-1000	100	Maximum number of new lines the agent will send per second to Zabbix server or proxy processing 'log', 'logrt' and 'eventlog' active checks. The provided value will be overridden by the parameter 'maxlines', provided in 'log', 'logrt' or 'eventlog' item keys.
PerfCounter	no			Syntax: <parameter_name>,"<perf_counter_path>",<period> Defines new parameter <parameter_name> which is an average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds). For example, if you wish to receive average number of processor interrupts per second for last minute, you can define new parameter "interrupts" as following: PerfCounter = interrupts,"\\Processor(0)\\Interrupts/sec",60 Please note double quotes around performance counter path. The parameter name (interrupts) is to be used as the item key when creating an item. Samples for calculating average value will be taken every second. You may run "typeperf -qx" to get list of all performance counters available in Windows.
RefreshActiveChecks	no	60-3600	120	How often list of active checks is refreshed, in seconds. Note that after failing to refresh active checks the next refresh will be attempted after 60 seconds.
Server	no			List of comma delimited IP addresses (or hostnames) of Zabbix servers. No spaces allowed. Incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then '127.0.0.1', '::127.0.0.1', '::ffff:127.0.0.1' are treated equally.

Parameter	Mandatory	Range	Default	Description
ServerActive	no			IP:port (or hostname:port) of Zabbix server or Zabbix proxy for active checks. Multiple comma-delimited addresses can be provided to use several independent Zabbix servers in parallel. No spaces allowed. If port is not specified, default port is used. IPv6 addresses must be enclosed in square brackets if port for that host is specified. If port is not specified, square brackets for IPv6 addresses are optional. If this parameter is not specified, active checks are disabled.
SourceIP	no			Source IP address for outgoing connections.
StartAgents	no	0-63 (*)	3	Number of pre-forked instances of zabbix_agentd that process passive checks. If set to 0, disables passive checks and the agent will not listen on any TCP port. The upper limit used to be 16 before version 1.8.5.
Timeout	no	1-30	3	Spend no more than Timeout seconds on processing
UnsafeUserParameters	no	0-1	0	Allow all characters to be passed in arguments to user-defined parameters. 0 - do not allow 1 - allow
UserParameter				User-defined parameter to monitor. There can be several user-defined parameters. Format: UserParameter=<key>,<shell command> Note that shell command must not return empty string or EOL only. Example: UserParameter=system.test,echo 1

Note:

(*) The number of active servers listed in ServerActive plus the number of pre-forked instances for passive checks specified in StartAgents must be less than 64.

Note:

In Zabbix agent 2.0.0 version configuration parameters related to active and passive checks have been changed. See the **"See also"** section at the bottom of this page to read more details about these changes.

Note:

Zabbix supports configuration files only in UTF-8 encoding without [BOM](#).

See also

1. [Differences in the Zabbix agent configuration for active and passive checks starting from version 2.0.0.](#)

5 Zabbix Java gateway

If you use `startup.sh` and `shutdown.sh` scripts for starting **Zabbix Java gateway**, then you can specify the necessary configuration parameters in file `settings.sh`. The startup and shutdown scripts source the settings file and take care of converting shell variables (listed in the first column) to Java properties (listed in the second column).

If you start Zabbix Java gateway manually by running `java` directly, then you specify the corresponding Java properties on the command line.

Variable	Property	Mandatory	Range	Default	Description
LISTEN_IP	zabbix.listenIP	no		0.0.0.0	IP address to listen on.
LISTEN_PORT	zabbix.listenPort	no	1024-32767	10052	Port to listen on.
PID_FILE	zabbix.pidFile	no		/tmp/zabbix_java.pid	Name of PID file. If omitted, Zabbix Java Gateway is started as a console application.
START_POLLERS	zabbix.startPollers	no	1-1000	5	Number of worker threads to start.
TIMEOUT	zabbix.timeout	no	1-30	3	How long to wait for network operations. This parameter is supported since Zabbix 2.0.15.

Warning:

Port 10052 is not [IANA registered](#).

6 Archive: Zabbix agent (UNIX, Inetd version)

The parameters supported in a Zabbix agent (UNIX, Inetd version) configuration file:

Parameter	Mandatory	Default value	Description
Alias	no		Sets an alias for an item key. The alias can be used instead of the full item key in item configuration. For example: <i>Alias</i> =system.cpu.load[percpu,5]:system.cpu.load[percpu,5] <i>Item key</i> : system.cpu.load[percpu,5] <i>Item name</i> : Processor load (\$2 min average per core) An alias can be especially useful to substitute a long and complex item key with a smaller and simpler one. Multiple aliases can be defined.
Include	no		You may include individual files or all files in a directory in the configuration file. See special notes about limitations.

HP-UX							▼▼				
Solaris							▼▼				
FreeBSD					▼▼						
Linux				▼▼							
2.6											
(and later)											
Linux				▼▼							
2.4											
Windows		▼▼									
Parameter	▼▼										
/ sys-tem											
▼▼	1	2	3	4	5	6	7	8	9	10	11
agent.hostname	X	X	X	X	X	X	X	X	X	X	X
agent.ping	X	X	X	X	X	X	X	X	X	X	X
agent.version	X	X	X	X	X	X	X	X	X	X	X
kernel.maxfiles	-	X	X	X	-	-	-	?	X	X	X
kernel.maxproc	-	-	X	X	X	-	-	?	X	X	X
log[file,<regexp>,<encoding>,<maxlines>,<mode>]						X	X	X	X	X	X
logrt[file_format,<regexp>,<encoding>,<maxlines>,<mode>]							X	X	X	X	X
net.dns[<ip>,<zone>,<type>,<timeout>,<count>]	X				X	X	X	X	X	X	X
net.dns.record[<ip>,<zone>,<type>,<timeout>,<count>]						X	X	X	X	X	X
net.if.collisions[if]		X	X	X	X	-	X	-	X	X	r
net.if.discovery	X	X	X	-	X	X	X	-	-	-	-
net.if.in[if,<mode>]		X	X	X	X	-	X	-	X	X	r
mode bytes	X	X	X	X	X ¹	-	X	-	X	X	r
▲ (de-fault)											
packets	X	X	X	X	X	-	X	-	X	X	r
errors	X	X	X	X	X ¹	-	X	-	X	X	r
dropped	X	X	X	X	-	-	-	-	X	X	r
net.if.out[if,<mode>]		X	X	X	X	-	X	-	X	X	r
mode bytes	X	X	X	X	X ¹	-	X	-	X	X	r
▲ (de-fault)											
packets	X	X	X	X	X	-	X	-	X	X	r
errors	X	X	X	X	X ¹	-	X	-	X	X	r
dropped	X	X	X	-	-	-	-	-	-	-	-
net.if.total[if,<mode>]		X	X	X	X	-	X	-	X	X	r
mode bytes	X	X	X	X	X ¹	-	X	-	X	X	r
▲ (de-fault)											
packets	X	X	X	X	X	-	X	-	X	X	r
errors	X	X	X	X	X ¹	-	X	-	X	X	r
dropped	X	X	X	-	-	-	-	-	-	-	-
net.tcp.listen[port]		X	X	X	X	-	-	-	X	-	-
net.tcp.port[<ip>,<port>]	X		X	X	X	X	X	X	X	X	X
net.tcp.service[service,<ip>,<port>]				X	X	X	X	X	X	X	X
net.tcp.service.perf[service,<ip>,<port>]	X			X	X	X	X	X	X	X	X
net.udp.listen[port]		X	X	X	X	-	-	-	X	-	-
1	2	3	4	5	6	7	8	9	10	11	
proc.mem[<name>,<user>,<mode>,<cmdline>]	X ²						X	X	-	X	X
mode sum	-	X	X	X	X	-	X	X	-	X	X
▲ (de-fault)											
avg	-	X	X	X	X	-	X	X	-	X	X
max	-	X	X	X	X	-	X	X	-	X	X
min	-	X	X	X	X	-	X	X	-	X	X
proc.num[<name>,<user>,<state>,<cmdline>]	X ²						X	X	-	X	X

state	all	-	X	X	X	X	-	X	X	-	X	X
▲	(de-fault)											
	sleep	-	X	X	X	X	-	X	X	-	X	X
	zomb	-	X	X	X	X	-	X	X	-	X	X
	run	-	X	X	X	X	-	X	X	-	X	X
cmdline		-	X	X	X	X	-	X	X	-	X	X
▲												
	sensor[device,sensor,<mode>]	-	-	-	-	-	-	-	-	-	X	-
	system.boottime		X	X	X	X	-	-	-	X	X	X
	system.cpu.intr	-	X	X	X	X	-	X	-	-	X	X
	system.cpu.load[<cpu>,<mode>]	X	X	X	X	X	X	X	X	X	X	X
cpu ▲	all	X	X	X	X	X	X	X	X	X	X	X
	(de-fault)											
	percpu	X	X	X	X	X	X	X	-	X	X	X
mode	avg1	X	X	X	X	X	X	X	X	X	X	X
▲	(de-fault)											
	avg5	X	X	X	X	X	X	X	X	X	X	X
	avg15	X	X	X	X	X	X	X	X	X	X	X
	system.cpu.num[<type>]	X	X	X	X	X	X	X	-	X	X	X
type	online	X	X	X	X	X	X	X	-	X	X	X
▲	(de-fault)											
	max	-	X	X	X	X	-	-	-	X	-	-
	system.cpu.switches		X	X	X	X	-	X	-	-	X	X
	system.cpu.util[<cpu>,<type>,<mode>]	X	X	X	X	X	X	X	X	-	X	X
type	user	-	X	X	X	X	X	X	X	-	X	X
▲	(non-Windows de-fault)											
	nice	-	X	X	X	-	X	-	X	-	X	X
	idle	-	X	X	X	X	X	X	X	-	X	X
	system	X	X	X	X	X	X	X	X	-	X	X
	(Windows de-fault)											
	iowait	-	-	X	-	X	-	X	-	-	-	-
	interrupt	-	-	X	X	-	-	-	-	-	X	-
	softirq	-	-	X	-	-	-	-	-	-	-	-
	steal	-	-	X	-	-	-	-	-	-	-	-
mode	avg1	X	X	X	X	X	X	X	X	-	X	X
▲	(de-fault)											
	avg5	X	X	X	X	X	X	X	-	-	X	X
	avg15	X	X	X	X	X	X	X	-	-	X	X
	1	2	3	4	5	6	7	8	9	10	11	
	system.hostname[<type>]		X	X	X	X	X	X	X	X	X	X
	system.hw.chassis[<info>]		X	-	-	-	-	-	-	-	-	-
	system.hw.cpu[<cpu>,<info>]		X	-	-	-	-	-	-	-	-	-
	system.hw.devices[<type>]		X	-	-	-	-	-	-	-	-	-
	system.hw.macaddr[<interface>,<format>]		-	-	-	-	-	-	-	-	-	-
	system.localtime[<type>]		X	X	X	X	X	X	X	X	X	X
type	utc	X	X	X	X	X	X	X	X	X	X	X
▲	(de-fault)											
	local	X	X	X	X	X	X	X	X	X	X	X
	system.run[command,<mode>]	X	X	X	X	X	X	X	X	X	X	X

mode	wait	X	X	X	X	X	X	X	X	X	X	X
▲	(de-fault)											
	nowait	X	X	X	X	X	X	X	X	X	X	X
	system.stat[resource,<type>]	-	-	-	-	X	-	-	-	-	-	-
	system.sw.arch	-	X	X	-	-	-	-	-	-	-	-
	system.sw.os[<info>]	X	X	-	-	-	-	-	-	-	-	-
	system.sw.packages[<package>,<manager>,<format>]	-	-	-	-	-	-	-	-	-	-	-
	system.swap.in[<device>,<type>]	-	X	-	-	-	-	-	-	X	-	-
	(specifying a device is only supported under Linux)											
type	count	-	X	X	-	X	-	-	-	-	X	-
▲	(pages will only work if device was not specified)											
	sectors	-	X	X	-	-	-	-	-	-	-	-
	pages	-	X	X	-	X	-	-	-	-	X	-
	(de-fault under Linux)											
	system.swap.out[<device>,<type>]	-	X	-	-	-	-	-	-	X	-	-
	(specifying a device is only supported under Linux)											
type	count	-	X	X	-	X	-	-	-	-	X	-
▲	(pages will only work if device was not specified)											
	sectors	-	X	X	-	-	-	-	-	-	-	-
	pages	-	X	X	-	X	-	-	-	-	X	-
	(de-fault under Linux)											

system.swap.size[<device>,<type>]												
(specifying a device is only supported under FreeBSD, for other platforms must be empty or "all")												
type	free	X	X	X	X	X	-	-	X	-	X	-
▲	(default)											
	total	X	X	X	X	X	-	-	X	-	X	-
	used	-	X	X	X	-	-	-	-	-	X	-
	pfree	-	X	X	X	X	-	-	-	-	X	-
	pusd	-	X	X	X	X	-	-	-	-	X	-
system.uname		X	X	X	X	X	X	X	X	X	X	X
system.uptime		X	X	X	X	X	-	X	?	X	X	X
system.users.num			X	X	X	X	X	X	X	X	X	X
	1		2	3	4	5	6	7	8	9	10	11
vfs.dev.read[<device>,<type>,<mode>]												
type	sectors	-	X	X	-	-	-	-	-	-	-	-
▲												
(defaults are different under various OSes)												
	operations		X	X	X	X	-	-	-	-	X	-
	bytes	-	-	-	X	X	-	-	-	-	X	-
	sps	-	X	X	-	-	-	-	-	-	-	-
	ops	-	X	X	X	-	-	-	-	-	-	-
	bps	-	-	-	X	-	-	-	-	-	-	-
mode	avg1	-	X	X	X	-	-	-	-	-	i	-
▲	(default)											
(compatible only with type in: sps, ops, bps)												
	avg5	-	X	X	X	-	-	-	-	-	i	-
	avg15	-	X	X	X	-	-	-	-	-	i	-
vfs.dev.write[<device>,<type>,<mode>]												
		X			X		-	-	-	-	X	-

type	sectors	-	X	X	-	-	-	-	-	-	-	-
▲												
(defaults are different under various OSes)												
	operations		X	X	X	X	-	-	-	-	X	-
	bytes	-	-	-	X	X	-	-	-	-	X	-
	sps	-	X	X	-	-	-	-	-	-	-	-
	ops	-	X	X	X	-	-	-	-	-	-	-
	bps	-	-	-	X	-	-	-	-	-	-	-
mode	avg1	-	X	X	X	-	-	-	-	-	i	-
▲	(de- fault)											
(compatible only with type in: sps, ops, bps)												
	avg5	-	X	X	X	-	-	-	-	-	i	-
	avg15	-	X	X	X	-	-	-	-	-	i	-
	vfs.file.cksum[file]		X	X	X	X	X	X	X	X	X	X
	vfs.file.contents[file,<encoding>]		X	X	X	X	X	X	X	X	X	X
	vfs.file.exists[file]		X	X	X	X	X	X	X	X	X	X
	vfs.file.md5sum[file]		X	X	X	X	X	X	X	X	X	X
	vfs.file.regexp[file,regexp,<encoding>]		X	X	X	X	X	X	X	X	X	X
	vfs.file.regmatch[file,regexp,<encoding>]		X	X	X	X	X	X	X	X	X	X
	vfs.file.size[file]	X	X	X	X	X	X	X	X	X	X	X
		1	2	3	4	5	6	7	8	9	10	11
	vfs.file.time[file,<mode>]	X	X	X	X	X	X	X	X	X	X	X
mode	modify	X	X	X	X	X	X	X	X	X	X	X
▲	(de- fault)											
	access	X	X	X	X	X	X	X	X	X	X	X
	change	X	X	X	X	X	X	X	X	X	X	X
	vfs.fs.discovery	X	X	X	X	X	X	X	-	X	X	X
	vfs.fs.inode[fs,<mode>]	X	X	X	X	X	X	X	X	X	X	X
mode	total	-	X	X	X	X	X	X	X	X	X	X
▲	(de- fault)											
	free	-	X	X	X	X	X	X	X	X	X	X
	used	-	X	X	X	X	X	X	X	X	X	X
	pfree	-	X	X	X	X	X	X	X	X	X	X
	pused	-	X	X	X	X	X	X	X	X	X	X
	vfs.fs.size[fs,<mode>]	X	X	X	X	X	X	X	X	X	X	X
mode	total	X	X	X	X	X	X	X	X	X	X	X
▲	(de- fault)											
	free	X	X	X	X	X	X	X	X	X	X	X
	used	X	X	X	X	X	X	X	X	X	X	X
	pfree	X	X	X	X	X	X	X	X	X	X	X
	pused	X	X	X	X	X	X	X	X	X	X	X
	vm.memory.size[*mode*]		X	X	X	X	X	X	X	X	X	X
mode	total	X	X	X	X	X	X	X	X	X	X	X
▲	(de- fault)											
	active	-	-	-	X	-	X	-	-	X	X	X

anon	-	-	-	-	-	-	-	-	-	-	X
buffers	-	X	X	X	-	-	-	-	-	X	X
cached	X	X	X	X	-	-	X	-	-	X	X
exec	-	-	-	-	-	-	-	-	-	-	X
file	-	-	-	-	-	-	-	-	-	-	X
free	X	X	X	X	X	X	X	X	X	X	X
inactive	-	-	-	X	-	-	-	-	X	X	X
pinned	-	-	-	-	-	-	X	-	-	-	-
shared	-	X	-	X	-	-	-	-	-	X	X
wired	-	-	-	X	-	-	-	-	X	X	X
used	X	X	X	X	X	X	X	X	X	X	X
pused	X	X	X	X	X	X	X	X	X	X	X
available	X	X	X	X	X	X	X	X	X	X	X
pavailable	X	X	X	X	X	X	X	X	X	X	X
web.page.get[host,<path>,<port>]				X	X	X	X	X	X	X	X
web.page.perf[host,<path>,<port>]				X	X	X	X	X	X	X	X
web.page.regex[host,<path>,<port>,<regexp>,<length>]				X	X	X	X	X	X	X	X
	1	2	3	4	5	6	7	8	9	10	11

Footnotes

¹ These values for these items are not supported for loopback interfaces on Solaris systems up to and including Solaris 10 6/06 as byte, error and utilisation statistics are not stored and/or reported by the kernel. However, if you're monitoring a Solaris system via net-snmp, values may be returned as net-snmp carries legacy code from the cmu-snmp dated as old as 1997 that, upon failing to read byte values from the interface statistics returns the packet counter (which does exist on loopback interfaces) multiplied by an arbitrary value of 308. This makes the assumption that the average length of a packet is 308 octets, which is a very rough estimation as the MTU limit on Solaris systems for loopback interfaces is 8232 bytes.

These values should not be assumed to be correct or even closely accurate. They are guestimates. The Zabbix agent does not do any guess work, but net-snmp will return a value for these fields.

² The command line on Solaris, obtained from /proc/pid/psinfo, is limited to 80 bytes and contains the command line as it was when the process was started.

Note:

See also a description of [vm.memory.size parameters](#).

2 vm.memory.size parameters

- **total** - total physical memory available.
- **free** - memory that is readily available to any entity requesting memory.
- **active** - memory currently in use or very recently used, and so it is in RAM.
- **inactive** - memory that is marked as not used.
- **wired** - memory that is marked to always stay in RAM. It is never moved to disk.
- **pinned** - same as 'wired'.
- **anon** - memory not associated with a file (cannot be re-read from it).
- **exec** - executable code, typically from a (program) file.
- **file** - cache for contents of recently accessed files.
- **buffers** - cache for things like file system metadata.
- **cached** - cache for various things.
- **shared** - memory that may be simultaneously accessed by multiple processes.
- **used** - active + wired memory.
- **pused** - active + wired memory in relation to total.
- **available** - inactive + cached + free memory.
- **pavailable** - inactive + cached + free memory in relation to 'total'.

Attention:

The sum of *vm.memory.size[used]* and *vm.memory.size[available]* does not necessarily equal total. For instance, on FreeBSD active, inactive, wired, cached memories are considered used, because they store some useful information. At the same time inactive, cached, free memories are considered available, because these kinds of memories can be given instantly to processes that request more memory. So inactive memory is both used and available simultaneously. Because of this, item *vm.memory.size[used]* is designed for informational purposes only, while item *vm.memory.size[available]* is designed to be used in triggers.

Note:

See the **"See also"** section at the bottom of this page to find more detailed information about memory calculation in different OS.

Platform-specific notes

- on AIX and Solaris **available** and **free** are the same
- on Linux **shared** works only on kernel 2.4

See also

1. [Detailed information about memory calculation in different OS](#)

3 Passive and active agent checks

Overview

This section provides details on passive and active checks performed by Zabbix agent.

Zabbix uses a JSON based communication protocol for communicating with Zabbix agent.

There are some definitions used in the details of protocols used by Zabbix:

<HEADER> - "ZBXD\x01" (5 bytes)

<DATALEN> - data length (8 bytes). 1 will be formatted as 01/00/00/00/00/00/00/00 (eight bytes in HEX, 64

To not exhaust memory (potentially) Zabbix server is limited to accept only 128MB in one connection (unlimited before Zabbix 2.0.3) when using the Zabbix protocol.

Passive checks

A passive check is a simple data request. Zabbix server or proxy asks for some data (for example, CPU load) and Zabbix agent sends back the result to the server.

Server request

<item key>\n

Agent response

<HEADER><DATALEN><DATA>

For example:

1. Server opens a TCP connection
2. Server sends **agent.ping\n**
3. Agent reads the request and responds with **<HEADER><DATALEN>1**
4. Server processes data to get the value, '1' in our case
5. TCP connection is closed

Active checks

Active checks require more complex processing. The agent must first retrieve from the server(s) a list of items for independent processing.

Note:

The servers to get the active checks from are listed in the 'ServerActive' parameter of the agent **configuration file**. The frequency of asking for these checks is set by the 'RefreshActiveChecks' parameter in the same configuration file. However, if refreshing active checks fails, it is retried after hardcoded 60 seconds.

The agent then periodically sends the new values to the server(s).

Getting the list of items

Agent request

```
<HEADER><DATALEN>{
  "request": "active checks",
  "host": "<hostname>"
}
```

Server response

```
{
  "response": "success",
  "data": [
    {
      "key": "log[\\home\\zabbix\\logs\\zabbix_agentd.log]",
      "delay": "30",
      "lastlogsize": "0"
    },
    {
      "key": "agent.version",
      "delay": "600"
    }
  ]
}
```

The server must respond with success. For each returned item, **key** and **delay** must exist. For items having type "Log", the **lastlogsize** must exist as well.

Attention:

If sending of some values fails on the server (for example, because host or item has been disabled or deleted), agent will not retry sending of those values.

For example:

1. Agent opens a TCP connection
2. Agent asks for the list of checks
3. Server responds with a list of items (item key, delay)
4. Agent parses the response
5. TCP connection is closed
6. Agent starts periodical collection of data

Sending in collected data

Agent sends

```
<HEADER><DATALEN>{
  "request": "agent data",
  "data": [
    {
      "host": "<hostname>",
      "key": "log[\\home\\zabbix\\logs\\zabbix_agentd.log]",
      "value": " 13039:20090907:184546.759 zabbix_agentd started. ZABBIX 1.6.6 (revision {7836}).",
      "lastlogsize": 80,
      "clock": 1252926015
    },
    {
      "host": "<hostname>",
      "key": "agent.version",
      "value": "1.6.6",
      "clock": 1252926015
    }
  ],
  "clock": 1252926016
}
```

Server response


```
<HEADER><DATALEN>{
  "response": "success",
  "info": "Processed 2 Failed 0 Total 2 Seconds spent 0.002070"
}
```

For example:

1. Agent opens a TCP connection
2. Agent sends a list of values
3. Server processes the data and sends the status back
4. TCP connection is closed

Older XML protocol

Note:

Zabbix server will take up to 16 MB of XML Base64-encoded data in one connection, but a single decoded value should be no longer than 64 KB otherwise it will be truncated to 64 KB while decoding.

See also

1. [More details on Zabbix agent protocol](#)

4 Encoding of returned values

Zabbix server expects every returned text value in the UTF8 encoding. This is related to any type of checks: zabbix agent, ssh, telnet, etc.

Different monitored systems/devices and checks can return non-ASCII characters in the value. For such cases, almost all possible zabbix keys contain an additional item key parameter - **<encoding>**. This key parameter is optional but it should be specified if the returned value is not in the UTF8 encoding and it contains non-ASCII characters. Otherwise the result can be unexpected and unpredictable.

A description of behavior with different database back-ends in such cases follows.

MySQL

If a value contains a non-ASCII character in non UTF8 encoding - this character and the following will be discarded when the database stores this value. No warning messages will be written to the *zabbix_server.log*.

Relevant for at least MySQL version 5.1.61

PostgreSQL

If a value contains a non-ASCII character in non UTF8 encoding - this will lead to a failed SQL query (PGRES_FATAL_ERROR:ERROR invalid byte sequence for encoding) and data will not be stored. An appropriate warning message will be written to the *zabbix_server.log*.

Relevant for at least PostgreSQL version 9.1.3

5 Large file support

Large file support, often abbreviated to LFS, is the term applied to the ability to work with files larger than 2 GB on 32-bit operating systems. Since Zabbix 2.0 support for large files has been added. This change affects at least **log file monitoring** and all **vfs.file.* items**. Large file support depends on the capabilities of a system at Zabbix compilation time, but is completely disabled on a 32-bit Solaris due to its incompatibility with procfs and swapctl.

6 Unreachable/unavailable host settings

Overview

Several configuration **parameters** define how Zabbix server should behave when an agent check (Zabbix, SNMP, IPMI, JMX) fails and a host becomes unreachable.

Unreachable host

A host is treated as unreachable after a failed agent check (network error, timeout).

From that moment **UnreachableDelay** defines how often a host is rechecked using one of the items (including LLD rules) in this unreachability situation and such rechecks will be performed already by unreachable pollers. By default it is 15 seconds before the next check.

In the Zabbix server log unreachability is indicated by messages like these:

```
Zabbix agent item [system.cpu.load[percpu,avg1]] on host [New host] failed: first network error, wait for  
Zabbix agent item [system.cpu.load[percpu,avg15]] on host [New host] failed: another network error, wait f
```

Note that the exact item that failed is indicated and the item type (Zabbix agent).

Note:

The *Timeout* parameter will also affect how early a host is rechecked during unreachability. If the Timeout is 20 seconds and UnreachableDelay 30 seconds, the next check will be in 50 seconds after the first attempt.

The **UnreachablePeriod** parameter defines how long the unreachability period is in total. By default UnreachablePeriod is 45 seconds. UnreachablePeriod should be several times bigger than UnreachableDelay, so that a host rechecked more than once before a host becomes unavailable.

If the unreachable host reappears, the monitoring returns to normal automatically:

```
resuming Zabbix agent checks on host [New host]: connection restored
```


Unavailable host

After the UnreachablePeriod ends and the host has not reappeared, the host is treated as unavailable.

In the server log it is indicated by messages like these:

```
temporarily disabling Zabbix agent checks on host [New host]: host unavailable
```

and in the frontend the host availability icon goes from green to red (note that on mouseover a tooltip with the error description is displayed):



The **UnavailableDelay** parameter defines how often a host is checked during host unavailability.

By default it is 60 seconds (so in this case "temporarily disabling", from the log message above, will mean disabling checks for one minute).

When the connection to the unavailable host is restored, the monitoring returns to normal automatically, too:

```
enabling Zabbix agent checks on host [New host]: host became available
```

7 Implementation details of net.tcp.service checks

Overview

Implementation of net.tcp.service checks is detailed in this section for various services specified in the service parameter.

ftp

Creates a TCP connection and expects the first 3 characters of the response to be "220" response, then sends "QUIT\n". Default port 21 is used if not specified.

http

Creates a TCP connection without expecting and sending anything. Default port 80 is used if not specified.

https

Uses (and only works with) libcurl, does not verify the authenticity of the certificate, does not verify the host name in the SSL certificate, only fetches the response header (HEAD request). Default port 443 is used if not specified.

imap

Creates a TCP connection and expects the first 4 characters of the response to be "* OK", then sends "a1 LOGOUT\n". Default port 143 is used if not specified.

ldap

Opens a connection to an LDAP server and performs an LDAP search operation with filter set to (objectClass=*). Expects successful retrieval of the first attribute of the first entry. Default port 389 is used if not specified.

nntp

Creates a TCP connection and expects the first 3 characters of the response to be "200", then sends "QUIT\n". Default port 119 is used if not specified.

ntp

Sends an SNTP packet over UDP and validates the response according to [RFC 4330, section 5](#). Default port 123 is used if not specified.

pop

Creates a TCP connection and expects the first 3 characters of the response to be "+OK", then sends "QUIT\n". Default port 110 is used if not specified.

smtp

Creates a TCP connection and expects the first 3 characters of the response to be "220". Then sends "QUIT\r\n". Default port 25 is used if not specified.

ssh

Creates a TCP connection. If the connection has been established, both sides exchange an identification string (SSH-major.minor-XXXX), where major and minor are protocol versions and XXXX is a string. Zabbix checks if the string matching the specification is found and then sends back the string "SSH-major.minor-zabbix_agent\r\n" or "0\n" on mismatch. Default port 22 is used if not specified.

tcp

Creates a TCP connection without expecting and sending anything. Unlike the other checks requires the port parameter to be specified.

telnet

Creates a TCP connection and expects a login prompt (':' at the end). Default port 23 is used if not specified.

5 Triggers

1 Supported trigger functions

All functions supported in **trigger expressions** are listed here:

FUNCTION	Description	Parameters	Comments
abschange			

FUNCTION

The amount of absolute difference between last and previous values.

Supported value types: float, int, str, text, log

For example:
(previous value;last value=abschange)
1;5=4
3;1=2
0;-
2.5=2.5

For strings returns:
0 - values are equal
1 - values differ

avg (sec|#num,<time_shift>)

FUNCTION

Average value of an item within the defined evaluation period.	sec or #num - evaluation period in seconds or in latest collected values (preceded by a hash mark) time_shift (optional) - evaluation point is moved the number of seconds back in time	Supported value types: float, int Examples: => avg(#5) → average value for the five latest values => avg(3600) → average value for an hour => avg(3600,86400) → average value for an hour one day ago. The <code>time_shift</code> parameter is supported since Zabbix 1.8.2. It is useful when there is a need to compare the current average value with the average value <code>time_shift</code> seconds back.

change

FUNCTION	The amount of difference between last and previous values.	Supported value types: float, int, str, text, log For example: (previous value;last value=change) 1;5=+4 3;1=-2 0;-2.5=-2.5 For strings returns: 0 - values are equal 1 - values differ
count (sec #num,<pattern>,<operator>,<time_shift>)		

Number of values within the defined evaluation period.	sec or #num - evaluation period in seconds or in latest collected values (preceded by a hash mark) pattern (optional) - required pattern operator (optional) Supported operators: => <i>eq</i> - equal <i>ne</i> - not equal <i>gt</i> - greater <i>ge</i> - greater or equal <i>lt</i> - less <i>le</i> - less or equal <i>like</i> - matches if contains pattern (case-sensitive) Note that: <i>eq</i> (default), <i>ne</i> , <i>gt</i> , <i>ge</i> , <i>lt</i> , <i>le</i> are supported for integer items <i>eq</i> (default), <i>ne</i> , <i>gt</i> , <i>ge</i> , <i>lt</i> , <i>le</i> are supported for float items <i>like</i> (default), <i>eq</i> , <i>ne</i> are supported for string, text and log items	Supported value types: float, integer, string, text, log Float items match with the precision of 0.000001. Examples: => count(600) → number of values for last 10 minutes => count(10m,"error",eq) → number of values for last 10 minutes that equal 'error' => count(600,12) → number of values for last 10 minutes that equal '12' => count(600,12,gt) → number of values for last 10 minutes that are over '12' => count(#10,12,gt) → number of values within last 10 values that are over '12' => count(600,12,gt,86400) → number of values for preceding 10 minutes up to 24 hours ago that were

FUNCTION			
date		Current date in YYYYM-MDD format.	Supported value types: <i>any</i> Example of returned value: 20150731
dayofmonth		Day of month in range of 1 to 31.	Supported value types: <i>any</i> This function is supported since Zabbix 1.8.5.
dayofweek		Day of week in range of 1 to 7 (Mon - 1, Sun - 7).	Supported value types: <i>any</i>
delta (sec #num,<time_shift>)		Difference between the maximum and minimum values within the defined evaluation period ('max()' minus 'min()').	sec or #num - evaluation period in seconds or in latest collected values specified (preceded by a hash mark) time_shift (optional) - see avg()
diff			Supported value types: float, int The time_shift parameter is supported since Zabbix 1.8.2.

FUNCTION		
	Checking if last and previous values differ.	Supported value types: float, int, str, text, log Returns: 1 - last and previous values differ 0 - otherwise
fuzzytime (sec)		

FUNCTION

Checking how much an item times- tamp value differs from the Zabbix server time.	sec - seconds	Supported value types: float, int Returns: 0 - if difference between item times- tamp value and Zabbix server times- tamp is over T seconds 1 - other- wise. Usually used with sys- tem.localtime to check that local time is in sync with local time of Zabbix server. Can be used also with vfs.file.time[/path/file,mode] key to check that file didn't get updates for long time. Example: => fuzzy- time(60)=0 → detect a problem if time difference is over 60 seconds

iregexp (pattern,<sec|#num>)

FUNCTION			
last (sec #num,<time_shift>)	This function is a non case-sensitive analogue of regexp().	see regexp()	Supported value types: str, log, text

FUNCTION

	The most recent value.	sec (ignored) or #num - the Nth most recent value time_shift (optional) - see avg()	Supported value types: float, int, str, text, log Take note that #num works differently here than with many other functions. For example: last() is always equal to last(#1) last(#3) - third most recent value (<i>not</i> three latest values) Zabbix does not guarantee exact order of values if more than two values exist within one second in history. The #num parameter is supported since Zabbix 1.6.2. The time_shift parameter is supported since Zabbix 1.8.2.
logeventid (pattern)			

FUNCTION

	Check if event ID of the last log entry matches a regular expression.	pattern - regular expression describing the required pattern, POSIX extended style.	Supported value types: log Returns: 0 - does not match, 1 - matches This function is supported since Zabbix 1.8.5.
logseverity	Log severity of the last log entry.		Supported value types: log Returns: 0 - default severity, N - severity (integer, useful for Windows event logs: 1 - Information, 2 - Warning, 4 - Error, 7 - Failure Audit, 8 - Success Audit, 9 - Critical, 10 - Verbose). Zabbix takes log severity from Information field of Windows event log.
logsource (pattern)			

FUNCTION

	Checking if log source of the last log entry matches parameter.	pattern - required string	Supported value types: log Returns: 0 - does not match 1 - matches Normally used for Windows event logs. For example, log-source("VMware Server").
max (sec #num,<time_shift>)	Highest value of an item within the defined evaluation period.	sec or #num - evaluation period in seconds or in latest collected values (preceded by a hash mark) time_shift (optional) - see avg()	Supported value types: float, int The time_shift parameter is supported since Zabbix 1.8.2.
min (sec #num,<time_shift>)	Lowest value of an item within the defined evaluation period.	sec or #num - evaluation period in seconds or in latest collected values (preceded by a hash mark) time_shift (optional) - see avg()	Supported value types: float, int The time_shift parameter is supported since Zabbix 1.8.2.
nodata (sec)			

FUNCTION

Checking for no data received.	sec - eval- uation period in seconds. The period should not be less than 30 seconds.	Supported value types: <i>any</i> Returns: 1 - if no data received during the defined period of time 0 - otherwise Note that this function will display an error if, within the period of the 1st pa- rameter: - there's no data and Zabbix server was restarted - there's no data and main- tenance was com- pleted - there's no data and the item was added or re- enabled Errors are displayed in the <i>Info</i> column in trigger <i>configura-</i> <i>tion</i> .

now

FUNCTION			
prev	Number of seconds since the Epoch (00:00:00 UTC, January 1, 1970).		Supported value types: <i>any</i>
	Previous value.		Supported value types: float, int, str, text, log
regexp (pattern,<sec #num>)			Returns the same as last(#2).
	Checking if the latest (most recent) value matches regular expression.	pattern - regular expression, POSIX extended style. sec or #num (optional) - evaluation period in seconds or in latest collected values (preceded by a hash mark). In this case, more than one value may be processed.	Supported value types: str, text, log Returns: 1 - found 0 - otherwise This function is case-sensitive.
str (pattern,<sec #num>)			

FUNCTION

	Finding a string in the latest (most recent) value.	pattern - required string sec or #num (optional) - evaluation period in seconds or in latest collected values (preceded by a hash mark). In this case, more than one value may be processed.	Supported value types: str, text, log Returns: 1 - found 0 - otherwise This function is case-sensitive.
strlen	(sec #num,<time_shift>)		

	Length of the latest (most recent) value in characters (not bytes).	sec (ignored) or #num - the Nth most recent value time_shift (optional) - see avg()	Supported value types: str, text, log Take note that #num works differently here than with many other functions. Examples: => strlen()(is equal to strlen(#1)) → length of the latest value => strlen(#3) → length of the third most recent value => strlen(,86400) → length of the most recent value one day ago. This function is supported since Zabbix 1.8.4.
sum (sec #num,<time_shift>)	Sum of collected values within the defined evaluation period.	sec or #num - evaluation period in seconds or in latest collected values (preceded by a hash mark) time_shift (optional) - see avg()	Supported value types: float, int The time_shift parameter is supported since Zabbix 1.8.2.

FUNCTION		
time	Current time in HHMMSS format.	Supported value types: <i>any</i> Example of returned value: 123055

Attention:

- 1) All functions return numeric values only. Comparison to strings is not supported, for example.
- 2) Some of the functions cannot be used for non-numeric parameters!
- 3) String arguments should be double quoted. Otherwise, they might get misinterpreted.

6 Macros

1 Macros supported by location

Overview

The table contains a complete list of macros supported by Zabbix:

- **X** means "supported" in that location
- The numbered macro syntax of {MACRO<**1-9**>} is used to reference hosts in the order in which they appear in a trigger expression. Thus, macros like {HOST.IP1}, {HOST.IP2}, {HOST.IP3} will expand to the IP of the first, second and third host in the trigger expression, providing the expression contains those hosts.

	DB	▼	▼DESCRIPTION▼
moni-			
tor			
addi-			
tional			
pa-			
ram-			
eters,			
SSH			
and			
Telnet			
scripts			
Interface		▼	
IP/DNS			
Item		▼	
names			
Trigger		▼	
names			
Trigger		▼	
ex-			
pres-			
sions			
Map	▼		
URLs			

{DISCOVERY.DEVIC&DNS}

DNS
name
of the
dis-
cov-
ered
de-
vice.
Avail-
able
al-
ways,
does
not
de-
pend
on
host
being
added.
Status
of the
dis-
cov-
ered
de-
vice:
can
be
either
UP or
DOWN.

{DISCOVERY.DEVIC&STATUS}

{DISCOVERY.DEVICE&UPTIME}	Time since the last change of discovery status for a particular device. For example: 1h 29m. For devices with status DOWN, this is the period of their downtime.
{DISCOVERY.RULE.NAME}	Name of the discovery rule that discovered the presence or absence of the device or service.

{DISCOVERY.SERVICE.NAME}	Name of the service that was discovered. For example: HTTP.
{DISCOVERY.SERVICE.PORT}	Port of the service that was discovered. For example: 80.
{DISCOVERY.SERVICE.STATUS}	Status of the discovered service: can be either UP or DOWN.

{DISCOVERY.SERVICE.UPTIME}

Time since the last change of discovery status for a particular service. For example: 1h 29m. For services with status DOWN, this is the period of their down-time.

{ESC.HISTORY}			Escalation history. Log of previously sent messages. Shows previously sent notifications, on which escalation step they were sent and their status (<i>sent, in progress or failed</i>).
{EVENT.ACK.HISTORY}			Age of the event. Useful in escalated messages. Date of the event. Numeric event ID which triggered this action. Time of the event.
{EVENT.ACK.STATUS}			
{EVENT.AGE}	X	X	
{EVENT.DATE}	X	X	Date of the event. Numeric event ID which triggered this action. Time of the event.
{EVENT.ID}	X	X	
{EVENT.TIME}	X	X	

{HOST.CONN<1-9>}		X	X ²	X			X		X	X ⁵	IP or host DNS name depending on host settings ³ . Supported in trigger names since 2.0.0.
{HOST.DNS<1-9>}			X	X ²	X		X		X	X ⁵	Host DNS name ³ . Supported in trigger names since 2.0.0.
{HOST.HOST<1-9>}	X	X	X	X			X		X	X ⁵	Host name. {HOSTNAME<1-9>} is deprecated.
{HOST.ID}											Host ID.
{HOST.IP<1-9>}	X	X	X ²	X			X		X	X ⁵	Host IP address ³ . Supported since 2.0.0. {IPADDRESS<1-9>} is deprecated.
{HOST.NAME<1-9>}		X	X	X			X		X	X ⁵	Visible host name. Supported since 2.0.0.

{HOST.PORT}												Host (agent) port ³ . Supported since 2.0.0. Host group ID.
{HOSTGROUP.ID}												
	1	2	3	4	5	6	7	8	9	10	11	12
{INVENTORY.ALIAS<1-9>}												Alias field in host inventory.
{INVENTORY.ASSET.TAG<1-9>}												Asset tag field in host inventory.
{INVENTORY.CHASSIS<1-9>}												Chassis field in host inventory.
{INVENTORY.CONTACT<1-9>}												Contact field in host inventory.
{INVENTORY.CONTRACT.NUMBER<1-9>}												{PROFILE.CONTRACT.NUMBER<1-9>} is deprecated. Contract number field in host inventory.
{INVENTORY.DEPLOYMENT.STATUS<1-9>}												Deployment status field in host inventory.

{INVENTORY.HARDWARE<1-9>}		Hardware field in host inventory. {PROFILE.HARDWARE<1-9>} is deprecated. Hardware (Full details) field in host inventory. Host subnet mask field in host inventory. Host networks field in host inventory. Host router field in host inventory. Hardware architecture field in host inventory.
{INVENTORY.HARDWARE.FULL<1-9>}		
{INVENTORY.HOST.NETMASK<1-9>}		
{INVENTORY.HOST.NETWORKS<1-9>}		
{INVENTORY.HOST.ROUTER<1-9>}		
{INVENTORY.HW.ARCH<1-9>}		

{INVENTORY.HW.DATE.DECOMM<1-9>}	Date hardware de-commissioned field in host inventory.
{INVENTORY.HW.DATE.EXPIRY<1-9>}	Date hardware maintenance expires field in host inventory.
{INVENTORY.HW.DATE.INSTALL<1-9>}	Date hardware installed field in host inventory.
{INVENTORY.HW.DATE.PURCHASE<1-9>}	Date hardware purchased field in host inventory.
{INVENTORY.INSTALLER.NAME<1-9>}	Installer name field in host inventory.

{INVENTORY.LOCATION<1-9>}		Location field in host inventory. {PROFILE.LOCATION<1-9>} is deprecated.
{INVENTORY.LOCATION.LAT<1-9>}		Location latitude field in host inventory.
{INVENTORY.LOCATION.LON<1-9>}		Location longitude field in host inventory.
1 2 3 {INVENTORY.MACADDRESS.A<1-9>}	4 5 6 7 8 9 10 11 12	MAC address A field in host inventory. {PROFILE.MACADDRESS.A<1-9>} is deprecated.
{INVENTORY.MACADDRESS.B<1-9>}		MAC address B field in host inventory.
{INVENTORY.MODEL<1-9>}		Model field in host inventory.

{INVENTORY.NAME<1-9>}	Name field in host inventory. {PROFILE.NAME<1-9>} is deprecated.
{INVENTORY.NOTES<1-9>}	Notes field in host inventory. {PROFILE.NOTES<1-9>} is deprecated.
{INVENTORY.OOB.IP<1-9>}	OOB IP address field in host inventory.
{INVENTORY.OOB.NETMASK<1-9>}	OOB subnet mask field in host inventory.
{INVENTORY.OOB.ROUTER<1-9>}	OOB router field in host inventory.
{INVENTORY.OS<1-9>}	OS field in host inventory. {PROFILE.OS<1-9>} is deprecated.

{INVENTORY.OS.FULL<1-9>}	OS (Full details) field in host inventory.
{INVENTORY.OS.SHORT<1-9>}	OS (Short) field in host inventory.
{INVENTORY.POC.PRIMARY.CELL<1-9>}	Primary POC cell field in host inventory.
{INVENTORY.POC.PRIMARY.EMAIL<1-9>}	Primary POC email field in host inventory.
{INVENTORY.POC.PRIMARY.NAME<1-9>}	Primary POC name field in host inventory.
{INVENTORY.POC.PRIMARY.NOTES<1-9>}	Primary POC notes field in host inventory.
{INVENTORY.POC.PRIMARY.PHONE.A<1-9>}	Primary POC phone A field in host inventory.

{INVENTORY.POC.PRIMARY.PHONE.B<1-9>}												Primary POC phone B field in host inventory.
{INVENTORY.POC.PRIMARY.SCREEN<1-9>}												Primary POC screen name field in host inventory.
{INVENTORY.POC.SECONDARY.CELL<1-9>}												Secondary POC cell field in host inventory.
{INVENTORY.POC.SECONDARY.EMAIL<1-9>}												Secondary POC email field in host inventory.
1	2	3	4	5	6	7	8	9	10	11	12	
{INVENTORY.POC.SECONDARY.NAME<1-9>}												Secondary POC name field in host inventory.
{INVENTORY.POC.SECONDARY.NOTES<1-9>}												Secondary POC notes field in host inventory.
{INVENTORY.POC.SECONDARY.PHONE.A<1-9>}												Secondary POC phone A field in host inventory.

{INVENTORY.POC.SECONDARY.PHONE.B<1-9>}	Secondary POC phone B field in host inventory.
{INVENTORY.POC.SECONDARY.SCREEN<1-9>}	Secondary POC screen name field in host inventory.
{INVENTORY.SERIALNO.A<1-9>}	Serial num- ber A field in host inventory.
{INVENTORY.SERIALNO.B<1-9>}	{PROFILE.SERIALNO.A<1-9>} is dep- re- cated. Serial num- ber B field in host inventory.
{INVENTORY.SITE.ADDRESS.A<1-9>}	Site ad- dress A field in host inventory.
{INVENTORY.SITE.ADDRESS.B<1-9>}	Site ad- dress B field in host inventory.

{INVENTORY.SITE.ADDRESS.C<1-9>}	Site address C field in host inventory.
{INVENTORY.SITE.CITY<1-9>}	Site city field in host inventory.
{INVENTORY.SITE.COUNTRY<1-9>}	Site country field in host inventory.
{INVENTORY.SITE.NOTES<1-9>}	Site notes field in host inventory.
{INVENTORY.SITE.RACK<1-9>}	Site rack location field in host inventory.
{INVENTORY.SITE.STATE<1-9>}	Site state / province field in host inventory.
{INVENTORY.SITE.ZIP<1-9>}	Site ZIP / postal field in host inventory.

{INVENTORY.SOFTWARE<1-9>}												Software field in host inventory. {PROFILE.SOFT 9>} is deprecated.
{INVENTORY.SOFTWARE.APP.A<1-9>}												Software application A field in host inventory.
{INVENTORY.SOFTWARE.APP.B<1-9>}												Software application B field in host inventory.
{INVENTORY.SOFTWARE.APP.C<1-9>}												Software application C field in host inventory.
1	2	3	4	5	6	7	8	9	10	11	12	
{INVENTORY.SOFTWARE.APP.D<1-9>}												Software application D field in host inventory.
{INVENTORY.SOFTWARE.APP.E<1-9>}												Software application E field in host inventory.
{INVENTORY.SOFTWARE.FULL<1-9>}												Software (Full details) field in host inventory.

{INVENTORY.TAG<1-9>}																Tag field in host inven- tory. {PROFILE.TAG-9>} is dep- re- cated.	
{INVENTORY.TYPE<1-9>}																Type field in host inven- tory. {PROFILE.DEVI-9>} is dep- re- cated.	
{INVENTORY.TYPE.FULL<1-9>}																Type (Full de- tails) field in host inven- tory. URL A field in host inven- tory. URL B field in host inven- tory. URL C field in host inven- tory. Vendor field in host inven- tory.	
{INVENTORY.URL.A<1-9>}																	
{INVENTORY.URL.B<1-9>}																	
{INVENTORY.URL.C<1-9>}																	
{INVENTORY.VENDOR<1-9>}																	
1	2	3	4	5	6	7	8	9	10	11	12						

{ITEM.ID<1-9>}		Numeric ID of the Nth item in the trigger expression that caused a notification. Supported since 1.8.12.
{ITEM.DES<1-9>}		Description of the Nth item in the trigger expression that caused a notification. Supported since 2.0.0.

{ITEM.KEY<1-9>}		Key of the Nth item in the trigger expression that caused a notification. Supported since 2.0.0.
{ITEM.KEY<1-9>}		{TRIGGER.KEY} is deprecated. Original key (with macros not expanded) of the Nth item in the trigger expression that caused a notification. Supported since 2.0.6.

{ITEM.LASTVALUE<1-9>}	X	The latest value of the Nth item in the trigger expression that caused a notification. Supported since 1.4.3. It is alias to {{HOST.HOST}}:{{
{ITEM.LOGAGE<1-9>}		
{ITEM.LOGDATE<1-9>}		
{ITEM.LOGEVENTID<1-9>}		
{ITEM.LOGNSEVERITY<1-9>}		
{ITEM.LOGSEVERITY<1-9>}		
{ITEM.LOGSOURCE<1-9>}		
{ITEM.LOGTIME<1-9>}		
{ITEM.NAME<1-9>}		Name of the Nth item in the trigger expression that caused a notification.

{ITEM.NAME.ORIG<1-9>}	Original name (with macros not expanded) of the Nth item in the trigger expression that caused a notification. Supported since 2.0.6.

{ITEM.VALUE<1-9>}

{MAP.ID}												X													Network map ID.		
{NODE.ID<1-9>}												X	X														
{NODE.NAME<1-9>}												X	X														
														1	2	3	4	5	6	7	8	9	10	11	12		
{PROXY.NAME<1-9>}												X	X													Proxy name of the Nth item in the trigger expression that caused a notification. Supported since 1.8.4. Current time in hh:mm:ss. Trigger description. Supported since 2.0.4. {TRIGGER.COMMENT} is deprecated.	
{TIME}												X	X	X													
{TRIGGER<DESCRIPTION>}																											

{TRIGGEREVENTS.ACK}	X	Number of acknowledged events for a map element in maps, or for the trigger which generated current event in notifications. Supported since 1.8.3.
{TRIGGEREVENTS.PROBLEM.ACK}	X	Number of acknowledged PROBLEM events for all triggers disregarding their state. Supported since 1.8.3.

{TRIGGEREVENTS.PROBLEM.UNACK}	X	Number of unacknowledged PROBLEM events for all triggers disregarding their state. Supported since 1.8.3.
{TRIGGEREVENTS.UNACK}	X	Number of unacknowledged events for a map element in maps, or for the trigger which generated current event in notifications. Supported in map element labels since 1.8.3.

{TRIGGER.PROBLEM.EVENTS.PROBLEM.ACK}												X													Number of acknowledged PROBLEM events for triggers in PROBLEM state. Supported since 1.8.3.															
{TRIGGER.PROBLEM.EVENTS.PROBLEM.UNACK}												X													Number of unacknowledged PROBLEM events for triggers in PROBLEM state. Supported since 1.8.3.															
{TRIGGER.EXPRESSION}													1	2	3	4	5	6	7	8	9	10	11	12	Trigger expression. Supported since 1.8.12.															

{TRIGGER\HOSTGROUP.NAME}		Sorted (by SQL query) comma- space sepa- rated list of host groups in which the trig- ger is de- fined. Sup- ported since 2.0.6.
{TRIGGER\ID}	X	Numeric trig- ger ID which trig- gered this ac- tion.
{TRIGGER\NAME}		Name of the trig- ger.
{TRIGGER\NAME.ORIG}		Original name (with macros not ex- panded) of the trig- ger. Sup- ported since 2.0.6.

{TRIGGER\SEVERITY}	Numerical trigger severity. Possible values: 0 - Not classified, 1 - Information, 2 - Warning, 3 - Average, 4 - High, 5 - Disaster. Supported starting from Zabbix 1.6.2.
{TRIGGER\SEVERITY}	Trigger severity name. Can be defined in Administration → General → Trigger severities.
{TRIGGER\STATUS}	Trigger state. Can be either PROBLEM or OK. {STATUS} is deprecated.

{TRIGGERXTEMPLATE.NAME}	Sorted (by SQL query) comma- space sepa- rated list of tem- plates in which the trig- ger is de- fined or *UN- KNOWN* if the trig- ger is de- fined in a host. Sup- ported since 2.0.6. Trigger URL.
{TRIGGERXURL}	

{TRIGGER\VALUE}	X	Current trigger value: 0 - trigger is in OK state, 1 - trigger is in PROBLEM state, 2 - trigger UNKNOWN. This macro can also be used in trigger expressions.

{TRIGGERS.UNACK}	X	Number of unacknowledged triggers for a map element, disregarding trigger state. Trigger is considered to be unacknowledged if at least one of its PROBLEM events is unacknowledged.

{TRIGGERS.PROBLEM.UNACK}	X	Number of unacknowledged PROBLEM triggers for a map element. Trigger is considered to be unacknowledged if at least one of its PROBLEM events is unacknowledged. Supported since 1.8.3.

{TRIGGERS.ACK}	X	Number of acknowledged triggers for a map element, disregarding trigger state. Trigger is considered to be acknowledged if all of it's PROBLEM events are acknowledged. Supported since 1.8.3.

Note:

[1] Macros for map labels are supported since 1.8.

[2] Since 2.0.3 HOST.* macros supported in item keys will only work for item types that have interfaces. E. g. they will not work for types "Zabbix agent (active)", "Calculated" etc.

[3] Depending on the context, the macro will be expanded differently.

In 2.0.1 or earlier version in the context of remote command execution, GUI scripts, item key parameters, and interface IP/DNS fields only the main agent interface will be considered as the source of information.

Since 2.0.2 the macro will use the main agent interface, however, if it is not present, main SNMP interface will be used. If SNMP is also not present, main JMX interface will be used. If JMX is not present either, main IPMI interface will be used.

In item key parameters the interface that is selected for the item will be used, since 2.0.3.

[4] This macro is supported in icon labels and link labels in maps. Only functions **last**, **avg**, **max** and **min** with seconds as arguments are supported in this macro.

[5] Supported since 2.0.3.

[6] Only in trigger expression constants and function parameters.

[7] While supported to build trigger expressions, simple macros may not be used inside each other.

Macro {TRIGGER.ID} is supported in trigger URL since Zabbix 1.8.8.

Additional support for user macros

In addition to the locations listed, user definable macros since Zabbix 2.0 are supported in numerous other locations.

- Hosts
 - Interface IP/DNS
 - Interface port
- Passive proxy
 - Interface port
- Items
 - SNMPv3 security name
 - SNMPv3 auth pass
 - SNMPv3 priv pass
 - SNMPv1/v2 community
 - SNMP OID
 - SNMP port
 - DB monitor additional parameters
 - SSH username
 - SSH public key
 - SSH private key
 - SSH password
 - SSH script
 - Telnet username
 - Telnet password
 - Telnet script
 - Calculated item **formula**
- Discovery
 - * SNMPv3 security name
 - * SNMPv3 auth pass
 - * SNMPv3 priv pass
 - * SNMPv1/v2 community
 - * SNMP OID

Macros used in low-level discovery

There is a type of macro used within the **low-level discovery** function - **{#MACRO}**. It is a macro that is used in an LLD rule and returns real values of file system names, network interfaces and SNMP OIDs.

These macros can be used for creating item, trigger and graph *prototypes*. Then, when discovering real file systems, network interfaces etc., these macros are substituted with real values and are the basis for creating real items, triggers and graphs.

LLD macros can be used:

- for item prototypes in
 - names
 - key parameters
 - SNMP OIDs
 - calculated item formulas
 - SSH and Telnet scripts
 - database monitoring SQL queries
- for trigger prototypes in
 - names
 - expressions (insofar as when referencing an item key prototype)
- for graph prototypes in
 - names

Some low-level discovery macros come "pre-packaged" with the LLD function in Zabbix - `{#FSNAME}`, `{#FSTYPE}`, `{#IFNAME}`, `{#SNMPINDEX}`, `{#SNMPVALUE}`. However, adhering to these names is not compulsory when creating a **custom** low-level discovery rule. Then you may use any other LLD macro name and refer to that name.

7 Setting time periods

1 Format

To set a time period, the following format has to be used:

`d-d, hh:mm-hh:mm`

You can specify more than one time period using a semicolon (;) separator:

`d-d, hh:mm-hh:mm; d-d, hh:mm-hh:mm . . .`

2 Description

Symbol	Description
d	Day of the week: 1 - Monday, 2 - Tuesday ,... , 7 - Sunday
hh	Hours: 00-24
mm	Minutes: 00-59

3 Default

Empty time period specification equals 01-07,00:00-24:00, which is the default value.

Attention:

The upper limit of a time period is not included. Thus, if you specify 09:00-18:00 the last second included in the time period is 17:59:59. This is true starting from version 1.8.7, for everything, while **Working time** has always worked this way.

4 Examples

Working hours. Monday - Friday from 9:00 till 18:00:

`1-5,09:00-18:00`

Working hours plus weekend. Monday - Friday from 9:00 till 18:00 and Saturday, Sunday from 10:00 till 16:00:

`1-5,09:00-18:00;6-7,10:00-16:00`

8 Command execution

Zabbix uses common functionality to execute user parameters, remote commands, `system.run[]` items without the "nowait" flag, scripts (alert, external and global) and some internal commands.

The command/script is executed similarly on both Unix and Windows platforms:

1. Zabbix (the parent process) creates a pipe for communication
2. Zabbix sets the pipe as the output for the to-be-created child process
3. Zabbix creates the child process (runs the command/script)
4. A new process group (in Unix) or a job (in Windows) is created for the child process
5. Zabbix reads from the pipe until timeout occurs or no one is writing to the other end (ALL handles/file descriptors have been closed). Note that the child process can create more processes and exit before they exit or close the handle/file descriptor.
6. If the timeout has not been reached, Zabbix waits until the initial child process exits or timeout occurs
7. At this point it is assumed that everything is done and the whole process tree (i.e. the process group or the job) is terminated

Attention:

Steps 5-7 do not refer to remote commands as they are executed with a "nowait" flag.

Attention:

Zabbix assumes that a command/script has done processing when the initial child process has exited AND no other process is still keeping the output handle/file descriptor open. When processing is done, ALL created processes are terminated.

All double quotes and backslashes in the command are escaped with backslashes and the command is enclosed in double quotes.

Read more about [user parameters](#), [remote commands](#), [alert scripts](#).

9 Recipes for monitoring

General

Monitoring server availability

At least three methods (or combination of all methods) may be used in order to monitor availability of a server.

- ICMP ping ("icmpping" key)
- "zabbix[host,agent,available]" item
- trigger function nodata() for monitoring the availability of hosts that use active checks only

Sending alerts via WinPopUps

WinPopUps maybe very useful if you're running Windows OS and want to get quick notification from Zabbix. It could be good addition for email-based alert messages. Details about enabling of WinPopUps can be found at <http://www.zabbix.com/forum/showthread.php?t=2147>.

Monitoring specific applications

AS/400

IBM AS/400 platform can be monitored using SNMP. More information is available at <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg244504.html?Open>.

MySQL

Several user parameters can be used for the monitoring of MySQL in the agent configuration file: /usr/local/etc/zabbix_agentd.conf

```
### Set of parameters for monitoring MySQL server (v3.23.42 and later)
### Change -u and add -p if required
#UserParameter=mysql.ping,mysqladmin -uroot ping|grep alive|wc -l
#UserParameter=mysql.uptime,mysqladmin -uroot status|cut -f2 -d":"|cut -f2 -d" "
#UserParameter=mysql.threads,mysqladmin -uroot status|cut -f3 -d":"|cut -f2 -d" "
#UserParameter=mysql.questions,mysqladmin -uroot status|cut -f4 -d":"|cut -f2 -d" "
#UserParameter=mysql.slowqueries,mysqladmin -uroot status|cut -f5 -d":"|cut -f2 -d" "
#UserParameter=mysql.qps,mysqladmin -uroot status|cut -f9 -d":"|cut -f2 -d" "
#UserParameter=mysql.version,mysql -V
```

- *mysql.ping*

Check whether MySQL is alive.

Result: 0 - not started 1 - alive

- *mysql.uptime*

Number of seconds MySQL is running.

- *mysql.threads*

Number of MySQL threads.

- *mysql.questions*

Number of processed queries.

- *mysql.slowqueries*

Number of slow queries.

- *mysql.qps*

Queries per second.

- *mysql.version*

Version of MySQL. For example: mysql Ver 14.14 Distrib 5.1.53, for pc-linux-gnu (i686)

For additional information see also the `userparameter_mysql.conf` file in `conf/zabbix_agentd` directory.

Mikrotik routers

Use SNMP agent provided by Mikrotik. See <http://www.mikrotik.com> for more information.

WIN32

Use Zabbix W32 agent included (pre-compiled) into Zabbix distribution.

Novell

Use MRTG Extension Program for NetWare Server (MRTGEXT.NLM) agent for Novell. The agent is compatible with protocol used by Zabbix. It is available from <http://forge.novell.com/modules/xfmod/project/?mrtgext>.

Items have to be configured of type Zabbix Agent and must have keys according to the MRTGEXT documentation.

For example:

- **UTIL1** - 1 minute average CPU utilization
- **CONNMAX** - max licensed connections used
- **VFKSys** - bytes free on volume Sys:

Full list of parameters supported by the agent can be found in `readme.txt`, which is part of the software.

Tuxedo

Tuxedo command line utilities `tmadmin` and `qmadm` can be used in definition of a `UserParameter` in order to return per server/service/queue performance counters and availability of Tuxedo resources.

Informix

Standard Informix utility **onstat** can be used for monitoring of virtually every aspect of Informix database. Also, Zabbix can retrieve information provided by Informix SNMP agent.

HP OpenView

Zabbix can be configured to send messages to OpenView server. The following steps must be performed:

Step 1

Define new media.

The media will execute a script which will send required information to OpenView.

Step 2

Define new user.

The user has to be linked with the media.

Step 3

Configure actions.

Configure actions to send all (or selected) trigger status changes to the user.

Step 4

Write media script.

The script will have the following logic. If trigger is ON, then execute OpenView command `opcmmsg -id application=<application> msg_grp=<msg_grp> object=<object> msg_text=<text>`. The command will return unique message ID which has to be stored somewhere, preferably in a new table of ZABBIX database. If trigger is OFF then `opcmack <message id>` has to be executed with message ID retrieved from the database.

Refer to OpenView official documentation for more details about `opcmmsg` and `opcmack`. The media script is not given here.

10 Performance tuning

Attention:

This is a work in progress.

Overview

It is very important to have Zabbix system properly tuned for maximum performance.

Hardware

General advice on hardware:

- Use fastest processor available
- SCSI or SAS is better than IDE (performance of IDE disks may be significantly improved by using utility `hdparm`) and SATA
- 15K RPM is better than 10K RPM which is better than 7200 RPM
- Use fast RAID storage
- Use fast Ethernet adapter
- Having more memory is always better

Operating system

- Use latest (stable!) version of OS
- Exclude unnecessary functionality from kernel
- Tune kernel parameters

Zabbix configuration parameters

Many parameters may be tuned to get optimal performance.

zabbix_server

StartPollers

General rule - keep value of this parameter as low as possible. Every additional instance of `zabbix_server` adds known overhead, in the same time, parallelism is increased. Optimal number of instances is achieved when queue, on average, contains minimum number of parameters (ideally, 0 at any given moment). This value can be monitored by using internal check `zabbix[queue]`.

Note:

See the ["See also"](#) section at the bottom of this page to find out how to configure optimal count of zabbix processes.

DebugLevel

Optimal value is 3.

DBSocket

MySQL only. It is recommended to use `DBSocket` for connection to the database. That is the fastest and the most secure way.

Database engine

This is probably the most important part of Zabbix tuning. Zabbix heavily depends on the availability and performance of database engine.

- use fastest database engine, i.e. MySQL
- use stable release of a database engine
- rebuild MySQL or PostgreSQL from sources to get maximum performance
- follow performance tuning instructions taken from MySQL or PostgreSQL documentation
- for MySQL, use InnoDB table structure
- ZABBIX works at least 1.5 times faster (comparing to MyISAM) if InnoDB is used. This is because of increased parallelism. However, InnoDB requires more CPU power.
- tuning the database server for the best performance is highly recommended.

- keep database tables on different hard disks
- 'history', 'history_str', 'items', 'functions', 'triggers', and 'trends' are most heavily used tables.
- for large installations keeping MySQL temporary files in tmpfs is:
 - MySQL $\geq$ 5.5: not recommended ([MySQL bug #58421](#))
 - MySQL $<$ 5.5: recommended

GUI debugging

Problems related to the frontend performance may be diagnosed using the frontend **debug mode**.

General advice

- monitor required parameters only
- tune 'Update interval' for all items. Keeping a small update interval may be good for nice graphs, however, this may overload Zabbix
- tune parameters for default templates
- tune housekeeping parameters
- do not monitor parameters which return the same information.
- avoid the use of triggers with long period given as function argument. For example, max(3600) will be calculated significantly slower than max(60).

See also

1. [How to configure optimal count of zabbix processes](#)

11 Version compatibility

Supported agents

Older agents from Zabbix 1.0, Zabbix 1.1.x, Zabbix 1.4.x, Zabbix 1.6.x and Zabbix 1.8.x can still be used with Zabbix 2.0. It does not require any configuration changes on agent side.

However, to take full advantage of new and improved items, improved performance and reduced memory usage, use the latest 2.0 agent.

Supported Zabbix proxies

Only Zabbix 2.0 proxies may be used with Zabbix server. Zabbix 1.6 and 1.8 proxies are not supported with Zabbix 2.0 server.

Zabbix 2.0 proxies may only be used with Zabbix 2.0 server. They will not work with 1.8 or older Zabbix server.

Supported distributed monitoring nodes

All distributed monitoring nodes must be of the same major version. 1.8 nodes are not supported together with 2.0 nodes.

Supported XML files

XML files, exported with 1.8, are supported for import in Zabbix 2.0.

Attention:

In Zabbix 1.8 XML export format, trigger dependencies are stored by name only. If there are several triggers with the same name (for example, having different severities and expressions) that have a dependency defined between them, it is not possible to import them. Such dependencies must be manually removed from the XML file and re-added after import.

Zabbix manpages

These are Zabbix manpages for Zabbix processes.

zabbix_agentd

Section: Maintenance Commands (8)

Updated: 10 November 2011

[Index](#) [Return to Main Contents](#)

NAME

`zabbix_agentd` - Zabbix agent daemon.

SYNOPSIS

`zabbix_agentd [-hpV] [-c <config-file>] [-t <item key>]`

DESCRIPTION

`zabbix_agentd` is a daemon for monitoring of various server parameters.

Options `-c, --config <config-file>`

Use the alternate *config-file* instead of the default one. Absolute path should be specified.

`-p, --print`

Print known items and exit.

`-t, --test <item key>`

Test single item and exit.

`-h, --help`

Display this help and exit.

`-V, --version`

Output version information and exit.

FILES

/usr/local/etc/zabbix_agentd.conf

Default location of Zabbix agent configuration file (if not modified during compile time).

SEE ALSO

`zabbix_get`(8), `zabbix_proxy`(8), `zabbix_sender`(8), `zabbix_server`(8)

AUTHOR

Alexei Vladishev <alex@zabbix.com>

Index

[NAME](#)

[SYNOPSIS](#)

[DESCRIPTION](#)

[Options](#)

[FILES](#)

[SEE ALSO](#)

[AUTHOR](#)

This document was created by man2html, using the manual pages.

Time: 14:45:37 GMT, July 23, 2012

zabbix_get

Section: Maintenance Commands (8)

Updated: 5 July 2011

[Index](#) [Return to Main Contents](#)

NAME

zabbix_get - Zabbix get utility.

SYNOPSIS

zabbix_get [-hV] [-s <host name or IP>] [-p <port number>] [-I <IP address>] [-k <item key>]

DESCRIPTION

zabbix_get is a command line utility for getting data from a remote Zabbix agent.

Options -s, --host <host name or IP>

Specify host name or IP address of a host.

-p, --port <port number>

Specify port number of agent running on the host. Default is 10050.

-I, --source-address <IP address>

Specify source IP address.

-k, --key <item key>

Specify key of item to retrieve value for.

-h, --help

Display this help and exit.

-V, --version

Output version information and exit.

EXAMPLES

zabbix_get -s 127.0.0.1 -p 10050 -k system.cpu.load[all,avg1]

SEE ALSO

[zabbix_agentd\(8\)](#), [zabbix_proxy\(8\)](#), [zabbix_sender\(8\)](#), [zabbix_server\(8\)](#)

AUTHOR

Alexei Vladishev <alex@zabbix.com>

Index

[NAME](#)

[SYNOPSIS](#)

[DESCRIPTION](#)

[Options](#)

[EXAMPLES](#)

[SEE ALSO](#)

This document was created by man2html, using the manual pages.
Time: 14:47:43 GMT, July 23, 2012

zabbix_proxy

Section: Maintenance Commands (8)
Updated: 10 November 2011
[Index Return to Main Contents](#)

NAME

zabbix_proxy - Zabbix proxy daemon.

SYNOPSIS

zabbix_proxy [-hV] [-c <config-file>] [-R <option>]

DESCRIPTION

zabbix_proxy is a daemon used for remote data collection.

Options -c, --config <config-file>

Use the alternate *config-file* instead of the default one. Absolute path should be specified.

-R, --runtime-control <option>

Perform administrative functions according to *option*.

Runtime control options

config_cache_reload

Reload configuration cache. Ignored if cache is being currently loaded. Active Zabbix proxy will connect to the Zabbix server and request configuration data. Default configuration file (unless -c option is specified) will be used to find PID file and signal will be sent to process, listed in PID file.

-h, --help

Display this help and exit.

-V, --version

Output version information and exit.

FILES

/usr/local/etc/zabbix_proxy.conf

Default location of Zabbix proxy configuration file (if not modified during compile time).

SEE ALSO

zabbix_agentd(8), **zabbix_get**(8), **zabbix_sender**(8), **zabbix_server**(8)

AUTHOR

Alexei Vladishev <alex@zabbix.com>

Index

[NAME](#)

[SYNOPSIS](#)

[DESCRIPTION](#)

[Options](#)

[FILES](#)

[SEE ALSO](#)

[AUTHOR](#)

This document was created by man2html, using the manual pages.

Time: 14:48:15 GMT, July 23, 2012

zabbix_sender

Section: Maintenance Commands (8)

Updated: 10 November 2011

[Index Return to Main Contents](#)

NAME

zabbix_sender - Zabbix sender utility.

SYNOPSIS

zabbix_sender [-hpzvIV] {-kso | [-T] -i <inputfile>} [-c <config-file>]

DESCRIPTION

zabbix_sender is a command line utility for sending data to a remote Zabbix server. On the Zabbix server an item of type **Zabbix trapper** should be created with corresponding key. Note that incoming values will only be accepted from hosts specified in **Allowed hosts** field for this item.

Options -c, --config <config-file>

Use *config-file*. Zabbix sender reads server details from the agent configuration file. By default Zabbix sender does not read any configuration file. Absolute path should be specified. Only parameters **Hostname**, **ServerActive** and **SourceIP** are supported (starting with Zabbix 2.0.5). First entry from the **ServerActive** parameter is used.

-z, --zabbix-server <server>

Hostname or IP address of Zabbix server. If a host is monitored by a proxy, proxy hostname or IP address should be used instead.

-p, --port <port>

Specify port number of server trapper running on the server. Default is 10051.

-s, --host <host>

Specify host name as registered in Zabbix frontend. Host IP address and DNS name will not work.

-I, --source-address <IP>

Specify source IP address.

-k, --key <key>

Specify item key to send value to.

-o, --value <value>

Specify value.

-i, --input-file <inputfile>

Load values from input file. Specify - for standard input. Each line of file contains whitespace delimited: <hostname> <key> <value>. Specify - in <hostname> to use hostname from configuration file or --host argument. All entries are sent in a sequential order top-down.

-T, --with-timestamps

Each line of file contains whitespace delimited: <hostname> <key> <timestamp> <value>. This can be used with --input-file option. Timestamp should be specified in Unix timestamp format. If target item has triggers referencing it, all timestamps must be in an increasing order, otherwise event calculation will not be correct.

-r, --real-time

Send values one by one as soon as they are received. This can be used when reading from standard input.

-v, --verbose

Verbose mode, -vv for more details.

-h, --help

Display this help and exit.

-V, --version

Output version information and exit.

EXAMPLES

zabbix_sender -c /usr/local/etc/zabbix_agentd.conf -s Monitored Host -k mysql.queries -o 342.45

Send **342.45** as the value for **mysql.queries** key in **Monitored Host** host using Zabbix server defined in agent daemon configuration file.

zabbix_sender -z 192.168.1.113 -i data_values.txt

Send values from file **data_values.txt** to server with IP **192.168.1.113**. Host names and keys are defined in the file.

echo - hw.serial.number 1287872261 SQ4321ASDF | zabbix_sender -c /usr/local/etc/zabbix_agentd.conf -T -i -

Send a timestamped value from the commandline to Zabbix server, specified in the agent daemon configuration file. Dash in the input data indicates that hostname also should be used from the same configuration file.

SEE ALSO

zabbix_agentd(8), **zabbix_get(8)**, **zabbix_proxy(8)**, **zabbix_server(8)**

AUTHOR

Alexei Vladishev <alex@zabbix.com>

Index

NAME

SYNOPSIS

DESCRIPTION

Options

EXAMPLES

SEE ALSO

AUTHOR

This document was created by man2html, using the manual pages.

Time: 14:48:41 GMT, July 23, 2012

zabbix_server

Section: Maintenance Commands (8)

Updated: 10 November 2011

[Index](#) [Return to Main Contents](#)

NAME

zabbix_server - Zabbix server daemon.

SYNOPSIS

zabbix_server [-hV] [-c <config-file>] [-n <nodeid>] [-R <option>]

DESCRIPTION

zabbix_server is the core daemon of Zabbix software.

Options -c, --config <config-file>

Use the alternate *config-file* instead of the default one. Absolute path should be specified.

-n, --new-nodeid <nodeid>

Convert database data to new *nodeid*. Does not start the server. Run this only once.

-R, --runtime-control <option>

Perform administrative functions according to *option*.

Runtime control options

config_cache_reload

Reload configuration cache. Ignored if cache is being currently loaded. Default configuration file (unless -c option is specified) will be used to find PID file and signal will be sent to process, listed in PID file.

-h, --help

Display this help and exit.

-V, --version

Output version information and exit.

FILES

/usr/local/etc/zabbix_server.conf

Default location of Zabbix server configuration file (if not modified during compile time).

SEE ALSO

zabbix_agentd(8), **zabbix_get**(8), **zabbix_proxy**(8), **zabbix_sender**(8)

AUTHOR

Alexei Vladishev <alex@zabbix.com>

Index

[NAME](#)

[SYNOPSIS](#)

[DESCRIPTION](#)

[Options](#)

[FILES](#)

[SEE ALSO](#)

[AUTHOR](#)

This document was created by man2html, using the manual pages.

Time: 14:49:10 GMT, July 23, 2012

Zabbix API

Zabbix API provides programmable interface to Zabbix for mass manipulations, 3rd party software integration and other purposes.

Currently Zabbix API specification is in draft state. All objects marked as 'draft' are experimental and should be used with a great care. We do not guarantee compatibility with future releases.

Object specifications without the 'draft' mark are stable and can be used for production purposes.

See [Zabbix wiki](#) for community provided solutions around the API.

See full list of API methods in the [Appendixes section](#).