

Server Monitoring

release date. 2026. 09. 11.



Table Of Contents

• Server Monitoring	5
◦ Getting started	7
◦ Support Environment	9
◦ Installation	
■ Overview	15
■ Environment-specific guides	
■ Linux	17
■ Unix	20
■ Windows	30
■ Cases other than the server computer	33
■ Elastic Beanstalk	34
■ Verify Installation	37
■ Agent Update	39
■ Agent deletion	42
◦ Configuration	
■ Basic Configuration	44
■ Configuring the Windows performance counter	49
■ Advanced functions	50
■ GPU Feature	76
◦ Dashboard	
■ Resource board	78
■ Compound Eye	84
■ Server Inventory Map	87
■ Server inventory map details	100

■ Server inventory map grouping	106
■ Flex Board	111
■ Using Flex board	114
■ Managing the widgets	124
◦ Server List	
■ Server Inventory	129
■ Server List	138
■ Server details	148
■ Adding the process group	152
◦ GPU	156
■ GPU Dashboard	161
■ GPU Inventory	165
■ GPU Performance Summary	167
■ GPU Detail	170
◦ Analysis	
■ CUBE	175
■ Metric Anomaly Detection	179
◦ Metrics	182
■ Server metrics	185
■ Server Metrics V2	199
■ Inventory collection items	267
■ VMware collection items	273
■ GPU Metrics	311
■ Metrics Chart	369
■ Metrics Search	372

◦ Log Monitoring	374
◦ Report.....	0
■ Server Report.....	0
■ Integrated Report	0
◦ Alert	0
■ Event Configuration	0
■ Metrics	0
■ Log File.....	0
■ Composite Metrics.....	0
■ Anomaly Detection	0
■ Real-time logs	0
■ Composite log	0
■ Event Configuration	0
■ Event Notification Settings.....	0
■ Event Notification Format.....	0
■ Event History and Monitoring	0
■ Real-time Notifications.....	0
■ System Events	0
◦ Laboratory	0
■ Metrics Explorer	0
■ Process Group Configuration	0
■ File System Operation Info	0
■ NIC Configuration Info	0
■ Time Synchronization Info	0
■ OS Parameter information.....	0

Server Monitoring

WhaTap Server Monitoring deploys in less than a minute, enabling immediate visibility into your infrastructure. Monitor core metrics—such as CPU and Disk I/O—at 5-second intervals for high-fidelity insights, overcoming the limitations of traditional averages. Processes are tracked every 20 seconds to ensure no blind spots.

Access unified dashboards for a comprehensive view of your servers, whether on-premises or in the cloud. WhaTap delivers actionable alerts in real time, so you can respond to incidents before they impact your users.

Automated inventory management, flexible reporting, and multi-channel notifications (email, SMS, and integrations) help your team stay proactive. Scale effortlessly across hybrid environments and optimize system resources with predictive sizing tools.

With WhaTap, ensure reliable, hands-off server operations—no dedicated admin required.

Key features:

- **Fast, unified deployment:**
Install in less than a minute and start monitoring all your servers immediately.
- **Real-time dashboards:**
View comprehensive system health and resource usage across all nodes.
- **Automated inventory management:**
Customize server lists in table view and monitor in real time.
- **Instant, integrated alerts:**
Get notified about anomalies, resource thresholds, and incidents with multi-channel alerts (email, SMS, integrations).
- **Process-level monitoring:**
Automatically collect CPU, memory, and I/O usage for all processes every 20 seconds.
- **Actionable reporting:**
Reduce manual work with built-in reports—add new report types as your needs grow.
- **Cloud and hybrid support:**
Monitor physical, cloud, and hybrid environments with scale-out support and cloud-native metrics.
- **Proactive resource planning:**
Predict and optimize server capacity in real time.

- **Detection of cloud performance vulnerabilities**

WhaTap Server Monitoring provides automatic detection and notification of vulnerabilities that may occur in cloud applications.

- **Determining the appropriate system size**

The required server size differs every year. If you are using the cloud environment, this may differ in real time on a monthly or weekly basis. This is predicted through monitoring, and an appropriate size is calculated.

- **Automatic collection of process performance**

WhaTap Server Monitoring automatically collects the CPU, memory, and IO usage of all processes every 20 seconds.

- **Multi-channel and personalized notification**

We provide personalized notification by using email, SMS, and 3rd party plugin.

- **Cloud, physical server, and hybrid environments** It provides metrics for monitoring cloud environments beyond physical servers, and can be applied to scale-out environments as well.

Getting Started

WhaTap monitoring starts with the following steps:

[Step 1. Sign up](#)

[Step 2. Select a product](#)

[Step 3. Create a project](#)

[Step 4. Check the access key](#)

After confirming the supported environment, proceed to the agent installation step.

Sign up

To use the WhaTap monitoring service, you must first [sign up](#). For more information about signing up, see the [Account Management > Sign Up](#) documentation.

Select a product

1. Log in to the [WhaTap Monitoring Service](#).
2. Select  **All Projects** > **+ Project**.
3. On the **Select product** screen, select the Go product to monitor and then click the Create button.

Create a project

After selecting the product to monitor, create a project before installing the agent.

1. On the **Create project** screen, enter or select the following items to create a project:
 - **Project name** (required): Enter the project name.
 - **Data server region** (required): Select the region where the data server is located. A region is a collection of data centers that provide cloud services. Your data is stored in the selected region.

- **Time zone** (required): Set the time zone that serves as the reference for generating alerts and reports.
- **Notification language setting** (required): Set the language for alert notification messages generated in the project (supports Korean and English).
- **Project groups**: You can group and manage multiple projects. Select a group if there is one to join.
- **Project description**: Enter an additional description for the project.

2. After completing all settings, click the **Creating a project** button.

ⓘ When adding a project with an organization selected, you must set the **Groups of organization** as required. For more information about groups, see [the following documentation](#).

Check the access key

The **access key** is a unique ID for activating the WhaTap service.

In the installation guide section, select the **Getting the access key** button. After automatically issuing the **access key**, proceed to the next step.

ⓘ After creating a project, you are automatically redirected to the **Agent installation** page. If you are not redirected to the **Agent installation** page, select **Management > Agent installation** from the left menu.

Support Environment

Support environments for each OS

Before installing the WhaTap server monitoring agent, check the supported environment.

ⓘ Last updated: December 11, 2025

Linux

Agent Version	RHEL 5.8+	CentOS 6+	Rocky Linux 8.6+	SUSE 10.1+	Amazon Linux	Ubuntu 12.04+	Debian 6.0+
v2.9.5	✓	✓	✓	✓	✓	✓	✓

The minimum supported version of CentOS has been raised to 6. CentOS 6 supports only some features, and all features are supported on CentOS 7 or later.

The supported version of RHEL has not changed; 5.8 or later is supported as before.

Unix

Agent Version	AIX 6.1+	Oracle Solaris 10+	Oracle Solaris 11.2+	HP-UX 11.23+
v1.4.4	✓	✓	✓	✓

FreeBSD has been excluded from the support list. The agent installation screen no longer provides the FreeBSD item.

Windows

Agent Version	Windows 2008 R2+
v2.9.5	

ⓘ If the environment is not included in the support list, contact us at support@whatap.io.

Supporting the browser

WhaTap Monitoring is available in web browsers and mobile Apps.

Browser	Recommended	Supported version
Google Chrome	O	111 or later
Mozilla FireFox	X	Latest version
Edge	X	Latest version
Safari	X	Latest version

ⓘ Support guidance

- For browser compatibility and performance reasons, the latest version of Chrome is recommended.
- The user interface (UI) is implemented using HTML5 standard technologies, so Internet Explorer is not supported.

ⓘ Limitations

WhaTap's web interface does not support mobile browsers. To access WhaTap on a mobile device, install the WhaTap app for Android or iOS. The WhaTap mobile app is optimized for mobile



environments. For more information about the WhaTap mobile app, see [this document](#).

Service deployment models

You can use WhaTap Monitoring as a SaaS or on-premises (self-hosted) service. Whichever you choose, you use the same agents and the same monitoring features.

Deployment	Support	Description
SaaS		Connects to collection servers operated by WhaTap. See the firewall section for regional addresses
On-premises (self-hosted)		Collection servers you build and operate yourself (air-gapped environments supported). Contact your sales representative for the scope



You can also deploy WhaTap on-premises in air-gapped environments without internet connectivity.

Feature differences by deployment model

Whichever you choose, the monitoring features are the same. The mobile app is the only exception and is available on the [service.whatap.io](#) SaaS service only.



To use certain features on-premises, a separate agreement is required.


Firewall

The WhaTap agent must be able to access the collection server through **TCP port 6600**. Allow access to the collection server address that is closest to the monitored target.

Source: WhaTap agent

Cloud	Destination	Destination IP	Port
AWS	WhaTap collection server (Seoul)	13.124.11.223 / 13.209.172.35	TCP 6600
	WhaTap collection server (Tokyo)	52.68.36.166 / 52.193.60.176	TCP 6600
	WhaTap collection server (Singapore)	18.138.0.93 / 18.139.67.236	TCP 6600
	WhaTap collection server (Jakarta)	108.136.91.69 / 108.137.158.44	TCP 6600
	WhaTap collection server (Thailand)	43.209.109.219 / 43.209.48.86	TCP 6600
	WhaTap collection server (Mumbai)	13.127.125.69 / 13.235.15.118	TCP 6600
	WhaTap collection server (California)	52.8.223.130 / 52.8.239.99	TCP 6600
	WhaTap collection server (Virginia)	107.23.220.101 / 54.236.221.105	TCP 6600
	WhaTap collection server (Frankfurt)	3.125.142.162 / 3.127.76.140	TCP 6600
Azure	WhaTap collection server (Seoul)	52.231.66.38 / 20.194.5.115	TCP 6600
	WhaTap collection server (Tokyo)	52.246.169.54 / 20.210.27.232	TCP 6600
Kakao	WhaTap collection server (Seoul)	61.109.237.237 / 61.109.238.166	TCP 6600

If the agent cannot directly access the collection server, use the provided proxy module to access it.



Mobile app

The WhaTap mobile app supports both Android and iOS environments. You can install the WhaTap app by following the link or scanning the QR code. For detailed information about the mobile app, see [the following](#).

iOS	Android
 iOS 12 or later	 Android 5.0 or later

Overview

This document guides you through the installation process for the WhaTap server monitoring agent.

Pre-installation checklist

Before installing the agent, verify the following:

- Confirm that you have completed project creation and access key issuance in [Getting Started](#).
- Check that your operating system is supported in [Supported Environment](#).

Basic installation procedure

1. In the **Agent Installation** screen, select the operating system of the target server to monitor.
2. Download and install the agent package according to the detailed guide for your operating system.
3. Enter the project access key and collection server information in the configuration file ([whatap.conf](#)).
4. Start the agent.

✔ If you have not navigated to the **Agent Installation** screen, select **All projects** > your project from the left menu.

Agent file configuration

The agent file configuration for each operating system is as follows:

- **Common files**
 - [ChangeLog.txt](#): Agent's revision history
 - [whatap.conf](#): In this file, the address of the server to collect data and the server's project access key are entered.

- **Linux/FreeBSD**
 - [whatap_infrad](#): Agent that collects and transmits data.
 - [whatap_infrad.pid](#): File in which the PID of the running agent has been written.
 - [VERSION](#): File in which the versions of the installed agents have been written.
- **Windows**
 - [whatap_infrad.exe](#): Agent that collects and transmits data.
 - [unins000.](#): File to delete any agent
 - [whatap.ico](#): Image of the WhaTap infrastructure icon.

Installation guide by environment

Select the installation guide that matches your environment.

Linux

ⓘ When installing the server monitoring agent, proceed with the installation with a **user account** with the `sudo` role or a **root account**.

If you select a newly generated server monitoring project, the **Agent Installation** screen appears. It is also possible in **Management > Agent Installation** for the project.

Agent Installation

< Choose your operating system

Install guide Unfold Chat contact Video tutorials

- ▶ 1. Access Verification
- ▶ 2. Add the WhaTap Repository.
- ▶ 3. Install the WhaTap Server Monitor Package.
- ▶ 4. Run the Configure Script and Start the Server Monitor Daemon.
- ▶ 5. If you want to update the agent:

Select a desired OS on the **Agent Installation** screen, check the project access key, and then proceed with the installation according to the following instructions.

1. Register the WhaTap repository.

In the **Add the WhaTap Repository** tab, copy the command and enter it.

- [Debian / Ubuntu](#) SUSE CentOS Amazon Linux

```
curl -s https://repo.whatap.io/debian/release.gpg | gpg --dearmor | sudo tee /etc/apt/trusted.gpg.d/whatap-release.gpg > /dev/null
echo "deb [signed-by=/etc/apt/trusted.gpg.d/whatap-release.gpg] https://repo.whatap.io/debian unstable/" | sudo tee /etc/apt/sources.list.d/whatap.list
sudo apt-get update
```

```
sudo rpm --import https://repo.whatap.io/suse/release.gpg
echo "[whatap]" | sudo tee /etc/zypp/repos.d/whatap.repo > /dev/null
echo "name=whatap packages for enterprise linux" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
echo "baseurl=https://repo.whatap.io/suse/12/x86_64" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
echo "enabled=1" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
echo "gpgcheck=1" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
sudo zypper refresh
```

```
sudo rpm --import https://repo.whatap.io/centos/release.gpg
sudo rpm -Uvh https://repo.whatap.io/centos/5/noarch/whatap-repo-1.0-1.noarch.rpm
```

```
sudo rpm --import https://repo.whatap.io/centos/release.gpg
echo "[whatap]" | sudo tee /etc/yum.repos.d/whatap.repo > /dev/null
echo "name=whatap packages for enterprise linux" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
echo "baseurl=https://repo.whatap.io/centos/latest/\$basearch" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
echo "enabled=1" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
echo "gpgcheck=0" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
```

2. Install the **package** with the following command.

In the **Install the WhaTap Server Monitor Package** tab, copy the command and enter it.

- [Debian / Ubuntu](#) [SUSE](#) [CentOS / Amazon Linux](#)

```
sudo apt-get install whatap-infra
```

```
sudo zypper install whatap-infra
```

```
sudo yum install whatap-infra
```

3. Run the configuration script.

In the **Run the Configure Script and Start the Server Monitor Daemon.** tab, get the command or execute the command below after adding the project access key and server IP.

```
echo "license=[Granted Project Access Key]" |sudo tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host=[Granted WhaTap Server IP]" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%N`" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo service whatap-infra restart
```

ⓘ Cautions

TCP outbound **6600 PORT** must be open for data transmission.

ⓘ Server agent installation file configuration



• Common files

- **ChangeLog.txt**: Agent's revision history
- **whatap.conf**: In this file, the address of the server to collect data and the server's project access key are entered.

• Linux/FreeBSD

- **whatap_infrad**: Agent that collects and transmits data.
- **whatap_infrad.pid**: File in which the PID of the running agent has been written.
- **VERSION**: File in which the versions of the installed agents have been written.

Unix

FreeBSD

! End of FreeBSD support

Support for the FreeBSD server monitoring agent has ended. The operating system selection list on the **Agent Installation** screen no longer provides the FreeBSD item, and the server inventory feature is not supported either. The procedure below is reference material for environments where the agent was installed previously.

If you select a newly generated server monitoring project, the **Agent Installation** screen appears. You can also perform this through **Management > Agent Installation** of the project.

Agent Installation



< Choose your operating system



WhaTap Monitor

Supported versions are Ubuntu 12.04 LTS, Ubuntu 14.04 LTS and Debian 7.0. For more information, please refer to the WhaTap Server Monitor User Guide.

Install guide

▼ Unfold

Chat contact

Video tutorials

▶ 1. Access Verification

▶ 2. Add the WhaTap Repository.

▶ 3. Install the WhaTap Server Monitor Package.

▶ 4. Run the Configure Script and Start the Server Monitor Daemon.

▶ 5. If you want to update the agent:

Select a desired OS on the **Agent Installation** screen, check the project access key, and then proceed with the installation according to the following instructions.

1. Register the WhaTap repository.

In the **Add the WhaTap Repository** tab, copy the command and enter it.

```
wget https://repo.whatap.io/freebsd/10/whatap-infra.txz
```

2. Install the package with the following command.

In the **Install the WhaTap Server Monitor Package** tab, copy the command and enter it.

```
pkg install whatap-infra.txz
```

3. Run the configuration script.

In the **Run the Configure Script and Start the Server Monitor Daemon.** tab, get the command or execute the command below after adding the project access key and server IP.

```
echo "license=[Granted Project Access Key]" |sudo tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host=[Granted WhaTap Server IP]" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%N`" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo service whatap-infra restart
```

ⓘ Cautions

TCP outbound **6600 PORT** must be open for data transmission.

ⓘ Server agent file configuration



• Common files

- **ChangeLog.txt**: Agent's revision history
- **whatap.conf**: In this file, the address of the server to collect data and the server's project access key are entered.

• Linux/FreeBSD

- **whatap_infrad**: Agent that collects and transmits data.
- **whatap_infrad.pid**: File in which the PID of the running agent has been written.
- **VERSION**: File in which the versions of the installed agents have been written.

Aix

- It is available in the AIX 6.1 or later.
- Additional options need to be set when installing WhaTap AIX agent 1.3.3 or earlier. See [the following](#).

If you select a newly generated server monitoring project, the **Agent Installation** screen appears. You can also perform this through **Management > Agent Installation** of the project.

Agent Installation

< Choose your operating system

AIX

Supported versions: AIX 6.1 and above

The 6600 port must be opened to transfer data to WhaTap Server (127.0.0.1). (TCP Outbound)

Install guide

Unfold

Chat contact

Video tutorials

▶ 1. Access Verification

▼ 2. Download installation file

▶ 3. Run the setup file

▶ 4. Get started with server monitoring

▶ 5. If you want to update your agent

After you have been issued an access key, you can download the file.

For server security reasons, it is provided as a download only.

whatap-infra-latest.noarch.rpm

After selecting [AIX](#) on the operating system selection screen in **Agent Installation** and confirming the project access key, proceed with the installation according to the following instructions.

1. In the **Download installation file** tab, select [whatap-infra-latest.noarch.rpm](#) to download the agent installation file.
2. Go to the **Run the setup file** tab, check the following command, and then run the package.

```
rpm -Uvh whatap-infra-latest.noarch.rpm
```

ⓘ For server security purposes, the **installation files** are provided only in the download format.

Running the agent

Execute the following command to start the server monitoring. When the application server runs, it begins collecting monitoring information.

```
echo "license=x4t2322hcqshs-x1q9rv0b9rjvs6-x4jnrqfl5ce5g" |tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host=127.0.0.1" |tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=date +%s%S" |tee -a /usr/whatap/infra/conf/whatap.conf
stopsrc -s whatap-infra
startsrc -s whatap-infra
```

Precautions

When installing WhaTap AIX agent 1.3.3 or earlier, the following additional options must be set:

1. In the [Server List](#) menu, select the [Column settings](#) button and then activate [OID](#) in the server information.
2. In [Server list](#), check the [OID](#) column of the target server.
3. Add the OID's fixed option as follows:

```
echo "oid=[oid of server]" >> /usr/whatap/infra/conf/whatap.conf
```

Oracle Solaris

- Additional options need to be set when installing WhaTap Oracle Solaris agent 1.3.3 or earlier. See [the following](#).
- Supported versions: Oracle Solaris 10, 11.2, 11.3, 11.4

If you select a newly generated server monitoring project, the **Agent Installation** screen appears. You can also perform this through **Management > Agent Installation** of the project.

Agent Installation

< Choose Operating System

Oracle Solaris
Supported Version : Oracle Solaris 11.2+
The 6600 port must be opened to transfer data to WhaTap Server (127.0.0.1). (TCP Outbound)

Install guide

Unfold
Chat contact
Video tutorials

1. Access Verification

2. Download installation file

Downloading the installation file is only possible if you have been issued a project access key.
Run the download command below or download the installation file.

Download command

```
wget https://repo.whatap.io/sunos/11/whatap-infra.latest.SPARC.pkg.tar.gz
```

Copy

Download installation file

whatap_infra.zip

3. Run the setup file

4. Get started with server monitoring

5. If you want to update your agent

After selecting [Oracle Solaris](#) on the operating system selection screen in **Agent Installation** and confirming the project access key, proceed with the installation according to the following instructions.

1. Execute the download command or select a version in the **Download installation file** tab, and then download the installation file.

- Execute the following command to download the agent installation file.

```
wget https://repo.whatap.io/sunos/11/whatap-infra.latest.SPARC.pkg.tar.gz
```

- Select [whatap_infra.zip](#) to download the agent installation file.

2. Go to the **Run the setup file** tab, check the following command, and then run the package.

```
tar xzvf whatap-infra.latest.SPARC.pkg.tar.gz
pkgadd -d .
```

Running the agent

Execute the following command to start the server monitoring. When the application server runs, it begins collecting monitoring information.

```
echo "license=x4t2322hcqshs-x1q9rv0b9rjvs6-x4jnrqfl5ce5g" |tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host=127.0.0.1" |tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%S`" |tee -a /usr/whatap/infra/conf/whatap.conf
/etc/init.d/whatap-infra start
```

Precautions

Additional options need to be set as follows when installing WhaTap Oracle Solaris agent 1.3.3 or earlier.

1. In the **Server List** menu, select the **Column settings** button and then activate **OID** in the server information.
2. In **Server list**, check the **OID** column of the target server.
3. Add the OID's fixed option as follows:

```
echo "oid=[oid of server]" >> /usr/whatap/infra/conf/whatap.conf
```

! Server running on Oracle Enterprise Edition

If you install on the server running on Oracle Enterprise Edition, disable the process collection option by referencing the following:

```
echo "task_proc=false" >> /usr/whatap/infra/conf/whatap.conf
/etc/init.d/whatap-infra restart
```

Addressing the agent incidents

The following provides guidance on how to address the cases where no data is incorrectly detected in the Oracle Solaris agent.

Checking the OID settings

Check that the `oid` option is set in the `/usr/whatap/infra/conf/whatap.conf` file.

- For the agent 1.3.3 or earlier, set the OID manually. For the setting method, see [the following](#).
- For the agent 1.3.4 or later, the OID is set automatically.

In case the version of the agent is 1.3.3 or earlier and it is running on Oracle Enterprise Edition

Process collection is not supported on the servers running on Oracle Enterprise Edition.

1. Check that the `task_proc=false` option is set in the `/usr/whatap/infra/conf/whatap.conf` file.
2. If it is not set, execute the following command to set the option and then restart the agent.

```
echo "task_proc=false" >> /usr/whatap/infra/conf/whatap.conf
/etc/init.d/whatap-infra restart
```

This process disables the process collection option, ensuring that the agent operates in environments running on Oracle Enterprise Edition.

HP-UX

- It is supported from HP-UX Itanium 11.23 and later.
- It is supported from HP PA-RISC-UX 11.23 and later.

If you select a newly generated server monitoring project, the **Agent Installation** screen appears. You can also

perform this through **Management > Agent Installation** of the project.

Agent Installation

< Choose Operating System

HP-UX
 Supported Version : HP-UX Itanium : 11.23, HP PA-RISC-UX : 11.23
 The 6600 port must be opened to transfer data to WhaTap Server (127.0.0.1). (TCP Outbound)

Install guide [Unfold] [Chat contact] [Video tutorials]

IA64 PA-RISC

► 1. Access Verification

▼ 2. Download installation file

Downloading the installation file is only possible if you have been issued a project access key.
 Download the installation file below.

11.23 Version **11.31 Version**

↓ whatap_infra.zip ↓ whatap_infra.zip

► 3. Run the setup file

► 4. Get started with server monitoring

► 5. If you want to update your agent

In the OS selection screen in **Agent Installation**, select **HP-UX**. Depending on your HP-UX server architecture, select **IA64** or **PA-RISC** from the tab at the top and then check the project access key. Proceed with installation according to the following instructions.

- **IA64 PA-RISC**

1. In the **Download installation file** tab, select the appropriate version of **whatap_infra.zip** to download the agent installation file.
2. Go to the **Run the setup file** tab, check the following command, and then run the package.

```
gunzip -c whatap-infra.latest.ia64.tar.gz | tar xvf -
./whatap-infra/install.sh x4t2322hcqshs-x1q9rv0b9rjvs6-x4jnrqfl5ce5g 127.0.0.1
```


1. In the **Download installation file** tab, select [whatap_infra.zip](#) to download the agent installation file.
2. Go to the **Run the setup file** tab, check the following command, and then run the package.

```
gunzip -c whatap-infra.latest.pa.tar.gz | tar xvf -  
./whatap-infra/install.sh x4t2322hcqshs-x1q9rv0b9rjvs6-x4jnrqfl5ce5g 127.0.0.1
```

Running the agent

Execute the following command to start the server monitoring. When the application server runs, it begins collecting monitoring information.

```
/sbin/init.d/whatap-infra start
```

Windows

If you select a newly generated server monitoring project, the **Agent Installation** screen appears. It is also possible in **Management > Agent Installation** for the project.

Agent Installation

< Choose your operating system



Windows

Supported versions are Windows Server 2008 R2 and above
The 6600 port must be opened to transfer data to WhaTap Server (127.0.0.1). (TCP Outbound)

Install guide

Unfold

Chat contact

Video tutorials

▶ 1. Access Verification

▼ 2. Download the WhaTap Installer for Windows.

whatap_infra.exe

whatap_infra.zip

By click the link below, you can download the WhaTap installer file (whatap.exe / whatap.zip) for windows servers.
setup.zip file is also provided for customers using browsers whose security settings prevent downloading .exe file.
* For your server security, it is recommended that running the downloaded installer file than direct installation through a browser.

▶ 3. Run the WhaTap Installer.

▶ 4. If you want to update the agent

After selecting [Windows](#) on the operating system selection screen in **Agent Installation** and confirming the project access key, proceed with the installation according to the following instructions.

1. Download the **agent file**.

In the **Download the WhaTap Installer for Windows** tab, select [whatap_infra.exe](#) or [whatap_infra.zip](#) to download the agent installation file.

- ❗ • The [.zip](#) file is also provided for the users who cannot accept [.exe](#) files due to security issues.
- For server security, it is recommended to run the downloaded installation file rather than installing through the browser.

2. Access the server to be monitored and then upload the **agent file**.
3. Double-click the downloaded **installation file** to run it.
4. **3** Go to **Run the WhaTap Installer** tab and then check the granted project access key and server IP.
5. According to the on-screen instructions, enter the project access key in [License Key](#) and the IP address of the WhaTap server in [WhatTap IP](#).
6. Click [Next](#).
7. In the following screen, click [Install](#). Proceed with the installation.
8. Click [Finish](#) to finish the installation.

Signature verification of the installation file in a closed network

On closed-network servers isolated from the internet, Windows cannot receive automatic root certificate updates, so signature verification sometimes failed even for properly signed installation files. The installer verifies the signature by directly comparing the root certificate fingerprint included in the installer with the signing certificate chain, without relying on the trust store of the operating system. It does not change the certificate store of the server, so you can proceed with the installation in a closed network without any separate action.

Files with a corrupted signature or hash, and files signed with an expired or revoked certificate, are still rejected for installation.

Server agent installation file configuration

- ❗ • **Common files**
 - [ChangeLog.txt](#): Agent's revision history
 - [whatap.conf](#): In this file, the address of the server to collect data and the server's project



access key are entered.

- **Windows**

- [whatap_infrad.exe](#): Agent that collects and transmits data.
- [unins000.](#): File to delete any agent
- [whatap.ico](#): Image of the WhaTap infrastructure icon.

Cases Other Than The Server Computer

Apply the following option in an environment where a personal computer (not a server computer, such as a laptop) and Wi-Fi with a floating IP address are used.

Default value change

Since agent 2.9.21, the default value of the `ipreg` option is `false`. Therefore, even in environments where the IP address changes frequently, the agent identifier determined at first is kept as is without any separate setting. Use the settings below only when you want to specify the default value explicitly.

Autoscale environments are an exception. If you use the `autoscale` option or specify the `autoscale.group` value, the agent switches the `ipreg` value to `true`. Servers in an autoscale environment are short-lived and must be registered again based on the current IP address every time they start.

- **Linux** Windows

```
echo "ipreg=false" | sudo tee -a /usr/whatap/infra/conf/whatap.conf  
sudo service whatap-infra restart
```

SH

```
// Execute the following command in the powershell with an admin role.  
Add-Content "c:\Program Files\WhatapInfra\whatap.conf" "ipreg=false"  
Restart-Service "Whatap Infra"
```

ⓘ It can be easily used because it is not automatically registered whenever the IP is changed.

Elastic Beanstalk

The following explains how to apply the WhaTap agent when deploying with **Elastic Beanstalk**.

1. Depending on the Linux installation environment, add the following Linux to the Elastic Beanstalk's **installation script**.

❗ Enter the `whatap.server.host` and `license` that is the **Project access key** to your taste.

- [Debian / Ubuntu](#) [CentOS](#) [Amazon Linux](#)

```
```bash
wget https://repo.whatap.io/debian/release.gpg -O - | sudo apt-key add -
wget https://repo.whatap.io/debian/whatap-repo_1.0_all.deb
sudo dpkg -i whatap-repo_1.0_all.deb
sudo apt-get update
sudo apt-get install whatap-infra
echo "license={LICENSE_INFO}" | sudo tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host={HOST_IP}" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%N`" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo service whatap-infra restart
```
```

```
```bash
sudo rpm --import https://repo.whatap.io/centos/release.gpg
sudo rpm -Uvh https://repo.whatap.io/centos/5/noarch/whatap-repo-1.0-1.noarch.rpm
sudo yum install whatap-infra
echo "license={LICENSE_INFO}" | sudo tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host={HOST_IP}" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%N`" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo service whatap-infra restart
```
```

```
```bash
sudo rpm --import https://repo.whatap.io/centos/release.gpg
echo "[whatap]" | sudo tee /etc/yum.repos.d/whatap.repo > /dev/null
```
```

```

echo "name=whatap packages for enterprise linux" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
echo "baseurl=https://repo.whatap.io/centos/latest/\$basearch" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
echo "enabled=1" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
echo "gpgcheck=0" | sudo tee -a /etc/yum.repos.d/whatap.repo > /dev/null
sudo yum install whatap-infra
echo "license={LICENSE_INFO}" | sudo tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host={HOST_IP}" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=date +%s%N" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo service whatap-infra restart
'''

```

- **SUSE FreeBSD**

```

'''bash
sudo rpm --import https://repo.whatap.io/suse/release.gpg
echo "[whatap]" | sudo tee /etc/zypp/repos.d/whatap.repo > /dev/null
echo "name=whatap packages for enterprise linux" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
echo "baseurl=https://repo.whatap.io/suse/12/x86_64" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
echo "enabled=1" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
echo "gpgcheck=1" | sudo tee -a /etc/zypp/repos.d/whatap.repo > /dev/null
sudo zypper refresh
sudo zypper install whatap-infra
echo "license={LICENSE_INFO}" | sudo tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host={HOST_IP}" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%N`" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo /etc/init.d/whatap-infra restart
'''

```

```

'''bash
wget https://repo.whatap.io/freebsd/10/whatap-infra.txz
pkg install whatap-infra.txz
echo "license={LICENSE_INFO}" | tee /usr/whatap/infra/conf/whatap.conf
echo "whatap.server.host={HOST_IP}" | tee -a /usr/whatap/infra/conf/whatap.conf
echo "createdtime=`date +%s%N`" | tee -a /usr/whatap/infra/conf/whatap.conf
service whatap_infra restart
'''

```

1. Modify the following **command**.

```
echo "createdtime=date +%s%N" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

By including the server monitoring into the Beanstalk after installation, a new list for server monitoring may be generated repeatedly in every release of applications.

This is because the WhaTap agent's internal ID is newly issued each time the above command is executed. Modify as follows to prevent the new list from generating in a later release.

```
echo "createdtime=12345678" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
```


Verify Installation

The server monitoring has been installed completely. Check the checklist after installation.

Linux and Unix

Checking the log output

If the agent has been installed, you can check the log output as follows. After connection, the message, "connected" appears.

- [Debian/Ubuntu/SUSE](#) CentOS/Amazon Linux/Unix

```
tail -10 /var/log/syslog
```

```
...
```

```
Apr 8 16:28:08 localhost systemd[1]: Started LSB: Start whatap-infra daemon.
```

```
Apr 8 16:28:08 localhost whatap_infrad[32468]: 2019/04/08 16:28:08 Config file : /usr/whatap/infra/conf/whatap.conf
```

```
Apr 8 16:28:08 localhost whatap_infrad[32468]: 2019/04/08 16:28:08 OID: 1193741647 ONAME: localhost
```

```
Apr 8 16:28:08 localhost whatap_infrad[32468]: 2019/04/08 16:28:08 connected to {whatap ip}:{whatap port}
```

```
tail -10 /var/log/messages
```

```
...
```

```
Apr 8 16:28:08 localhost systemd[1]: Started LSB: Start whatap-infra daemon.
```

```
Apr 8 16:28:08 localhost whatap_infrad[32468]: 2019/04/08 16:28:08 Config file : /usr/whatap/infra/conf/whatap.conf
```

```
Apr 8 16:28:08 localhost whatap_infrad[32468]: 2019/04/08 16:28:08 OID: 1193741647 ONAME: localhost
```

```
Apr 8 16:28:08 localhost whatap_infrad[32468]: 2019/04/08 16:28:08 connected to {whatap ip}:{whatap port}
```

Checking the firewall and network settings

If connection is not made, check the network settings including the firewall settings. When the network settings are applied, the agent is automatically connected.

Windows

After installation is completed, you can see the following screen at the final step. The agent automatically starts the monitoring.

Cautions

TCP outbound **6600 PORT** must be open for data transmission.

Agent Update

! To install a previous version, contact us at support@whatap.io.

Linux

To update the installed agent to the latest version, execute the following command.

- **Debian/Ubuntu**

```
sudo apt-get update  
sudo apt-get install whatap-infra  
sudo service whatap-infra restart
```

To install a specific version, specify the version during installation as follows:

```
sudo apt-get install whatap-infra={version}
```

- **CentOS / Amazon Linux / SUSE**

```
sudo yum update whatap-infra  
sudo service whatap-infra restart
```

To install a specific version, specify the version during installation as follows:

```
sudo yum install whatap-infra-{version}
```

Unix

- **FreeBSD**

To update the installed agent to the latest version, execute the following command.

```
wget https://repo.whatap.io/freebsd/10/whatap-infra.txz
pkg install whatap-infra.txz
```

- **AIX**

To update the installed agent to the latest version, execute the following command.

```
stopsrc -s whatap-infra
cp /usr/whatap/infra/conf/whatap.conf /tmp/whatap.conf
rpm -e whatap-infra
rpm -Uvh https://repo.whatap.io/aix/noarch/whatap-infra-latest.noarch.rpm
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
startsrc -s whatap-infra
```

- **Oracle Solaris**

To update the installed agent to the latest version, execute the following command.

```
cp /usr/whatap/infra/conf/whatap.conf /tmp
/etc/init.d/whatap-infra stop
pkgrm whatap-infra
wget https://repo.whatap.io/sunos/11/whatap-infra.latest.SPARC.pkg.tar.gz
tar xzvf whatap-infra.latest.SPARC.pkg.tar.gz
pkgadd -d .
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
/etc/init.d/whatap-infra start
```

- **HP-UX**

To update the installed agent to the latest version, execute the following command.

- **IA64 PA-RISC**

IA64

```
cp /usr/whatap/infra/conf/whatap.conf /tmp
/sbin/init.d/whatap-infra stop
gunzip -c whatap-infra.latest.ia64.tar.gz | tar xvf -
yes n | ./whatap-infra/install.sh
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
```

```
/sbin/init.d/whatap-infra start
```

PA-RISC

```
cp /usr/whatap/infra/conf/whatap.conf /tmp
/sbin/init.d/whatap-infra stop
gunzip -c whatap-infra.latest.pa.tar.gz | tar xvf -
yes n | ./whatap-infra/install.sh
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
/sbin/init.d/whatap-infra start
```

Windows

To update the installed agent to the latest version, go to **Management > Agent Installation**. Download the installation file for your Windows Server OS by following the instructions in the **Download the WhaTap Installer for Windows** section of **Install guide**, and then run the `whatap_infra.exe` file. The service is restarted after updating the agent.

ⓘ The `.zip` file is also provided for the users who cannot accept `.exe` files due to security issues. For more information, see [the following](#).

Agent Deletion

Linux

To delete the installed agent, execute the following command.

- **Debian/Ubuntu**

```
sudo apt-get remove whatap-infra
```

- **CentOS / Amazon Linux / SUSE**

```
sudo yum remove whatap-infra
```

Unix

- **FreeBSD**

```
pkg delete whatap-infra
```

- **AIX**

```
stopsrc -s whatap-infra  
rpm -e whatap-infra  
rm -rf /usr/whatap/infra
```

- **Oracle Solaris**

```
pkgrm whatap-infra  
rm -rf /usr/whatap/infra
```

- **HP-UX**

IA64 / PA-RISC

```
swremove whatap-infra
```

Windows

- Execute the server monitoring agent deletion command in the command line with the admin privilege.

```
"C:\Program Files\WhatapInfra\unins000.exe"
```

- You can also use the Uninstall a Program function in Control Panel.

Basic Configuration

The WhaTap server monitoring agent controls its operations through the [whatap.conf](#) configuration file. This document guides you through the basic agent configuration.

! Configuration file location

- **Unix/Linux:** `$WHATAP_HOME/whatap.conf`
- **Windows:** `%WHATAP_HOME%\whatap.conf`

After changing the configuration, you must restart the agent for the changes to take effect.

Agent name identification

WhaTap basically uses the data collected from the server for identifying the infrastructure server that is the target of monitoring data collection. The default value is the server's hostname.

- ! • The agent names must be unique on a project basis.
- If the agent ID and name have been changed, previous data is not connected.

- ! To see the user-defined agent name in [WhaTap Monitoring Service](#), restart the server.

For a user-friendly name, you can specify it in the configuration file as follows.

- **Linux Shell**

```
export WHATAP_SERVER_NAME=user-friendly name

echo "oname=$WHATAP_SERVER_NAME" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

- **Windows Power Shell**


```
# The admin role is required.
```

```
$WHATAP_SERVER_NAME="user-friendly name"
```

```
Add-Content "C:\Program Files\WhatapInfra\whatap.conf" "oname=$WHATAP_SERVER_NAME"
```

Network configuration

WhaTap transfers monitoring data after TCP connection between the server and the agent.

WhaTap Agent Network

The agent handles data transfer and server control requests through a single TCP session. The node agent does not use UDP communication. Open the firewall from the node agent to the WhaTap collection server.

Collection server address and port

WhaTap servers are divided into data region servers, front servers, and Eureka. The data regions include Proxy, Yard, Keeper, etc. Among them, the agent communicates with the proxy server.

Set the proxy server address of the Whatap server in the Node agent (e.g. `whatap.server.host=10.0.3.1/10.0.3.2`). When setting the server address, enter as many addresses as the proxy server count. WhaTap server can use one or multiple proxy servers depending on the installation method.

WhaTap Proxy server listens on the port 6600. If there is no separate setting in the agent, the agent attempts to connect the port 6600.

whatap.conf

```
whatap.server.port=6600
```

⚠️ No proxy servers can use different ports. When using multiple proxy servers, the listening port must be the same.

Connection and security

WhaTap is designed to collect monitoring data over public networks. Therefore, all monitoring data is encrypted before being transmitted to the server. To minimize overhead, WhaTap selectively encrypts data. The communication process between the agent and server is as follows:

1. Generate and copy the **project access key** when creating a project.
2. The **project access key** contains a secret key, so ensure it is not exposed externally.
3. When the agent starts, it establishes a TCP session with the WhaTap collection server.
4. The agent encrypts data using the communication secret key included in the **project access key** to request a new session security key.
5. The server generates the requested session security key and delivers it to the agent.
6. Subsequently, the agent encrypts data such as text and numbers before transmitting them to the server.

! To apply the changes, restart the server.

Agent resource limits

The WhaTap server monitoring agent for Linux and Windows servers provides the following agent settings to help maintain server stability and minimize the resource usage impact when any overload occurs.

Linux / Windows

Memory usage protection

The WhaTap agent automatically restarts the process when the memory usage exceeds the threshold (100 MB) to protect available memory on your server.

```
memory.usage.limit=104857600 # Memory usage threshold (100 MB, unit of byte)
```

Automatic restart

The WhaTap agent automatically restarts once a day. The cycle can be adjusted via user settings. This feature is provided to address data collection failures and temporary process malfunctions.

```
housekeeping.enabled=true      # Activation of automatic restart (Default: true)
housekeeping.hour=2           # Restart time (based on the server time, Default: 2 O'clock)
housekeeping.exitoncheckfail=true # Automatic restart upon collection failure (Default: true)
```

The restart time can also be specified in minutes, not only in hours. When you also set the `housekeeping.minute` option, the agent restarts after the time pointed to by `housekeeping.hour` and `housekeeping.minute`. In an environment with many servers, you can assign a different minute value per server so that the restart times do not pile up on the hour.

```
housekeeping.hour=2           # Restart time (based on the server time, Default: 2 O'clock)
housekeeping.minute=30        # Restart minute (Default: 0)
```

| Item | Description |
|----------------------------------|---|
| <code>housekeeping.hour</code> | Hour of the restart time. Default <code>2</code> |
| <code>housekeeping.minute</code> | Minute of the restart time. Default <code>0</code> . With <code>0</code> , the agent restarts within the specified hour as before |

The supported agent versions are Linux and Windows 2.9.22 or later, and Unix (AIX, HP-UX, Solaris) 1.4.13 or later.

CPU usage protection

When the CPU usage exceeds the threshold, the WhaTap agent automatically adjusts the execution speed to protect the CPU availability of the server. It is disabled by default and can be enabled if necessary.

```
cpu.limit.enabled=false      # Disable the CPU usage protection (Default: false)
cpu.usage.limit=50          # CPU usage threshold (50% based on the single core)
```

Windows

Handle usage protection

In the Windows environment, if the number of handles used by the WhaTap agent exceeds the threshold (1,000), it is automatically restarted to protect the available handles of the user server.

```
windows.handle.threshold=1000 # Handle usage threshold (Default: 1,000)
```

Watchdog communication method

For an agent in a Windows environment, the monitoring process (watchdog) and the collection process exchange their status with each other. This communication uses a named pipe by default, and you can change it to a local TCP socket in environments where pipe writes fail and the communication is disconnected.

```
housekeeping.rpc.mode=socket # Watchdog communication method (Default: pipe)
housekeeping.rpc.port=6601 # Port used in the socket method (Default: 6601)
```

| Item | Description |
|-----------------------|---|
| housekeeping.rpc.mode | Watchdog communication method. With <code>socket</code> , a local TCP socket is used. The default is a named pipe |
| housekeeping.rpc.port | TCP port used in the <code>socket</code> method. Default <code>6601</code> |

The `socket` method is used only for internal (local) communication on the server, so you do not need to open an external firewall. However, make sure that the specified port is not in use by another program.

Configuring The Windows Performance Counter

The Windows performance count can be arbitrarily collected. Add the performance counter set name to the WhaTap server monitoring configuration file.

- The performance counter set name can be retrieved.

Windows Powershell

```
Get-Counter -ListSet *
```

- A performance counter set can be added in the configuration file.

Windows Powershell

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "performance.counter.$(Get-Date -uformat %s)=Performance counter set name to view"
```

- In **Analysis > Metrics Search**, the added performance counter can be searched by **Category**
`performance_counter_{performance counter set name}` .

! To apply the changes, restart the server.

Advanced Functions

Expansion pack settings

To use the extension pack function after installation, create the configuration file and then restart the agent as follows.

MySQL instances

The MySQL/MariaDB monitoring agent is available.

Checking the web URL performance

The WhaTap infrastructure agent accesses the given web address and then collects the access status and response time. The collected web URL performance can be checked in **Category** `server_webcheck`. Add web addresses to check with the following command. For more information about URL monitoring, see [the following](#).

- [Linux Shell](#) Windows Powershell

```
echo "web address to check" | sudo tee -a /usr/whatap/infra/conf/webcheck.conf
echo "webcheck.enabled=true" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
sudo service whatap-infra restart
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\webcheck.conf" -Value "web address to check"
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "webcheck.enabled=true"
Restart-Service "Whatap Infra"
```

Redis Instances

The WhaTap infrastructure agent accesses the given Redis address and collects multiple performance metrics including the number of connections. The collected performance data can be checked in **Category** `server_redis`. You can add the connection data to the Redis server to be checked with the following command.

- [Linux Shell](#) Windows Powershell

```
sudo echo "addr=Redis address:port" | sudo tee /usr/whatap/infra/conf/redis.conf
sudo echo "password=Redis password" | sudo tee -a /usr/whatap/infra/conf/redis.conf
sudo service whatap-infra restart
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\redis.conf" -Value "addr=Redis address:port"
Add-Content -Path "C:\Program Files\WhatapInfra\redis.conf" -Value "password=Redis password"
Restart-Service "Whatap Infra"
```

Memcached Server List

The WhaTap infrastructure agent accesses the given Memcached address and collects multiple performance metrics including the number of connections. The collected performance data can be checked in **Category** `server_memcached`. You can add the connection data to the Memcached server to be checked with the following command.

- [Linux Shell](#) [Windows Powershell](#)

```
sudo echo "addr=Memcached address:port" | sudo tee /usr/whatap/infra/conf/memcached.conf
sudo echo "password=Memcached password" | sudo tee -a /usr/whatap/infra/conf/memcached.conf
sudo service whatap-infra restart
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\memcached.conf" -Value "addr=Memcached address"
Add-Content -Path "C:\Program Files\WhatapInfra\memcached.conf" -Value "password=Memcached password"
Restart-Service "Whatap Infra"
```

TCP port monitoring

You can be notified when the network daemon's listening port has been disabled.

Select Home > Select Project > **Alert** > **Event Configuration** > **Server** tab, and then select **+ Add alert settings**.

The screenshot shows the 'Event Configuration' window with the 'Add alert settings' sub-window open. The 'Title' field is empty. Under 'Choose event type', 'Server' is selected. Under 'Event rule condition', a table lists various metrics and their configurations. The 'Port' metric is highlighted with a yellow box and a red circle with the number 1. The 'Save' button is highlighted with a red circle with the number 2.

| Indicators | Condition | Level | Duration | Enable |
|--------------|--|------------------|----------------------|--|
| Restart | - | Fatal | - | ☒ |
| Not received | - | Warning | 10 Minu | ☐ |
| Port | - | Warning | 1 Minute | ☒ |
| Network IOPS | Warning > 45,000 pps
Fatal > 50,000 pps | Warning
Fatal | 5 Minute
5 Minute | ☒
☒ |
| Network BPS | Warning > 150,000,000 bps
Fatal > 200,000,000 bps | Warning
Fatal | 5 Minute
5 Minute | ☒
☒ |
| Disk I/O | 90% 95% | Warning
Fatal | 5 Minute
5 Minute | ☒
☒ |
| Disk quota | 90% 95% | Warning
Fatal | 5 Minute
5 Minute | ☒
☒ |

Server list

- ☐ Select all
- ☐ demo-k8s-master (10.21.1.149)
- ☐ demo-k8s-worker-01 (10.21.1.229)
- ☐ demo-k8s-worker-02 (10.21.1.157)
- ☐ demo-k8s-worker-03 (10.21.1.23)
- ☐ demo-k8s-worker-04 (10.21.1.210)
- ☐ dotnet-official-demo (10.21.1.51)
- ☐ go-official-demo (10.21.1.118)
- ☐ mysql-official-demo-01 (10.21.1.86)
- ☐ mysql-official-demo-02 (10.21.1.136)
- ☐ mysql-official-demo-02 (10.21.1.136)
- ☐ mysql-official-demo-03 (10.21.1.236)
- ☐ node-official-demo (10.21.1.98)
- ☐ python-demo-attack (10.21.1.254)
- ☐ python-demo-target (10.21.1.43)
- ☐ rum-official-demo (10.21.1.151)

Save

- Click the **1** Port metric toggle button and then change it to **Use**.
- Select **2** Save to save the settings.
- Add the following command on the target server.
- **Linux Shell** **Windows Powershell**

```
#Specify the IP to be monitored in the following variable.
export TARGET_IP=127.0.0.1
#Specify the port to be monitored in the following variable.
export TARGET_PORT=80
```



```
echo "tcp.check.$TARGET_PORT=tcp://$TARGET_IP:$TARGET_PORT" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

#Specify the IP to be monitored in the following variable.

```
$TARGET_IP="127.0.0.1"
```

#Specify the port to be monitored in the following variable.

```
$TARGET_PORT="80"
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value
```

```
"tcp.check.${TARGET_PORT}=tcp://${TARGET_IP}:${TARGET_PORT}"
```

❗ Multi-port monitoring

To monitor multiple ports, additional configuration is required in the [whatap.conf](#) file. See the following example:

whatap.conf

```
# port 8080, port 80
```

```
tcp.check.8080=tcp://127.0.0.1:8080
```

```
tcp.check.80=tcp://127.0.0.1:80
```

Collecting process file descriptor information

You can view process file descriptor information in the **Process** widget on the **Server List > Server Detail** screen.

- [Linux Shell](#) Windows Powershell

```
echo "process.fdcheck=true" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

```
# The admin role is required.
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "process.fdcheck=true"
```

Collecting the server specifications

It collects the specifications for CPU, memory, disk, and NIC, installed packages, and meta data of virtual

machines. For each **Category**, you can check the following.

- `server_package` : Installed package.
- `server_inventory` : Specifications for CPU, memory, disk, and NIC.
- `server_pending_security_update` : Waiting list for Linux security patches.
- **Linux Shell** **Windows Powershell**

```
echo "inventory.product.enabled=true" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

```
# The admin role is required.
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "inventory.product.enabled=true"
```

Collecting the Linux container performance

It collects the performance for CPU, memory, disk, and NIC in the Linux containers. The collected performance data can be checked in **Category** `server_container` .

Linux Shell

```
echo "container.enabled=true" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

Collecting the Windows server service operations status

It collects the Windows service operation status. The collected service operation status can be checked in **Category** `server_windows_service` .

Windows Powershell

```
# The admin role is required.
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "winservice.enabled=true"
```

Excluding file systems from collection

You can specify file systems that do not need to be collected with a regular expression and exclude them.

Excluding file systems that always report a usage close to 100%, such as snapshot mounts, reduces unnecessary alerts and data.

Enter a regular expression in the `filesystem.exclude.pattern` option. File systems whose mount point or device name matches the regular expression are excluded from both the `server_disk` and `infra_filesystem` categories.

whatap.conf

```
filesystem.exclude.pattern=^/snap/
```

| Item | Description |
|---|--|
| <code>filesystem.exclude.pattern</code> | Regular expression for the mount point or device name of the file system to exclude from collection. The default is empty, and an empty value means no exclusion |

❗ If the regular expression is invalid, the agent logs an error and does not apply the option. File system collection itself keeps working.

The supported agent versions are Linux and Windows 2.9.23 or later, and Unix (AIX, HP-UX, Solaris) 1.4.14 or later. Restart the agent after changing the settings.

File monitoring

You can monitor the files and directories that exist or can exist on the file system.

- [Linux Shell](#) [Windows Powershell](#)

```
echo "filecheck.enabled=true" |sudo tee -a /usr/whatap/infra/conf/whatap.conf
echo "Full path of the file or directory to monitor" |sudo tee -a /usr/whatap/infra/conf/filesystem.conf
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "filecheck.enabled=true"
Add-Content -Path "C:\Program Files\WhatapInfra\filesystem.conf" -Value "Full path of the file or directory to monitor"
```

Date substitution in the monitored path

If you put date format characters in the monitored path, the agent substitutes them with the date at collection time and monitors that path. You can monitor targets whose path changes every day, such as logs whose file name changes by date, with a single setting. It works the same way as the log monitoring feature.

filesystem.conf

```
/var/log/app/access_%Y%m%d.log
```

| Item | Description |
|------|------------------|
| %Y | Year (4 digits) |
| %m | Month (2 digits) |
| %d | Day (2 digits) |
| %H | Hour (2 digits) |

When the date changes and the monitored path changes, monitoring starts anew without comparing with the state of the previous path.

The supported agent versions are Linux and Windows 2.9.23 or later, and Unix (AIX, HP-UX, Solaris) 1.4.14 or later.

WhaTap Agent Log Options

Log Location

Change the path of the WhaTap agent logs (`whatap_infrad.log` , `whatap_infra.log` , `whatap_infrad.fork.stdout.log`). If the specified path does not exist, the WhaTap agent attempts to create the path. It is available in the agent 2.0.6 or later.

- **Linux Shell** Windows Powershell

```
echo "agent.log.prefix=Desired Directory" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "agent.log.prefix=Directory to change"
```

ⓘ In case of Windows, enter "\" as the path delimiter.

Log Output Control Options

The server monitoring agent logs are stored under the `/var/log` directory by default. You can adjust the log file size limit and the number of backup files using the following options.

- **Linux / Windows** Unix (AIX, HPUX, Solaris)

- Linux / Windows
 - agent.log.size
 - Default: 10 (unit: MB)
 - agent.log.backupfile
 - Default: 2 (unit: count)
- Unix (AIX, HPUX, Solaris)
 - log.size
 - Default: 4 (unit: MB)
 - log.backup
 - Default: 3 (unit: count)
- AIX additional log options (errpt collection log)
 - aix.errpt.log.size=4
 - Default: 4 (unit: MB)
 - aix.errpt.log.backup=3

- Default: 3 (unit: count)

Netstat monitoring

If the netstat command is embedded, collection is uploaded to **Category** `server_netstat` (netstat connection statistics) and `server_netstat_s` (result of netstat -s command).

- [Linux Shell](#) [Windows Powershell](#)

```
echo "netstats.enabled=true" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

```
Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "netstats.enabled=true"
```

User script monitoring

- [Linux Shell](#) [Windows Powershell](#)

It periodically runs the executable user script on the Linux server and uploads the results. The user script monitoring feature runs the script at set intervals and uploads the results to the collection server. It collects the results by running the script at 30-second intervals.

Data is collected based on the script completion time, and there may be some difference between the actual collection time and the data reflection time depending on the content of the running script. Write your own script noting the following example.

1. Write the script as follows and then run it.

Linux Shell

```
#!/bin/bash

urls=("https://www.whatap.io" "https://blog.whatap.io" "https://service.whatap.io")

for url in ${urls[*]}; do
```

```

echo "M $url url $url"
curl -w @- -o /dev/null -s "$url" <<'EOF'
M %{url_effective} url %{url_effective}\n
H %{url_effective} time_namelookup %{time_namelookup}\n
H %{url_effective} time_connect %{time_connect}\n
H %{url_effective} time_appconnect %{time_appconnect}\n
H %{url_effective} time_pretransfer %{time_pretransfer}\n
H %{url_effective} time_redirect %{time_redirect}\n
H %{url_effective} time_starttransfer %{time_starttransfer}\n
H %{url_effective} time_total %{time_total}\n
EOF

done

```

2. Check whether the out is as follows:

Linux Shell

```

$ ./curl.sh
M https://www.whatap.io/ url https://www.whatap.io/
H https://www.whatap.io/ time_namelookup 0.007289
H https://www.whatap.io/ time_connect 0.010702
H https://www.whatap.io/ time_appconnect 0.033637
H https://www.whatap.io/ time_pretransfer 0.033831
H https://www.whatap.io/ time_redirect 0.000000
H https://www.whatap.io/ time_starttransfer 0.043230
H https://www.whatap.io/ time_total 0.043312
M https://blog.whatap.io/ url https://blog.whatap.io/
H https://blog.whatap.io/ time_namelookup 0.000000
H https://blog.whatap.io/ time_connect 0.000000
H https://blog.whatap.io/ time_appconnect 0.000000
H https://blog.whatap.io/ time_pretransfer 0.000000
H https://blog.whatap.io/ time_redirect 0.000000
H https://blog.whatap.io/ time_starttransfer 0.000000
H https://blog.whatap.io/ time_total 0.005654
M https://service.whatap.io/ url https://service.whatap.io/
H https://service.whatap.io/ time_namelookup 0.004296
H https://service.whatap.io/ time_connect 0.007282

```

```
H https://service.whatap.io/ time_appconnect 0.038404
H https://service.whatap.io/ time_pretransfer 0.038607
H https://service.whatap.io/ time_redirect 0.000000
H https://service.whatap.io/ time_starttransfer 0.055226
H https://service.whatap.io/ time_total 0.055307
```

3. Check whether the collection is normal and register the script to the WhaTap server monitoring agent as follows:

Linux Shell

```
export USER_SCRIPT=Absolute path to the script to run
suco chmod +x $USER_SCRIPT
sudo mkdir /usr/whatap/infra/ext/
sudo cp -r $USER_SCRIPT /usr/whatap/infra/ext/
export RUNAS_USER=Account to run the script. Root account is excluded.
sudo WHATAP_HOME=/usr/whatap/infra/conf /usr/whatap/infra/whatap_infrad --user=$RUNAS_USER init-script
sudo service whatap-infra restart
```

The collected data can be viewed in **Category** 'server_ext_folder_runscript' under **Analysis > Metrics Chart**.

It periodically runs the executable user script on the Windows server and uploads the results. The user script monitoring feature runs the script at set intervals and uploads the results to the collection server. It collects the results by running the script at 30-second intervals.

Data is collected based on the script completion time, and there may be some difference between the actual collection time and the data reflection time depending on the content of the running script. Write your own script noting the following example.

- **BAT** PowerShell

Running the BAT script

1. Write the script as follows and then run it.

BAT


```

@echo off
setlocal enabledelayedexpansion

set urls=https://www.whatap.io https://blog.whatap.io https://service.whatap.io

for %%u in (%urls%) do (
    echo M %%u url %%u

    rem Use curl to fetch the URL and get the total time taken for the request
    for /f "tokens=*" %%i in ('curl -o nul -s -w "M %{url_effective} url %{url_effective}" %%u') do (
        echo %%i
    )

    rem Measure the start time
    set startTime=!time!

    rem Perform the curl request and measure the end time
    curl -o nul -s %%u

    set endTime=!time!

    rem Parse the start time
    for /f "tokens=1-4 delims=.,," %%a in ("!startTime!") do (
        set /A "startH=%%a, startM=%%b, startS=%%c, startMS=%%d"
    )

    rem Parse the end time
    for /f "tokens=1-4 delims=.,," %%a in ("!endTime!") do (
        set /A "endH=%%a, endM=%%b, endS=%%c, endMS=%%d"
    )

    rem Calculate the elapsed time in milliseconds
    set /A "elapsedMS=(endH*3600000 + endM*60000 + endS*1000 + endMS) - (startH*3600000 + startM*60000 + startS*1000 + startMS)"

    rem Convert milliseconds to seconds and milliseconds part
    set /A "elapsedSec=elapsedMS/1000"
    set /A "elapsedMSPart=elapsedMS %% 1000"

```

```

    echo H %%u time_total !elapsedSec!.!elapsedMSPart!
)

endlocal

```

2. Check whether the out is as follows:

BAT

```

>webcheck.bat
M https://www.whatap.io url https://www.whatap.io
M https://www.whatap.io/ url https://www.whatap.io/
H https://www.whatap.io time_total 0.38
M https://blog.whatap.io url https://blog.whatap.io
M https://blog.whatap.io/ url https://blog.whatap.io/
H https://blog.whatap.io time_total 0.1
M https://service.whatap.io url https://service.whatap.io
M https://service.whatap.io/ url https://service.whatap.io/
H https://service.whatap.io time_total 0.922

```

3. Check whether the collection is normal and register the script to the WhaTap server monitoring agent as follows:

BAT

```

cd "\Program Files\WhatapInfra"
whatap_infra.exe init-script
sc stop "Whatap Infra"
sc start "Whatap Infra"

```

Running the PowerShell script

1. Write the script as follows and then run it.

PowerShell

```

$urls = @("https://www.whatap.io", "https://blog.whatap.io", "https://service.whatap.io")

```

```
# Hide progress bar
$ProgressPreference = 'SilentlyContinue'

foreach ($url in $urls) {
    Write-Output "M $url url $url"

    $startTime = Get-Date
    try {
        $response = Invoke-WebRequest -Uri $url -UseBasicParsing
        $endTime = Get-Date

        $timeTotal = ($endTime - $startTime).TotalSeconds

        Write-Output "M $($response.BaseResponse.ResponseUri) url $($response.BaseResponse.ResponseUri)"
        Write-Output ("H {0} time_total {1}" -f $url, $timeTotal)
    } catch {
        Write-Output "Failed to access $url"
    }
}
```

2. Check whether the out is as follows:

PowerShell

```
> .\webcheck.ps1
M https://www.whatap.io url https://www.whatap.io
M url
H https://www.whatap.io time_total 0.376979
M https://blog.whatap.io url https://blog.whatap.io
Failed to access https://blog.whatap.io
M https://service.whatap.io url https://service.whatap.io
Failed to access https://service.whatap.io
```

3. Check whether the collection is normal and register the script to the WhaTap server monitoring agent as follows:

PowerShell

```
cd "\Program Files\WhatapInfra"
.\whatap_infra.exe init-script
Restart-Service "Whatap Infra"
```

Server time synchronization monitoring

ⓘ Preliminary checking

- Operating system: Linux
- WhaTap Server Monitoring Agent version: 2.5.4 or later

You can enable the server time synchronization (NTP) monitoring through the following option:

```
echo "ntp.enabled=true" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

Guide to additional key metrics

The metrics to be added when activating the server time synchronization monitoring are as follows:

server_ntpq

Tags

| Tag name | Description | Remarks |
|----------|---|---------|
| refid | Time synchronization source of the corresponding remote (remote). | - |
| remote | Server to which the host performs time synchronization. | - |

Fields

| Field name | Unit | Description | Remarks |
|-------------|--------|--|---------|
| delay | ms | Time taken for the server to return a reply packet in response to a query. | - |
| jitter | - | Error threshold estimated as the difference in offset measurements between two samples. | - |
| offset | ms | Difference between the server setting time and client setting time. | - |
| poll | second | Polling interval to the remote host defined by the minpoll value in the ntp.conf file. | - |
| reach | - | You can see how successful the attempts to reach the server were. It records the success or failure of the most recent query at the position 2 ⁰ and operates as an 8-bit shift register. | - |
| st | - | Stratum level of the remote host. | - |
| sync_client | - | 1: The synchronization server exists. | - |
| | | 0: The synchronization server does not exist. | |
| t | - | Available types are as follows. | - |
| | | l: Local (e.g. GPS clock) | |
| | | u: Unicast (general type) | |
| | | m: Multicast | |
| | | b: Broadcast | |
| when | second | Seconds elapsed after the remote host responded. | - |

Linux system log monitoring

By uploading the Linux system logs in real time, you can check the logs generated upon failure. The log files to collect are as follows:

- `/var/log/syslog`
- `/var/log/messages`

The configuration is as follows:

Linux Shell

```
echo "systemlog.enabled=true" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

Log monitoring plug-in

You can arbitrarily collect the logs. The following configuration requires restart of the agent.

Linux Shell

- [Basic installation](#)

Linux Shell

```
export LOGFILES=/some/path/file1,/some/other/file2
echo "logsink.files=$LOGFILES" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

- [Install with Category](#)

Linux Shell

```
1 cd /usr/whatap/infra
2 sudo mkdir extension
3
4 cat <<EOL | sudo tee extension/logsink.conf >> /dev/null
```

```

5 [[inputs.logsink]]
6   category = "serverlog"
7   ## Whether to transfer the log generation statistics as separate data
8   stats_enabled = true
9   ## Log generation statistics category
10  stats_category = "logsink_stats"
11  ## Set the log file name to be excluded if the log file path contains asterisk (*).
12  excludeNames = [ ".gz", ".zip" ]
13  [[inputs.logsink.file]]
14    ## When setting a log file, you can set the date pattern (strftime.org).
15    path = "/some/path/%Y-%m-%d/.log"
16    disabled = false
17    encoding = "utf-8"
18
19  [[inputs.logsink.file]]
20    path = "/some/other/log"
21    disabled = false
22    encoding = "utf-8"
23
24  [[inputs.logsink.file]]
25    ## When searching for the keyword in the line-by-line log, merge it with the previous log.
26    nowrap_keywords = ["Caused by:", "Test"]
27
28  [[inputs.logsink.file]]
29    ## When entering the project code and project license, the logs can be transferred to a desiged project in the
    same region.
30    # pcode = Project code
31    # license = "Project license"
32
33  [[inputs.logsink.file]]
34    # whatap_host_ip = "Collection server IP1/Collection server IP2"
35    # whatap_host_port = 6600
36
37  [[inputs.logsink.file]]
38    path = "/tmp/mixed.log"
39    disabled = false
40    encoding = "utf-8"
41    ## The [{"A":1,"B":2},{"C":3,"D":4},...] format can be divided into individual logs.{"A":1,"B":2} {"C":3,"D":4}
42    multiline_pattern = "simple_json_array"

```

```

43 EOL
44
45 sudo service whatap-infra restart

```

Windows Powershell

- Basic installation

Windows Powershell

```

$LOGFILES="c:\\whatap\\logs\\%Y-%m-%d\\*.log,c:\\whatap\\logs\\*.log"
Add-Content "c:\\Program Files\\WhatapInfra\\whatap.conf" -Value "logsink.files=$LOGFILES"

```

- Install with Category

Windows Powershell

```

1  # The administrator role required
2  New-Item -type "Directory" -Path "C:\\Program Files\\WhatapInfra\\extension"
3
4  $contentToAdd = @"
5  [[inputs.logsink]]
6  category = "serverlog"
7  ## Whether to transfer the log generation statistics as separate data
8  stats_enabled = true
9  ## Log generation statistics category
10 stats_category = "logsink_stats"
11 ## Set the log file name to be excluded when the log file path contains asterisk (*).
12 excludeNames = [ ".gz", ".zip" ]
13 [[inputs.logsink.file]]
14 ## The date pattern (strftime.org) can be entered when setting the log file.
15 path = "c:\\whatap\\logs%Y-%m-%d_.log"
16 disabled = false
17 encoding = "euc-kr"
18
19 [[inputs.logsink.file]]
20 ## When searching for the keyword in the line-by-line logs, it merges with the previous log.
21 nowrap_keywords = ["Caused by:", "Test"]

```



```

22
23 [[inputs.logsink.file]]
24 ## When entering the project code and project license, logs can be transferred to a desired project in the same
    region.
25 # pcode = Project code
26 # license = "Project license"
27
28 [[inputs.logsink.file]]
29 # whatap_host_ip = "Collection server IP1/Collection server IP2"
30 # whatap_host_port = 6600
31
32 [[inputs.logsink.file]]
33 path = "/tmp/mixed.log"
34 disabled = false
35 encoding = "utf-8"
36 ## The [{"A":1,"B":2},{ "C":3,"D":4},...] format can be divided into individual logs.{"A":1,"B":2} {"C":3,"D":4}
37 multiline_pattern = "simple_json_array"
38
39 "@
40
41 New-Item -path "C:\Program Files\WhatapInfra\extension" -name "logsink.conf" -type "file" -value $contentToAdd
    -Force
42
43 Restart-Service "Whatap Infra"

```

Configuring options

- `stats_enabled` : Set whether or not to collect the collection status data. The default value is `false` . The value must be set to `true` . If set to `true` , statistical data is generated in the category (`logsink_stats`) set in the following `stats_category` .
- `stats_category` : Set the metrics category to store the collection status data in. The value must be set to `logsink_stats` . Statistical data fields are as follows:
 - `file`
 - `checkInterval`
 - `encoding`

- filepos
- checkedLocalTime
- lastupdatedLocalTime
- fileSize
- error
- firstCheck
- transferBytes
- excludeNames : By including asterisks (*) in the log file path, you can set the file name to exclude non log files. You can set multiple values using comma (,) as a separator.

Example

```
excludeNames = [ ".gz", ".zip" ]
```

- nowrap_keywords : While searching logs by line, if any keyword is detected by a specified option value, it is merged with the previous log.
- pcode : Enter the project code of the same region to which you want to send the logs.
- license : Enter the license code of the project in the same region to which you want to send the logs.

ⓘ In case of the pcode and license options, it is supported from Server Agent 2.7.4 or later.

Configuring the log options for Windows events

When collecting Windows event logs, you can set the options as follows:

```
1 # The administrator role required
2 New-Item -type "Directory" -Path "C:\Program Files\WhatapInfra\extension"
3
4 $contentToAdd = @"
5 [[inputs.win_eventlog]]
6     category = "win_event_log"
7     stats_category = "win_event_log_stats"
```

```

8  stats_enabled = true
9  enabled = true
10 [[inputs.win_eventlog.file]]
11   #true | false
12   enabled = true
13   # Application, Security, Setup, System, ForwardedEvents
14   file = "Application"
15   #1: Information, 2: Warning 3: Error 4: Audit Success 5 Audit Fail
16   #event_type =
17   #event_id
18   #event_id =
19   #event source name
20   #source_name = ""
21   #When entering the project code and project license, logs can be transferred to the desired project in the same region.
22   #pcode = Project code
23   #license = "Project license"
24   #whatap_host_ip = "Collection server IP1/Collection server IP2"
25   #whatap_host_port = 6600
26
27 "@
28
29 New-Item -path "C:\Program Files\WhatapInfra\extension" -name "win_eventlog.conf" -type "file" -value $contentToAdd
30 -Force
31 Restart-Service "Whatap Infra"

```

- It is required to specify the `category` .
e.g. `win_event_log`
- It is required to specify the statistics category (`stats_category`).
e.g. `win_event_log_stats`
- The statistics category On/Off (`stats_enabled`) is required.
e.g. `true` or `false`
- The collection On/Off (`enabled`) is required.
e.g. `true` or `false`

- The file-based collection On/Off (`enabled`) is required.

e.g. `true` or `false`

- The file (`file`) is required.

e.g. `Application` , `Security` , `Setup` , `System` , `Forwarded`

- The event type (`event_type`) is optional.

e.g. `1` , `2` , `3` , `4` , `5`

! Event type

1. Information
2. Warning
3. Error
4. Audit Success
5. Audit Fail

- The event ID (`event_id`) is optional.
- The event source name (`source_name`) is optional.
- Project code (`pcode`) optional
- Project license (`license`) optional

- ! • For Windows event log collection, the Server Agent 2.5.2 or later is supported.
- Among the Windows event log options, `pcode` and `license` are supported in the Server Agent 2.7.4 or later.

Specifying and separating multiple log files

Specifying multiple files

To categorize multiple log files under the same category, add each log file path into the `[[inputs.logsink.file]]` entry and specify the file path by using the `path` setting. Additionally, you can configure additional settings, such as log generation statistics and options such as `category`, `stats_enabled`, and `stats_category`. The example is where the `/root/test1` and `/root/test2` files are assigned the same category (`serverlog`).

```
[[inputs.logsink]]
  category = "serverlog"
  stats_enabled = true
  stats_category = "logsink_stats"
  excludeNames = [ ".gz", ".zip" ]
```

```
[[inputs.logsink.file]]
  path = "/root/test1"
  disabled = false
  encoding = "utf-8"
```

```
[[inputs.logsink.file]]
  path = "/root/test2"
  disabled = false
  encoding = "utf-8"
```

❗ For information about options such as `category`, `stats_enabled`, `stats_category`, and `nowrap_keywords`, see the guidance [above](#).

In this example, you can see the following in **Live Tail**.

| serverlog | | | |
|-----------|--------|-------------------------|---|
| | | | |
| ▼ | oname | | |
| ▼ | APPSVR | 2024-10-11 18:14:30.325 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709
kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh |
| ▼ | APPSVR | 2024-10-11 18:14:30.325 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709
kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh |
| ▼ | APPSVR | 2024-10-11 18:14:30.325 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709
kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh |
| ▼ | APPSVR | 2024-10-11 18:14:30.325 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709
kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh |
| ▼ | APPSVR | 2024-10-11 18:14:30.325 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709
kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh |
| ▼ | APPSVR | 2024-10-11 18:14:30.325 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709
kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh |
| ▼ | APPSVR | 2024-10-11 18:14:40.326 | file /root/test2 pcode 34 oname APPSVR agenttime 1728638079521 oid 1988910709
sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbvcmbn.bcmn |

File separation

To specify different categories for log files, create separate `logsink{number}.conf` files and then set different `category` values in each file. The following example shows how to set different log categories through `logsink.conf` and `logsink1.conf`.

logsink.conf

```
# cat logsink.conf
[[inputs.logsink]]
  category = "serverlog1"
  stats_enabled = true
  stats_category = "logsink_stats"
  excludeNames = [ ".gz", ".zip" ]

[[inputs.logsink.file]]
  path = "/root/test1"
  disabled = false
  encoding = "utf-8"
```

logsink1.conf

```
#cat logsink1.conf
[[inputs.logsink]]
  category = "serverlog2"
  stats_enabled = true
  stats_category = "logsink_stats"
  excludeNames = [ ".gz", ".zip" ]

[[inputs.logsink.file]]
  path = "/root/test2"
  disabled = false
  encoding = "utf-8"
```

❗ For information about options such as `category` , `stats_enabled` , `stats_category` , and `nowrap_keywords` , see the guidance [above](#).

In this example, you can see the following in **Live Tail**.

| FILTER | | VIEW | |
|----------|--------|-------------------------|--|
| All logs | | SHOW ALL LOGS | |
| ▼ | oname | TIME | LOG |
| ▼ | APPSVR | 2024-10-11 18:20:55.084 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638454313 oid 1988910709 category serverlog1
kdjfhakdhgldkjg ha,dfghad;fkjgha;dfkjgha;dfgjha;fkgjha;fgjh |
| ▼ | APPSVR | 2024-10-11 18:20:55.087 | file /root/test1 pcode 34 oname APPSVR agenttime 1728638454313 oid 1988910709 category serverlog1
kdjfhakdhgldkjg ha,dfghad;fkjgha;dfkjgha;dfgjha;fkgjha;fgjh |
| ▼ | APPSVR | 2024-10-11 18:21:05.103 | file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2
sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn |
| ▼ | APPSVR | 2024-10-11 18:21:05.104 | file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2
sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn |
| ▼ | APPSVR | 2024-10-11 18:21:05.104 | file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2
sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn |
| ▼ | APPSVR | 2024-10-11 18:21:05.104 | file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2
sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn |
| ▼ | APPSVR | 2024-10-11 18:21:05.104 | file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2
sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn |

GPU Feature

This guide provides information about the agent environment and GPU options required to enable GPU monitoring.

GPU Monitoring Support Environment

By default, performance and inventory data can be collected using `nvidia-smi`. However, if `nvidia-dcgm` is installed, additional metrics can be collected.

When `nvidia-dcgm` is present, the `nvidia-dcgm` service must be running.

- You can verify that the `nvidia-dcgm` service is running properly with the following command:
`systemctl status nvidia-dcgm`
- The list of collectable metrics may continue to grow over time.

For details on metrics by collection tool, see [GPU Metrics](#).

GPU Features

To use the GPU Inventory and GPU Performance Summary features, you must update the agent to the required version and apply the necessary settings.

✔ Supported Version

Linux minimum 2.9.3, recommended 2.9.16 or later

- Amazon Linux, RHEL, Fedora, Oracle Linux, CentOS, Rocky Linux, SUSE

```
sudo yum update whatap-infra
```

- Debian, Ubuntu


```
sudo apt-get update
sudo apt-get install whatap-infra
```

GPU Options

To enable the GPU Inventory feature, the following options must be enabled:

- **nvidiaiav2.enabled** Boolean

Default `false`

Option to collect NVIDIA GPU metrics including physical GPU configuration and resource usage.

- **nvidiaiav2.mig.enabled** Boolean

Default `false`

Option to collect metrics related to NVIDIA GPU MIG configuration and resource usage. Requires **nvidiaiav2.enabled** to be set to `true`.

- **nvidiaiav2.proc.enabled** Boolean

Default `false`

Option to collect information about processes assigned to NVIDIA GPUs. Requires **nvidiaiav2.enabled** to be set to `true`.

- **nvidiaiav2.dcp.enabled** Boolean

Default `true`

Option to collect NVIDIA DCP (Data Center Profiling) metrics. Requires **nvidiaiav2.enabled** to be set to `true`.

- **nvidiaiav2.interval.sec**


Default `30 seconds`

Interval for collecting NVIDIA GPU metrics.

Resource Board

Home > Select Project > **Dashboard** > **Resource Board**

On the initial screen of the WhaTap Monitoring service, select a project and then select **Dashboard** > **Resource Board**.



Resource Board allows you to monitor all servers registered in a project.

It provides the [CPU resource map](#) to check the summary of all servers in the project and changes in real-time resource usage. It displays [Top 5](#) total resource size, CPU, memory, and utilization processes. You can immediately recognize and respond to failures through the **Resource Board**.

1 Fixed menu

-  **Theme:** You can change the theme of the current page. The default value is [Light](#) mode.
-  **Alert Status:** It displays recent events. You can turn on or off the browser notifications by clicking the toggle menu.
-  **Customer Support:** Online customer support is available.
- **Profile icon:** Select the rightmost icon. **Manage Accounts**, **Create organization**, **Organization Management**, and function to select a language are provided.

2 Info panel

It provides the summary metrics for all servers registered in the project. They are updated every 10 seconds.

- **Server**

It provides the total number of servers together with the number of servers with problems. If you select an area, the **Server List** page appears.

- Exception status: [Danger](#), [Warning](#)
- Normal Server: [Medium](#), [Inactive](#)

- **OS**

Number of operating systems in the project such as Linux and Windows

- **Total Cores**

Sum of cores for all servers.

- **Avg CPU**

Average CPU usage (%) for all servers. The gauges below are Average, Maximum, and Minimum.

- **Avg Memory**

Average memory usage (%) for all servers. The gauges below are Average, Maximum, and Minimum.

- **Avg Disk**


Average disk usage (%) for all servers. The gauges below are Average, Maximum, and Minimum.

3 CPU Resource Map

This distribution chart displays the CPU usage of all servers. It displays data for 10 minutes and is updated every 10 seconds. The deepness of the cell color indicates that the server's resource usage is widely distributed in the area.

- **Provision of detailed info**

If you drag the cell area, the detailed information is provided.



- The server list is sorted by the maximum value of the data collected in the range.
- Five servers are selected by default, and the usage trend of the selected servers are also provided in a multi-line chart.
- Click **Refresh** to refresh the current data.

! Multi Line Chart

Multiple lines are displayed on a chart.

4 Resource Equalizer

It displays a list of the top 5 servers in real time for **CPU**, **Memory**, **Disk I/O**, and **Disk IOPS**.

For example, the disk I/O (%) metric displays the disk usage for 5 seconds. If the disk I/O usage (%) exceeds 80%, the system performance can be affected. The default alert value is 90%. If the disk I/O (%) is 100%, it means that the disk is working non-stop.

- If you select the equalizer chart area, the target **Server Detail** page appears.
- If you select > , the **Resource Equalizer** page appears. You can see the real-time usage for all servers.

5 Process Top 5

The list of top 5 processes is provided based on the maximum occupancy of CPU and memory. If you select > , the **Process list** for all servers appears.

| Process list | | | | | | | |
|--------------------|-----------------------|------------------------|-------|-------|--------|----|--|
| Process Name, User | | | | | | | |
| Name | User | Agent name | Count | CPU | Memory | IO | |
| w3wp.exe | IIS APPPOOL\smarts... | dotnet-official-demo | 1 | 47.9% | 16.4% | 0% | |
| java | root | demo-k8s-worker-04 | 3 | 36.9% | 11.1% | 0% | |
| java | whatap | virtual-java-agent | 6 | 29.6% | 30.9% | 0% | |
| java | whatap | virtual-java-agent | 6 | 26.4% | 29.7% | 0% | |
| java | root | demo-k8s-worker-02 | 7 | 25.2% | 30% | 0% | |
| java | root | demo-k8s-worker-02 | 7 | 22.8% | 30.4% | 0% | |
| w3wp.exe | IIS APPPOOL\demo4... | dotnet-official-demo | 1 | 14.4% | 24.1% | 0% | |
| node /home/ubun | ubuntu | node-official-demo | 3 | 14.1% | 19.9% | 0% | |
| rumsender | root | rum-official-demo | 1 | 10.3% | 2.7% | 0% | |
| rumstats | root | rum-official-demo | 1 | 9.7% | 3.6% | 0% | |
| whatap_infra.exe | NT AUTHORITY\SYS... | dotnet-official-demo | 2 | 6.6% | 0.3% | 0% | |
| whatap_agent | root | go-official-demo | 4 | 5.8% | 0.7% | 0% | |
| node.exe | EC2AMAZ-NO1BAVA... | dotnet-official-demo | 1 | 4.1% | 0.5% | 0% | |
| whatap_dotnet.exe | NT AUTHORITY\SYS... | dotnet-official-demo | 1 | 3.9% | 0.4% | 0% | |
| http_server | root | go-official-demo | 2 | 3.3% | 0.9% | 0% | |
| grpc_server | root | go-official-demo | 2 | 2.5% | 0.8% | 0% | |
| conhost.exe | EC2AMAZ-NO1BAVA... | dotnet-official-demo | 2 | 2.1% | 0.2% | 0% | |
| kuber | root | go-official-demo | 1 | 1.9% | 0.6% | 0% | |
| mariadb | mysql | mysql-official-demo... | 1 | 1.9% | 22% | 0% | |
| sqlservr.exe | NT SERVICE\MSSQL... | dotnet-official-demo | 1 | 1.8% | 5.1% | 0% | |
| echov4 | root | go-official-demo | 1 | 1.7% | 0.5% | 0% | |
| echo | root | go-official-demo | 1 | 1.6% | 0.5% | 0% | |
| mariadb | mysql | mysql-official-demo... | 1 | 1.6% | 22.4% | 0% | |
| kube-apiserver | root | demo-k8s-master | 1 | 1.6% | 9.4% | 0% | |

Compound Eye


Home > Select Project > **Dashboard** > **Compound Eye**

On the initial screen of the WhaTap Monitoring service, select a project and then select **Dashboard** > **Compound Eye**.



Compound Eye gives you a complete view of all servers where WhaTap agents have been installed. You can check the resource status of large-scale servers at a glance. It provides a visual view rather than sampling or summary for large-scale server resources.

Provided information




It displays information with an eye. It provides 5 different data.

- CPU: CPU usage
- Memory: Memory usage
- Disk: Disk usage
- Rx: Network Rx (received amount)
- Tx: Network Tx (transmitted amount)

⚠ Network Rx/Tx You can look at the network Rx/Tx usage metrics to see if external attacks like DDoS are coming from multiple servers at a time.

Anomaly occurrence



When an anomaly occurs on the server, each Eye expresses the condition with a color.

- **Red:** Fatal condition
- **Orange:** Alert condition
- **Gray:** Monitoring paused

In addition, when the mouse cursor is placed on the server, each individual pop-up message appears with numerical data. If you click on a server, the server's summary page appears for more information.

Traffic Max Value option

Because the amount of traffic may change depending on the network environment, the maximum traffic can be adjusted. If you set the desired maximum traffic value, the Rx/Tx graph changes depending on the value.

| | | | | | |
|-----------------|-------|--------------|---------|---------|-------|
| Chart max value | 1Kbps | 1Mbps | 100Mbps | 500Mbps | 1Gbps |
|-----------------|-------|--------------|---------|---------|-------|

Server Inventory Map

Home > Select Project > **Dashboard** > **Server Inventory Map**

On the initial screen of the WhaTap monitoring service, select a project, and then select the **Server Inventory Map** menu under the **Dashboard** menu. **Server Inventory Map** is a visualization dashboard that allows administrators to group and organize servers based on the classification criteria for monitoring purposes based on **Server Inventory**.

WhaTap allows users to intuitively check and analyze the server operation status through **Server Inventory Map**. The administrators can add and manage the user-defined information such as work classification, person in charge, and server location as custom columns, and easily create visualized dashboards based on this. **Server Inventory Map** also provides the grouping feature to meet different perspectives for monitoring purposes. Improve the server operation management efficiency through **Server Inventory Map**.

Requirements

❗ To use the **Server Inventory Map** dashboard, enable the **Server Inventory** and **Server Inventory Map** features respectively. Try out the new **Server Inventory Map** dashboard along with **Server Inventory**. For more information about enabling **Server Inventory**, see [the following](#).

Supported agents

The **Server Inventory Map** menu is supported in the following versions or later depending on the operating system.

- **Linux:** 2.7.2 or later

If an update is needed, check the following commands for each operating system.

- **Debian/Ubuntu** Amazon Linux/RHEL/Fedora/Oracle Linux/CentOS/Rocky Linux/SUSE

```
sudo apt-get update
sudo apt-get install whatap-infra
```

```
sudo yum update whatap-infra
```

- **Unix 1.3.7 or later**

- **AIX Solaris**

If an update is needed, check the following command.

```
stopsrc -s whatap-infra
cp /usr/whatap/infra/conf/whatap.conf /tmp/whatap.conf
rpm -e whatap-infra
rpm -Uvh https://repo.whatap.io/aix/noarch/whatap-infra-latest.noarch.rpm
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
startsrc -s whatap-infra
```

If an update is needed, check the following command.

```
cp /usr/whatap/infra/conf/whatap.conf /tmp
/etc/init.d/whatap-infra stop
gunzip -c whatap-infra.latest.SPARC.pkg.tar.gz | tar xvf -
yes n | ./whatap-infra/install.sh
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
/etc/init.d/whatap-infra start
```

- **Windows: 2.7.2 or later**

Download the installation file from the [link](#) and then run it. After the agent update, the service is restarted.

Activation option

The following settings are required to enable the **Server Inventory Map** menu viewing feature.

1. Enable the following options in the [whatap.conf](#) file.

```
categoryv2.enabled=true
```

2. After setting additional options according to the environment, restart the agent.

- **Linux**

- [Additional option commands](#) Restarting the agent

```
```bash
grep -q '^categoryv2.enabled' /usr/whatap/infra/conf/whatap.conf || echo 'categoryv2.enabled=true' | sudo tee -a /usr/
whatap/infra/conf/whatap.conf
```
```

```
```bash
sudo service whatap-infra restart
```
```

```
* **Unix**
```

- [Command to add options](#) Restart agent

```
```bash
grep -q '^categoryv2.enabled' /usr/whatap/infra/conf/whatap.conf || echo 'categoryv2.enabled=true' | tee -a /usr/whatap/
infra/conf/whatap.conf
```
```

```
* **AIX**
```

```
```bash title='AIX'
stopsrc -s whatap-infra
startsrc -s whatap-infra
```
```

```
* **Solaris**
```

```
```bash title='Solaris'
/etc/init.d/whatap-infra restart
```
```

```
* **Windows**
```


- [Add option command](#) Restart agent

```
```bash
if (-not (Select-String -Path "C:\Program Files\WhatapInfra\whatap.conf" -Pattern "^categoryv2.enabled" -Quiet)) { Add-
```

```
Content-Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "categoryv2.enabled=true" }
...
```

```
```bash
Restart-Service -Name "Whatap Infra"
```
```

## Basic screen guide



### 1 User configuration

Server Inventory Map provides several custom settings that allow you to configure the desired dashboard. You

can use the following features.


## Real-time inquiry and pausing

**Server Inventory Map** is basically based on the real-time inquiry. For real-time inquiry and pausing status change, use the time selector on the upper left of the screen.

| Lookup status            | Description                       |
|--------------------------|-----------------------------------|
| Real-time inquiry (Live) | Basic mode<br>Interval: 5 seconds |
| Pause inquiry (Pause)    | Real-time status change to OFF    |

! For more information on how to use the time selector.

## Grouping



**Server Inventory Map** provides the [grouping](#) feature that can systematically classify the servers. Manage servers effectively by setting the grouping criteria depending on the monitoring purpose and environment.

Based on **Server Inventory**, the default collection items and custom items among the columns that can be searched appear in the [Grouping](#) list. The default collection items are [server type](#), [default group](#), [OS type](#), [OS version](#), [model](#), and [serial](#). Other items that can be seen in the example image are custom columns ([custom](#)) that are user-defined items.

You can classify servers by selecting [1st group](#) and [2nd group](#) sequentially. According to the primary and secondary grouping, in the [3](#) area, the [primary group](#) classifies the data in card format, and the [secondary group](#) represents each grouped data in subgroup format in the chart area within cards.

❗ For more information about [grouping](#), see [the following](#).

## Label

**Server Inventory Map** provides the **Label** option that clearly displays the server information. When the desired **Label** option is selected, you can see the labels for each server in the chart.

| Label                        | Description                                                 |
|------------------------------|-------------------------------------------------------------|
| <a href="#">Server name</a>  | The actual hostname of each server appears as a label.      |
| <a href="#">Server alias</a> | It displays the user-specified alias as a label.            |
| <a href="#">IP Address</a>   | The IP address of the server appears as a label.            |
| <a href="#">CPU(%)</a>       | The current CPU utilization (%) appears as a label.         |
| <a href="#">Memory(%)</a>    | The current memory utilization (%) appears as a label.      |
| <a href="#">Swap(%)</a>      | The current swap memory utilization (%) appears as a label. |
| <a href="#">Disk Busy(%)</a> | The current disk utilization (%) appears as a label.        |
| <a href="#">Network</a>      | The network status appears as a label.                      |

## Filtering

**Server Inventory Map** provides the feature to filter the server types to display on the dashboard. Data applied through filtering is reflected into the server group chart, summary view, and detailed view.

- **Filter key:** Filters can be set based on the default collection columns provided in the **Server Inventory** menu and user-defined custom columns.



- Default collection columns: Items automatically collected from **Server Inventory** are **server type**, **default group**, **OS type**, **OS version**, **model**, and **serial** columns.
- Custom columns: They are the items added and managed on a project basis.
- **Condition:** Filter conditions can be set as numeric conditions such as  $\geq$ ,  $\leq$ ,  $>$ , and  $<$ , and string conditions such as Equal, Unequal, Includes, and Excludes.

## Inventory map configuration

**Server Inventory Map** provides the **Inventory map settings** option. Select the  icon on the upper right, and then enable or disable the following options.

- **Display no primary group:** Select whether to display server group cards with no settings for the **primary group** during **grouping**. The default settings appear.
- **Display the metric area at the bottom of group:** Select whether to display metric information at the bottom of the server group card. The default settings appear.


## 2 Server configuration status

**Server Inventory Map** provides the server configuration status that allows you to check the configurations for each server operating system of the project at a glance at the top of the dashboard. This allows users to get a quick and intuitive overview of key information about servers in the project.

- **OS:** This server operating system displays the OS name along with its OS icon. If an OS that is not predefined is included, the infrastructure icon appears by default.
- **Total / Active:** This indicates the total number of servers included in the project and the number of active servers.
- **Core:** It displays the sum of physical cores for all servers included in the project.

## 3 Server group card

**Server Inventory Map** provides individual cards based on the grouped server configuration. The **Server Inventory Map** card consists of three major elements as follows.



## a Group status

It provides the counts for each status of the servers that belong to the group name and group.

| Type  | Description                                                                                                                                   |
|-------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Total | It indicates the total number of servers included in the group. Servers are aggregated, categorized, and filtered depending on their purpose. |

| Type     | Description                                                                                                                                              |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Critical | It indicates the number of servers that generated events for the risk level of Critical.                                                                 |
| Warning  | It indicates the number of servers that generated events for the risk level of Warning.                                                                  |
| Normal   | It indicates the number of servers in normal status in which Critical and Warning events are not generated.                                              |
| Inactive | This indicates the number of servers where normal data collection is impossible due to server down, agent down, network communication failure, and such. |
| Cores    | It indicates the total number of cores that is the sum of physical core counts for each server within the group.                                         |

### **b** Server group chart

It provides visualization of the status and performance for servers within the group. When the [secondary grouping](#) is applied, data is additionally divided and displayed in subgroup format within the chart. Additionally, if you hover your mouse over a subgroup, you can briefly check the main status of the servers included in the corresponding group in popover format.

### **c** Metrics

It provides brief information on key metrics for each group. For more information about each metric, see the following table.


| Type      | Description                                                                                                        |
|-----------|--------------------------------------------------------------------------------------------------------------------|
| CPU(%)    | It displays the sum of the utilization in user mode (User) and kernel mode (Sys) out of the total CPU utilization. |
| Memory(%) | It displays the memory rate used by user processes and kernels, excluding the buffer and cache areas.              |
| Swap(%)   | It displays the size of the swap area in use as a percentage.                                                      |

| Type            | Description                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------|
| Disk<br>Busy(%) | It is the rate of the time processed for each disk per unit time, and displays the value for the most utilized disk on the server. |
| Network<br>RX   | It displays the amount of received (Rx) data in bps.                                                                               |
| Network<br>TX   | It displays the amount of transmitted (Tx) data in bps.                                                                            |

- [Avg](#): It displays the average metric value for all servers that belong to the group.
- [Max](#): It displays the maximum value of the metric.
- [Max Server](#): It displays the name of the server that recorded the maximum value.

## Summary view

When you select a group in the server inventory map chart, you can see the **summary view** panel on the right. The summary view panel provides the [Group Info](#) and [Metrics](#) tabs.



### Display Detail

Click [Display Detail](#) on the upper right of the summary view panel. For more information about [Display Detail](#), see [the following](#).

## Group Info

In the **Group Info** tab, you can effectively check the main operation status and resources of the selected server

group as follows.

## Server uptime

It divides the operation period into sections since the server booting and visualizes the number of servers in the section in bar chart format. When a bar chart is selected, the detailed information of servers belonging to the section is highlighted in the **Server Inventory** list below. If you click the same bar chart, the selection is released.

## Server Inventory

You can see key information about the servers in the selected server group in list format. You can filter the data by entering the hostname or IP address in the input fields on the right. It provides [server name](#), [IP address](#), [OS version](#), [CPU cores](#), [memory \(GB\)](#), and [server uptime](#) columns.

## File system usage information

You can intuitively see the status of the server's file system. Mount points for each server are represented as a block. The block colors indicate the utilization (Used, %) of the mount point. In the same color range, the color intensity increases as the usage rate rises.

- **Red:** It indicates that the usage rate is 90% or more.
- **Orange:** It indicates that the utilization is 80% or more and less than 90%.
- **Green:** It indicates that the utilization is less than 80%.

ⓘ When you hover the mouse over a mount point block, you can see the server name, mount point, utilization, and total capacity.

## Metrics

The **Metrics** tab allows you to view key server performance metrics in chart format. You can intuitively check the resource status by group.

| Type                 | Description                                                                                                                                                                |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [CPU] User + Sys (%) | It displays the sum of the utilization in user mode (User) and kernel mode (Sys) out of the total CPU utilization. This displays the overall utilization of CPU resources. |
| [MEM] Used (%)       | It displays the rate of memory used by the user processes or kernel. Only pure memory usage is reflected, excluding the buffer and cache areas.                            |
| [SWAP] Used (%)      | This value is the current swap memory size in use expressed in percentage.                                                                                                 |
| [NET] Rx (bps)       | It displays the amount of data received (Rx) from the network in bits per second (bps).                                                                                    |
| [NET] Tx (bps)       | It displays the amount of data transmitted (Tx) over the network in bits per second (bps).                                                                                 |
| [Disk] Read (Bytes)  | It displays the sum of the read speeds of the entire disk in bytes per second.                                                                                             |
| [Disk] Write (Bytes) | It displays the sum of write speeds per second of the entire disk in bytes.                                                                                                |


## Events in progress

You can see the list of active events in the server group selected in the **Active event status** tab.

# Server Inventory Map Details

Home > Select Project > **Dashboard** > **Server Inventory Map**

You can check detailed information about the server group selected from the **Server Inventory Map** menu on the [Display Detail](#) screen. Click [Display Detail](#) on the upper right of the **View summary** panel. The server group information in Display Detail is updated in real time.



- You can check the group data query time and refresh cycle through the **time selector** at the top, and you can stop or resume refreshing data.
- In the left area, you can see the **server group cards** and **server list**.
- In each tab on the right, you can see the [server inventory](#), [file system information](#), and [metrics information](#).



## Server group cards and list

Key summary information on the selected server group can be viewed through the server group cards and list.

- If you hover the mouse over a server in the chart area within the server group card, a popover with summary information of the selected server appears.
- You can specify a target server to query from the chart area and server list within the server group card.
- When selecting a server from the server list, you can select all or deselect all.
- When selecting a server, you can see the information of the selected server in individual tabs on the right.

## Server Inventory

You can check the detailed configuration of the selected server in list format in the **Server Inventory** tab. If no server is selected, the server configuration information for all servers in the server group appears. The information and configuration that can be viewed in the **Server Inventory** menu are the same.

### Filtering

The **Server Inventory** menu provides the server filtering feature based on the column list set in the menu.

- **Filter key:** Filters can be set based on the default collection columns provided in the **Server Inventory** menu and user-defined custom columns.
  - Default collection columns: Items automatically collected from **Server Inventory** are **server type**, **default group**, **OS type**, **OS version**, **model**, and **serial** columns.
  - Custom columns: They are the items added and managed on a project basis.
- **Condition:** Filter conditions can be set as numeric conditions such as  $\geq$ ,  $\leq$ ,  $>$ , and  $<$ , and string conditions such as Equal, Unequal, Includes, and Excludes.

### Column Settings

In the [column settings](#) of the **Server Inventory** tab in the [Display Detail](#) of **Server Inventory Map**, you can set the **display status** and **order** of the columns in the list. The display status and order are set separately from the

**Server Inventory** menu.

! Adding and modifying custom columns are possible through the **Server Inventory** menu.

## Download CSV

You can download the collected server configuration data in **CSV** file format by selecting **Download** at the top.

## File system information

In the **File system information** tab, you can see the file system information of the server selected from the server group cards or list on the left in list format. If no server is selected, the file system information for all servers in the server group appears.

| Type            | Description                                                                                                               |
|-----------------|---------------------------------------------------------------------------------------------------------------------------|
| Mount Point     | It displays the name of the mounted file system. In the Windows environment, it displays the name of the drive partition. |
| Total Space(GB) | It displays the total storage space allocated to the file system in GB.                                                   |
| Used Space(GB)  | It displays the storage space currently used on the file system in GB.                                                    |
| Used(%)         | It displays the percentage of storage space in use compared to the total storage space on the file system.                |
| i-node Used(%)  | It displays the percentage of i-nodes allocated to the file system that are in use.                                       |
| Free(GB)        | It displays the free space on the file system in GB.                                                                      |

### ! i-node

An i-node is a space in the file system that stores the data about files and directories. When this usage reaches 100%, you can no longer write files, even if **Used(%)** is not 100%.

## Filtering

The server filtering feature is provided. **Filter Key** provides the **Server Name** and **Mount Point**. The filter, **condition** can be set as a string such as Equal, Unequal, Includes, and Excludes.

## Download CSV

You can download the collected file system information in **CSV** file format by selecting **Download** at the top.

## Metrics

You can see the performance metrics charts for the server selected from the server group cards or list on the left in the **Metrics** tab. If no server is selected, the performance metrics chart for the top 50 servers within that server group appears.

| Name                 | Description                                                                                                                                             | Remarks |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|---------|
| [CPU] User + Sys (%) | It displays the sum of the usage in user mode (User) and kernel mode (Sys) out of the total CPU usage, indicating the utilization of all CPU resources. | -       |
| [CPU] User (%)       | The percentage of CPU consumed by processes running in user mode.                                                                                       | -       |
| [CPU] Sys (%)        | The percentage of CPU consumed by processes running in kernel mode.                                                                                     | -       |
| [CPU] Wait (%)       | The percentage of time when the CPU is idle, waiting for an I/O operation, such as the disk and network.                                                | -       |
| [System] Runqueue    | It is the number of processes (or threads) that are ready to run but have not yet been assigned CPU.                                                    | -       |

| Name                      | Description                                                                                                                        | Remarks                                    |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|
| [System]<br>Blockqueue    | This indicates the number of processes (or threads) in the queue for I/O, paging, and such.                                        | Windows excluded                           |
| [System] Forks            | It displays the number of tasks that generate new processes per second.                                                            | Windows excluded                           |
| [System]<br>Interrupts    | It displays the number of cases per second where the CPU handled hardware interrupts.                                              | -                                          |
| [System]<br>SystemCall    | It displays the number of cases per second where the kernel processed system calls.                                                | Linux excluded                             |
| [System]<br>ContextSwitch | It displays the number of cases per second where switching occurs between processes.                                               | -                                          |
| [MEM]<br>Used+Cache(%)    | It displays the total ratio of Used memory and Cached memory among the total memory.                                               | For Windows, <b>[MEM] Used(%)</b> appears. |
| [MEM] Used(%)             | It is the memory utilization indicating the ratio of memory use by user processes or kernel, excluding the buffer and cache areas. | -                                          |
| [MEM]<br>Cache(%)         | It displays the ratio of cached memory use to quickly serve frequently accessed files or data.                                     | Windows excluded                           |
| [Swap]<br>Used(%)         | It is the size of the swap in use expressed as a percentage.                                                                       | -                                          |
| [MEM] SwapIn              | It displays the number of pages moved from the swap (paging space) area of the disk into memory.                                   | -                                          |
| [MEM]<br>SwapOut          | It displays the number of pages moved from the memory to the disk swap (paging space) area.                                        | -                                          |
| [MEM]<br>FilesystemIn     | It displays the number of pages loaded into memory based on the read operation on the file system.                                 | -                                          |

| Name                           | Description                                                                                                        | Remarks |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------|---------|
| [MEM]<br>FilesystemOut         | It displays the number of pages written to disk based on the write operation on the file system.                   | -       |
| [NET] Rx+Tx<br>(bps)           | It is the value expressed in bps for the sum of the data received (Rx) and transmitted (Tx).                       | -       |
| [NET] Rx+Tx<br>(pps)           | It is the value expressed in packets per second (pps) for the sum of the data received (Rx) and transmitted (Tx).  | -       |
| [Disk] Read +<br>Write (bytes) | It displays the sum of the data read and written on the disk in bytes.                                             | -       |
| [Disk] Read +<br>Write (IOPS)  | It displays the number of disk read and write operations in the unit of IOPS (Input/Output Operations Per Second). | -       |
| [Process]<br>TotalCount        | It displays the current total number of processes.                                                                 | -       |
| [Thread]<br>TotalCount         | It displays the current total number of threads.                                                                   | -       |
| [Defunct]<br>TotalCount        | It displays the current total number of defunct processes.                                                         | -       |

## Events in progress

You can see the events in progress in the server selected from the server group cards or list on the left. You can see **Event title**, **Event time**, **Server name**, and **Message** in the list.


# Server Inventory Map Grouping

Home > Select Project > **Dashboard** > **Server Inventory Map**

In the **Server Inventory Map** dashboard, you can see the entire server information at a glance based on the columns set in the **Server Inventory** menu. However, when monitoring the server status and information, the criteria that each user considers important may differ. By reflecting this, **Server Inventory Map** provides the feature to group the data based on user concerns.

For example, application managers focus on the status for server groups or individual servers that relate to a specific service, while infrastructure administrators are primarily interested in server hardware configuration information and resource utilization. The administrators may want to analyze the server status based on resource changes or key events over a specific period.

To accommodate various perspectives and needs, **Server Inventory Map** provides the flexible grouping feature. This allows users to reorganize data according to their interests, allowing them to monitor and manage the server status more efficiently.



## Grouping & Card

**Server Inventory Map** provides cards by grouped server configuration information. When grouping is not applied, all server information is displayed within a single card. When applying grouping, cards are organized by each group unit.

When applying grouping, you can classify servers by sequentially selecting **1st group** and **2nd group**. According to the primary and secondary grouping, **1st group** classifies data in card format, and **2nd group** presents each grouped data in subgroup format in the chart area within cards.

- **1st group:** The primary group displays data based on the **group category** and **group name**.

e.g. `Infra_domain` **ALL**

- **2nd group:** Secondary groups are displayed based on the **group name** and **number of servers in the group**.

e.g. **Finance (31)**

- If you hover over a server group or server, you can briefly check the major status of the group and server in a popover format.
- When you select **Not set** among the grouping selection items, the grouping definition is initialized.

## Key grouping properties

Based on **Server Inventory**, the default collection items and custom items among the columns that can be searched appear in the **Grouping** list. The default collection items are **server type**, **default group**, **OS type**, **OS version**, **model**, and **serial**.

Other items such as **infrastructure domain** (`Infra_Domain`) and **business major category** (`Biz_MainCategory`) that can be seen in the example image, are user-defined items added through the custom column (`custom`) in the **Server Inventory** menu.

You can further refine the view by **primary grouping** servers with the **infrastructure domain** based on physical or logical configuration, and then **secondary grouping** them based on **business-based major classification** for business purposes, through custom items.



|           | Freq(MHz) | Memory(GB) | All IP                      | Time Zone |                     | Infra_Domain <small>custom</small> | Biz_MainCategory <small>custom</small> |
|-----------|-----------|------------|-----------------------------|-----------|---------------------|------------------------------------|----------------------------------------|
| server758 | 4,250     | 8          | 54.142.159.829,54.142.1...  | KST-9     | 2024-08-04 02:40:59 | ALL                                | External_System                        |
| server885 | 4,250     | 4          | 254.186.171.1093,254.1...   | KST-9     | 2024-08-23 12:59:43 | WEBWAS                             | Homepage                               |
| server902 | 4,250     | 32         | 135.175.32.952,135.175...   | KST-9     | 2024-08-30 08:53:12 | ALL                                | Homepage                               |
| server0   | 4,250     | 4          | 210.4.178.230,210.4.178...  | KST-9     | 2024-08-26 07:55:20 | DB                                 | Processing_System                      |
| server9   | 4,250     | 4          | 246.216.137.957,246.21...   | KST-9     | 2024-08-30 19:45:12 | WEBWAS                             | Finance                                |
| server958 | 4,250     | 8          | 118.4.94.1209,118.4.94.1... | KST-9     | 2024-08-22 20:36:20 | DB                                 | Finance                                |
| server975 | 4,250     | 16         | 248.61.3.1003,248.61.3.1... | KST-9     | 2024-08-05 22:08:48 | DB                                 | External_System                        |
| server993 | 4,250     | 4          | 153.144.222.1237,153.1...   | KST-9     | 2024-08-15 03:44:03 | WAS                                | Finance                                |
| server373 | 4,250     | 16         | 79.4.116.362,79.4.116.36... | KST-9     | 2024-08-05 05:54:49 | WEB                                | Processing_System                      |
| server659 | 4,250     | 32         | 241.234.162.736,241.23...   | KST-9     | 2024-08-05 19:08:40 | WEBWAS                             | Bancassurance                          |
| server667 | 4,250     | 16         | 220.224.107.784,220.22...   | KST-9     | 2024-08-04 19:36:34 | WEB                                | Homepage                               |

The example image displays the **infrastructure domains** in the **Server Inventory** menu sorted by the values such as **DB**, **WAS**, **WEB**, **WEBWAS**, and **ALL**, which includes custom infrastructure domains. Additionally, the **business-based major classification** items have been classified into values such as **External\_System**, **Processing\_System**, and **Finance**.

Based on this information, you can configure the **Server Inventory Map** dashboard to monitor the overall operation status.

## Metrics data by card

Below each card, a brief overview of key metrics information about the grouped servers is provided. For more information about each metric, see the following table.

| Type      | Description                                                                                                        |
|-----------|--------------------------------------------------------------------------------------------------------------------|
| CPU(%)    | It displays the sum of the utilization in user mode (User) and kernel mode (Sys) out of the total CPU utilization. |
| Memory(%) | It displays the memory rate used by user processes and kernels, excluding the buffer and cache areas.              |
| Swap(%)   | It displays the size of the swap area in use as a percentage.                                                      |


| Type            | Description                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------|
| Disk<br>Busy(%) | It is the rate of the time processed for each disk per unit time, and displays the value for the most utilized disk on the server. |
| Network<br>RX   | It displays the amount of received (Rx) data in bps.                                                                               |
| Network<br>TX   | It displays the amount of transmitted (Tx) data in bps.                                                                            |

- **Avg**: It displays the average metric value for all servers that belong to the group.
- **Max**: It displays the maximum value of the metric.
- **Max Server**: It displays the name of the server that recorded the maximum value.

# Flex Board

**Flex Board** is a custom integrated dashboard. You can create a real-time dashboard with the project data such as applications, servers, databases, and containers.

You can quickly complete the initial setup with preconfigured templates and build dashboards in your preferred format. By adding data widgets or modifying properties, you can display information in the form you need, and with filtering functions, you can narrow down monitoring targets. When you set a time range, you can check data from a specific point in time, and by using auxiliary charts, you can analyze data in various ways. Dashboards can be added to favorites for easy access, and personalized dashboards can be copied to other accounts for reuse.



The usage scope of the Flex board differs depending on the following entry:

### Home > Integrated Flex Board

- When creating a widget, all projects that can be viewed are provided as options.
- The dashboard is saved in your account and can be shared with other users with the "To copy" feature.
- Individual account dashboards are not affected by the role, but the shared dashboards are read-only and cannot be edited.

## Home > Select Project > Dashboard > Flex Board

- When creating a widget, the related project data is automatically entered.
- The created Flex board is automatically shared with other project members.
- The users with the project **Edit** role or **Project flexboard update** role, and **site admin** role can modify the board.

✓ Users with edit permissions for the Flex Board can use the following features:

- Export and import dashboards as JSON files
- Check data requests and responses within the dashboard
- View and edit widget configuration options in JSON format

ⓘ In the **Flex Board** menu within the project, only the users with the dashboard editing role can access **Edit Mode**, **Admin Mode**, and **Filter**.

- Accessible member permissions
  - Project **Edit** role
  - **Project flexboard update** role
  - **Site Admin** role
- For more information about the **View mode** and **Filter**, see [the following](#).
- For more information about the member roles, see [the following](#).

# Using Flex Board

- ✓ You can build dashboards using various widget templates, including metrics widgets.  
For details on the features and usage of each template, see the [Widget template](#) document.

## Apply a preset template

Use a preset template to speed up the initial setup, then modify it as needed.

1. Select the **Flex Board** menu.
2. From the **Template** list on the left, select the template you want.
3. When the **Flex board management** window appears, enter a name in **Rename**.
  - In **Layout**, select a widget and drag it to the position you want.
4. Click the **Save** button. A Flex board is created based on the selected template.
  - You can check the created board in the board list of **Flex Board**.

## Create a Flex board

You can create a Flex board and place widgets to build your own dashboard.

1. In **Flex Board**, click the **Create Dashboard** button at the top right.
2. When the **Create Dashboard** window appears, enter the dashboard **Name** in the input box.
3. Select a widget placement method, then click the **Create Dashboard** button.
  - **Fixed layout flexboard**: free widget placement with pixel basis
  - **Responsive flexboard**: responsive grid layout based on browser size (configure the grid, then click the **Create Dashboard** button)
4. In the **Widget templates** selection screen, choose the widgets to place on the dashboard.

- To add metrics widgets, click the **All Metrics** button on the right of **Widget templates**.

- ! • **Widget templates:** choose from a predefined widget list that summarizes key indicators commonly used in monitoring.
- **All metrics:** create widgets based on all metrics data collected in your project. For details about metrics widgets, see [the following document](#).

5. Move the placed widgets or adjust their size.
  - When a crosshair cursor appears at the top of a widget, click and drag the widget to move it.
  - When an arrow cursor appears at the bottom right of a widget, click and drag to resize it.
6. After you finish placing widgets, select **View mode** to review the layout. To edit the layout again, select **Edit Mode** and adjust widget placement.
7. When all steps are complete, click the  (**Go to list**) icon at the top right.
  - You can check the created board in **Flex Board** under **Dashboard list**. Select the created board to view the new Flex board.

- ! We plan to add dashboard templates or widgets based on customer feedback. If you need a specific dashboard template or widget, contact us at [support@whatap.io](mailto:support@whatap.io).

## Edit a Flex board

### ! Flex board edit permissions

- **Home > Integrated Flex Board:** this is a personal account dashboard, so permissions do not apply; however, dashboards shared as read only cannot be edited.
- **Home > Select project > Dashboard > Flex Board:** users with project edit permission, project Flex board edit permission, or Site Admin permission can edit.

For details on member permissions, see [the following document](#).

1. In **Flex Board** under **Dashboard list**, click the edit icon . The **Flex board management** window appears on the right.
2. In the **Flex board management** window, configure Name, Project, Layout, widget.json, and option.json.
  - **Rename:** edit the dashboard name.
  - **Project:** select the project from which widgets will fetch data.

 You can configure the **Project** option only in **Home > Integrated Flex Board**. If you entered from **Integrated Flex Board**, you can select the data source project for each widget.

1. Click the **+** button at the top right of a widget, then search and select one or multiple projects.
2. After selecting projects, click the **Apply** button.

- **Layout**
  - Drag widgets to move them.
  - Click and drag the bottom right of a widget to resize it.
  - To delete a widget, click the  icon at the top right of the widget.
- **Widget.json:** import dashboard widget settings in JSON format. Click the  icon to copy the JSON.
- **Option.json:** import dashboard option settings in JSON format. Click the  icon to copy the JSON.

## Delete a Flex board

1. In **Flex Board > Dashboard list**, click the  icon for the item to delete.
2. In the confirmation message, click the **Delete** button.

 Dashboard deletion behavior differs for owners and shared users.

- If the dashboard owner deletes a dashboard, it is deleted for all user accounts with whom it was shared.





- If a user who received a shared dashboard deletes it, it is deleted only from that user's **Dashboard list**. The original dashboard remains.


## Share a Flex board



- The Flex board sharing feature is available only in **Home > Integrated Flex Board**. It is not provided in a project's **Flex board** menu.



### Difference between share and copy



- Share: all members you share with view the same dashboard.
- Copy: if you copy a dashboard to another member or to yourself, you view a separate copied dashboard.
- Dashboards shared or copied as **Read only** cannot be edited. They can be edited only if shared as **Edit mode**.
- Members who receive a dashboard as **Read only** can also **re-share as read only** or **copy** that dashboard.



- You can share a dashboard only with members who belong to at least **one** of the projects included in the dashboard. Check whether the member to share or copy to belongs to the relevant project.

☆ test

K8S-JINS-OKD-TEST, K8S-JINS-TEST-MICROK8S, DB-M...

Updated: 2023/12/22 11:10:24 AM

You can share a Flex board created in **Integrated Flex Board** with other members. When the owner edits a shared dashboard, the changes apply to all members it was shared with.

- In **Integrated Flex Board** under **Dashboard list**, click the icon for the dashboard to share.
- When the **Sharing a dashboard** window appears, select one or multiple members to share with, or enter accounts directly.
  - If you select **Read only**, members cannot edit the shared dashboard.
  - Share targets must belong to at least one of the projects included in the dashboard.
- Click the **Share** button.
  - Shared items display the **Shared** tag. Hover over the **Shared** tag to view information about the members you shared with.
  - Items shared as read only display the **Read only** tag.

| When you receive a shared dashboard | When you share a dashboard (dashboard owner) |
|-------------------------------------|----------------------------------------------|
|                                     |                                              |

### ✓ Dashboard sharing conditions

- The dashboard sharing feature is available only in **Home > Integrated Flex Board**.



- Anyone can share their own dashboard, and you can re-share or copy a dashboard you received.
- Dashboards received as **Read only** cannot be edited; dashboards received as **Edit mode** can be edited.
- If the **owner** deletes a dashboard, it is deleted for all accounts that received it.
- If a **recipient** deletes a dashboard, it is removed only from that user's **Dashboard list** and the original remains.
- For projects in a received dashboard where you do not have **View** permission, you cannot view data.

## Copy a Flex board

Users can copy a Flex board and send it to other members.

1. In **Integrated Flex Board** under **Dashboard list**, click the  icon.
2. Choose a copy option.
  - **Copy to me:** copy to your own account.
    - a. An item with the `_copied` suffix is added to **Dashboard list**.
  - **Copy to others:** copy to another member's account.
    - a. Click **+ Account**.
    - b. Select the email of the member to copy to and click the **Copy dashboard** button.
    - c. An item with the `_copied` suffix is added to that member's **Dashboard list**.

## Share as a JSON file

You can export Flex board settings to a JSON file to share with other members, or import settings from other members.

## Export

1. Go to **Integrated Flex Board** or a project's **Flex board** menu.
2. In **Dashboard list**, click the  button.
3. Choose an export option.
  - **Include Project Information:** include dashboard project information in the JSON file when sharing.
  - **Except Project Information:** share without including dashboard project information in the JSON file.
4. Download the JSON file and send it to the member you want to share with.

## Import

1. Go to **Integrated Flex Board** or a project's **Flex board** menu.
2. Click the  **Import** button at the top right.
3. Select the downloaded JSON file.

 Only members with **Edit** permission for the project can use this feature.

## Flex board screen modes

Features vary by Flex board screen mode.

### Accessible member permissions

Only users with **Edit** permission in a project's **Flex Board** menu can access **Edit Mode**, **Admin**, and **Filter**. For details on member permissions, see [the following document](#).

- **Project Edit** permission
- **Project flexboard update** permission
- **Site Admin** permission

## Modes

Available features differ by **Flex Board** screen mode.

- **Edit Mode:** place widgets from the widget template list on the Flex board layout; edit, delete, or reposition widgets.
- **Admin:** manage widget properties in JSON format. Click the  button at the top right of a widget on the layout.
- **View mode:** after finishing layout, sizing, and detailed options, monitor each widget's data in real time like a dashboard.

## Common features

- Data filtering
- Real time data time range settings
- Print mode
- Flex board management
- Data merge options

## Data filtering

Use data filtering to narrow monitoring targets and build an efficient Flex board.

1. At the top of the screen, click the  **Filter** tab.
2. In the **Filter** window on the right, select filter targets. To select all filter targets, choose **Select all**.
  - **Agent** option: select agent types or an application list, with **Search** available
  - **Metrics** option: select a project or agent type; add **+ And** conditions, or auto select by string with **Contain** or **Exclude**
3. After selecting filter targets, click the **Apply** button.
4. Save frequently used conditions with the **Save filter** button.
  - a. Enter a filter name and select **Save**. **Saved filters** appears at the bottom of the **Filter** window.

## Filtering conditions

Filtering conditions are divided into agent type and metrics type.

Agents provide only basic monitoring targets as selectable options, while metrics provide selectable tags for all metrics data collected in the project. The metrics type is useful when you collect data in a non agent form and want to apply filtering. For details on metrics, see [the following document](#).

| Type              | Agent                                                                                                             | Metrics                           |
|-------------------|-------------------------------------------------------------------------------------------------------------------|-----------------------------------|
| Selection options | WhaTap project monitoring targets<br>(applications, servers, databases,<br>containers, Kubernetes PODs and NODEs) | All metrics tags including agents |
| Scope             | All widgets                                                                                                       | All widgets that use metrics data |
| Filtering method  | Checkboxes                                                                                                        | <b>Equal, Contain, Exclude</b>    |

## Set the lookup time

Use the time selector at the top right of **Flex Board** to set data lookup time for widgets.

- : stop real time mode and set a specific lookup time
- : switch back to real time mode

## Print

You can save the Flex board screen as a PDF and print it.

1. On the Flex board screen, click the  icon at the top right to switch to print mode.
2. Configure options.
  - **Viewing filter options:** toggle filter options on the screen
  - **Convert to {mode} orientation / Convert to {mode} orientation:** change orientation to horizontal or vertical
3. After finishing, select **PDFDownload**.

4. To exit print mode, click the  **Exit print mode** at the top right.

## Flex board management

You can change the Flex board layout; move, resize, or rename widgets; and use other management features. You can also open **Flex board management** by clicking the  icon in **Dashboard > Flex Board** under **Dashboard list**.

1. On the Flex board screen, click the  icon at the top right to open the **Flex board management** window.
  - **Rename:** change the Flex board name
  - **Layout:** move, resize, and delete widgets
  - **Don't Allow Other Users to Modify.:** restrict modifications
    - Members with edit permission can edit, delete, and modify; if you enable the option, it changes to **Don't Allow Other Users to Modify.: On**
  - **Widget.json:** edit and manage widgets on the Flex board in JSON
  - **Option.json:** edit and manage Flex board settings in JSON (project info, Flex board filters, time selection, and other settings)

 The **Widget.json** and **Option.json** features are available only to dashboard administrators.

2. After completing all settings, click the **Save** button.

## Add a dashboard to favorites

You can add a Flex board you created to favorites for quick access. In **Dashboard > Flex Board** under **Dashboard list**, click the  button for the dashboard to add. It changes to .

# Managing The Widgets

You can customize your dashboard by handling the widgets placed on the Flex board.

- ✓ You can build dashboards using various widget templates, including metrics widgets.  
For details on the features and usage of each template, see the [Widget template](#) document.

## Adding a widget

You can add any widget to the created Flex board.

1. In **Dashboard list** of **Flex Board**, select a dashboard to add the widget.
  2. On the upper left of the screen, select **Edit Mode**.
  3. Select a desired widget from the **Widget templates** list on the left. To add a metrics widget, select **All metrics** and then click a desired widget.
- If you have entered **Integrated Flex Board** on the Home screen and added a widget, click **Project selection** and then select target projects to collect.
  - If you have more widgets to add, repeat the same procedure.
4. After all steps are finished, select **View mode** to check the layout and then select  (**Go to list**).

## Changing the widget attributes

In **Dashboard list** of **Flex Board**, select a board to change the widget properties. It goes to **Edit Mode**.

### Changing the widget name

1. On the right of the widget name to change, select .
2. Set the widget values for **Title**, **Font size**, and **Color**, and then select **Save**.



- If the **Apply to the entire widget** toggle button is enabled, the font size and color are applied to other widgets.
- To lock the widget title, select [Configuration icon](#) and then click **Save**.

## Time Setting

1. On the upper right of the widget to set the time, select ⌚.
2. Set the time in the **Select a time** window.
  - To apply the set time to the dashboard, check **Use dashboard settings**. The time setting on the dashboard is located on the upper right.
  - To apply a desired time setting, uncheck **Use dashboard settings**. Use < or > to set the desired time.
3. After setting the time, select ✕.

ⓘ Depending on the data type of the widget, the time setting may not be available.

## Setting the output data options

By setting the output data options of the chart widget, you can select the desired data.

1. At the top right of the widget, click the ⋮ icon to configure the output data options.
  - You can filter data by selecting types for project, agent, application, etc.
  - You can change the type of the graph displayed on the chart by selecting individual or merged data.
  - Select ⋮ > ⚙️. You can change the option details through the **Chart settings** window. The options that can be set are as follows.
    - **Chart type:** This option allows you to display the current aggregated figures on the chart.
    - **y axis Chart max value:** The chart maximum value on the Y axis can be selected with one of **Auto** and **Fix**.
    - **Set horizontal line:** Select **Add** to display the threshold or range as a horizontal line. You can

apply multiple items added. After adding the horizontal line, select **Apply** to apply it to the chart.

- **Show chart max value:** The maximum value of the graph is displayed on the chart.
  - **Sub-chart:** You can add a sub-chart to the main charts of the widget. Select a type of sub-chart and then select its layout, size, and whether or not to merge data. To apply to other widgets, select **Apply to the entire widget**. Sub-charts are useful to view data in various ways.
2. After setting all options, on the left of the **Chart settings** window, click **×** or click the outside of the **Chart settings** window.

ⓘ Depending on the widget's data type, available options may differ.

## Changing the metrics widget data

You can change the data query conditions of the metrics widget.

1. In **Flex Board > Dashboard list**, select the dashboard where you want to change the widget data.
2. In **Edit Mode**, click the **⋮** icon at the top right of the widget you want to modify.
3. Click the **<>** icon to open the data settings edit window.
4. Change the settings referencing the following, and then select **< Apply to the widget**.
  - Metrics options that are created by default when creating a widget appear on the tab.
  - Metrics options are **Category**, **Field**, and **Tag**.
  - The **Category** currently restricts the editing function.
  - In **Field** and **Tag**, multiple items can be selected. The chart data is grouped using the selected tag options.
  - Select **🔍** to query the original data.
  - For more information about **Object Merge** and **Time merge**, see the [Data merge options](#).

## Data merge options

The metrics widget placed on the **All metrics** list of the Flex board provides the data merge options. Data merge provides the following methods: **Object Merge** and **Time merge**.

- **Object Merge**


It is used to group the chart data based on the tag (e.g. agent name) selected on the left. At this time, if the tags match among the data with different fields (e.g. CPU), the corresponding data is merged.

- **Time merge**

It is mainly used if the lookup time is too long. Data with the same fields in the original data are merged at regular time intervals (5 minutes or 1 hour) depending on the time range. This method is called "Time Merge."

- **AWS ColudWatch Metrics**

In case of AWS monitoring metrics, the statistical method is added after the field name in the form of suffix. Data merging defaults reference this suffix.



1. In **Edit Mode**, select **All metrics** and then change the **Widget templates** list to the **All metrics** list.
2. Place the metrics widget on the layout.
3. On the upper right of the metrics widget, select **:**.
4. If a tag (e.g. agent name) option is selected, the data merge options appear. Depending on the widget size, the metrics option's location may differ.

## Copying a widget

1. Click  on the upper right of the widget.
2. Click .
3. If the **Copy** window appears, click **+ Add widget**.
4. Change the widget name in the list.
5. Click **Apply**.
6. In **Edit Mode**, you can see the copied widget.

## Delete Widget

1. Click  on the upper right of the widget.
2. Click .
3. If the confirmation message appears, click **Delete**. To cancel the deletion, select **Cancel**.

# Server Inventory

Home > Select Project > **Server List** > **Server Inventory**

On the initial screen of the WhaTap monitoring service, select a project, and then select the **Server Inventory** menu under the **Server List** menu. **Server Inventory** provides configuration information required for server operations and allows you to manually add management items for management convenience.

The **Server Inventory** menu systematically manages the physical and virtual characteristics of the server, as well as the operating environment. This feature allows for detailed insight and management of the server configuration through automatically collected server components and custom items. **Server Inventory** provides essential data for performance analysis and troubleshooting of the server.

## ⓘ Supported agent versions

The **Server Inventory** menu is supported only in the following versions or later for each operating system.

- **Linux:** 2.6.6 or later

If an update is needed, check the following commands for each operating system.

- **Debian/Ubuntu** Amazon Linux/RHEL/Fedora/Oracle Linux/CentOS/Rocky Linux/SUSE

```
sudo apt-get update
sudo apt-get install whatap-infra
```

```
sudo yum update whatap-infra
```

- **Unix:** 1.3.4 or later

If an update is needed, check the following commands for each operating system.

- **AIX** Oracle Solaris

```
stopsrc -s whatap-infra
cp /usr/whatap/infra/conf/whatap.conf /tmp/whatap.conf
rpm -e whatap-infra
rpm -Uvh https://repo.whatap.io/aix/noarch/whatap-infra-latest.noarch.rpm
```



```
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
```

```
cp /usr/whatap/infra/conf/whatap.conf /tmp
/etc/init.d/whatap-infra stop
pkgrm whatap-infra
wget https://repo.whatap.io/sunos/11/whatap-infra.latest.SPARC.pkg.tar.gz
tar xzvf whatap-infra.latest.SPARC.pkg.tar.gz
pkgadd -d .
cp /tmp/whatap.conf /usr/whatap/infra/conf/whatap.conf
```

- **Windows:** 2.6.6 or later

Download the installation file from the [link](#) and then run it. After the agent update, the service is restarted.

## Activation options

The following settings are required to enable the **Server Inventory** menu viewing feature.

1. Enable the following options in the [whatap.conf](#) file.

```
hostinfov2.enabled=true
```

2. After setting additional options according to the environment, restart the agent.

- **Linux**

- **Additional option commands** Restarting the agent

```
```bash
grep -q '^hostinfov2.enabled' /usr/whatap/infra/conf/whatap.conf || echo 'hostinfov2.enabled=true' | sudo tee -a /usr/
whatap/infra/conf/whatap.conf
```
```

```
```bash
```

```
sudo service whatap-infra restart
'''
```

```
* **Unix**
```

- [Command to add options](#) Restart agent

```
'''bash title='AIX/Oracle Solaris'
grep -q '^hostinfov2.enabled' /usr/whatap/infra/conf/whatap.conf || echo 'hostinfov2.enabled=true' | tee -a /usr/whatap/infra/conf/whatap.conf
'''
```

```
'''bash title='AIX'
stopsrc -s whatap-infra
startsrc -s whatap-infra
'''
```

```
'''bash title='Oracle Solaris'
/etc/init.d/whatap-infra restart
'''
```

```
* **Windows**
```

- [Add option command](#) Restart agent

```
'''bash
if (-not (Select-String -Path "C:\Program Files\WhatapInfra\whatap.conf" -Pattern "^hostinfov2.enabled" -Quiet)) { Add-Content -Path "C:\Program Files\WhatapInfra\whatap.conf" -Value "hostinfov2.enabled=true" }
'''
```

```
'''bash
Restart-Service -Name "Whatap Infra"
'''
```

Basic screen guide

The **Server Inventory** menu provides key configuration items of the server in a list form. You can manage essential information required for server operation and management, as well as custom items that are automatically collected. This allows server administrators to monitor and operate their infrastructure more systematically and efficiently. Use the filters and sorting feature to easily find the information you need.

Server Inventory									
2025/09/23 📅 Total 5 cases 🔍 + There is no filter applied.									
OS Type ▼ > Model ▼ > Serial ▼ 📖 Guide to feature activation 🕒 Query change history ⚙️ Column settings 📄 Download CSV 📤 CSV upload ✎ Edit									
History	Connection	Host Name	Alias Name	Server Type	External IP	Default Group	OS Type	OS	
		ip-172-31-36-165.ap-north...		Virtual	172.31.36.165		Linux/A...	Amazon 2 5.10.237-230.94	
		rocky9arm	리눅스_운영_ARM_01	Virtual	192.168.12.128	production	Linux/R...	Rocky 9.4 5.14.0-427.28.1.ε	
		rocky9-vm	리눅스_운영_서버_01	Virtual	192.168.1.84	production	Linux/R...	Rocky 9.4 5.14.0-427.31.1.ε	
		nmstest		Virtual	192.168.1.76		Linux\U...	Ubuntu 22.04 5.15.0-144-ge	
		KHUNLEE-LAPTOP	테스트용_windows_랩탑	Physical	192.168.10.117	test	Windows	Windows 11 2009 (Version	

Filtering

Click the filter input field, select the desired **filter key** and **condition**, enter values, and then click **Apply**. In the server inventory list area, you can view the filtered list.

- **Filter key:** Filters can be set based on the default collection columns provided in the **Server Inventory** menu and user-defined custom columns.
 - Default collection columns: Items automatically collected from **Server Inventory** are **server type**, **default group**, **OS type**, **OS version**, **model**, and **serial** columns.
 - Custom columns: They are the items added and managed on a project basis.
- **Condition:** Filter conditions can be set as numeric conditions such as $\geq$, $\leq$, $>$, and $<$, and string conditions such as Equal, Unequal, Includes, and Excludes.

Align

Sorting options allow up to 3 levels of multi-sorting. You can sort by any desired criteria including custom columns. The default settings for sorting are **OS type** > **Model** > **Serial**.

 Sorting is reset when moving and refreshing the menu.

History Lookup

You can track server changes by checking past data in the server inventory, viewing change history, and comparing by date.

Past Data Lookup

You can load snapshot data saved at midnight each day to check past data or the change history of specific fields.

1. Click the YYYY/MM/DD  button.
2. Select the past date you want to query, then click the **OK** button to view the data.

 When viewing past data, the **Edit**, **CSV Upload**, and **Column Settings** functions are not available. The Connection field in the table is automatically hidden.

Change History Lookup

You can check how the value of a specific field in the selected server changed by date during the query period, as well as view creation and deletion history of that field.

1. In **Server Inventory**, click the  **Change History Lookup** button.
2. In the **Inventory Change History Search** window, select the date, server, and column name to query.
3. Review the change history by date according to the selected conditions.


Date Comparison

You can compare two dates of a selected server to check data changes and whether fields were created or deleted.

1. Click the **History**  icon of the server to compare.
2. In **History Lookup**, select the **Base Date** and **Comparison Date** to compare the data.

Column Settings

Click  **Column settings**. Select columns and add custom columns as shown below, and then select **Confirm** at the bottom to save the column settings.



• Select Column

Users can select the columns to display in the table. If the **Initialized** button is selected, it returns to the default selection state.

ⓘ The selected [columns](#) are saved as browser cookie values and remain in effect even if the page is refreshed. If cookies are deleted or a setting error occurs for any other reason, the default selection state is reset.

• Adding columns

Users can configure custom columns (**custom**) by selecting the **+ Add a column** button and adding their own desired items in addition to the default items. Add columns to your taste and then monitor the server status and performance in more detail.

CSV downloading and uploading

CSV downloading

By selecting the **Download** button, you can download the collected server inventory in **CSV** file format.

CSV uploading

When you need to enter server configuration information in bulk, manually entering individual pieces of data can be time-consuming. The **CSV file** uploading feature allows efficient handling of bulk data entry. Try modifying the downloaded **CSV** file and using the upload feature.

1. Click the **CSV Upload** button, then select the **CSV** file to upload.
2. Through the **CSV verification** window, check the verified results for the uploaded data.
 - You can check the numbers of changes and errors.

CSV verification														
Verification results														
Number of changes 1														
Error 0														
No.	id													
1	1906460792	virtual-java-a...	Virtual	10.21.1.26	APM DEMO	Linux/Ubuntu	Ubuntu 20.04...	Amazon EC2	ec2954f4-33...	-				
2	-1816960390	mysql-official...	Virtual	10.21.3.69	DPM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec21a8f7-07...	-				
3	-2142672185	mysql-official...	Virtual	10.21.1.247	DPM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec23a77b-e0...	-				
4	-780269224	mysql-official...	Virtual	10.21.1.223	DPM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec268308-e2...	-				
5	663452279	go-official-de...	Virtual	10.21.1.118	APM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec2c581e-d0f...	-				
6	1786631657	qa-php-demo	Virtual	10.21.3.73	APM QA	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec23280f-79...	-				
7	-1413588424	php-official-d...	Virtual	10.21.3.223	APM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec27d8ad-99...	-				
8	-978189497	ip-10-21-1-200	Virtual	10.21.1.200		Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec2618e8-ff...	-				
9	-846195603	rum-official-d...	Virtual	10.21.1.151	BROWSER DE...	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec2a72fe-f24...	-				
10	-868519645	wpm-official...	Virtual	10.21.1.181	WPM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec24c670-e5...	-				
11	-1393962412	qa-node-demo	Virtual	10.21.1.241	APM QA	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec230c90-1a...	-				
12	-1341918756	node-official...	Virtual	10.21.1.98	APM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec240246-05...	-				
13	332299567	python-demo...	Virtual	10.21.1.254	APM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Amazon EC2	ec2a269c-4d...	-				
14	-2078406413	python-demo...	Virtual	10.21.1.43	APM DEMO	Linux/Ubuntu	Ubuntu 22.04...	Xen (HVM do...	ec2e113a-c5...	-				
15	-1187373871	EC2AMAZ-J7...	Virtual	10.21.1.227		Windows	Microsoft Wi...	Amazon EC2	ec20420-90...	-				

Items not reflected

- Columns items that are not in existing inventory columns are not reflected.

- Columns can be added in the inventory list 'Add Column'. Server deletions and additions are not reflected.

Apply

- If you click an error, the item in question within the change data list on the right appears, and the erroneous item is highlighted.

CSV verification										
Verification results										
Number of changes 0										
Error 2										
No 8 / id This is a duplicate ID. (Found duplicate pk0.)										
No 16 / id This is a duplicate ID. (Found duplicate pk0.)										
1	1906460792	virtual-java-8...	Virtual	10.21.1.26	APM DEMO	Linux/Ubuntu	Ubuntu 20.04	Amazon EC2	ec295af4-53...	-
2	18161892090	mysql-official...	Virtual	10.21.3.69	DPM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec21a8f7-67...	-
3	2142672185	mysql-official...	Virtual	10.21.1.247	DPM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec23673b-e0...	-
4	780269224	mysql-official...	Virtual	10.21.1.223	DPM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec268326-c2...	-
5	663432279	go-official-de...	Virtual	10.21.1.118	APM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec3c581c-d5f...	-
6	1786431657	qa-php-demo	Virtual	10.21.3.73	APM QA	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec22280f-79...	-
7	1413588424	php-official-d...	Virtual	10.21.3.223	APM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec27d8ed-99...	-
8	978189497	ip-10-21-1-200	Virtual	10.21.1.200		Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec2b18e8-f8f...	-
9	846956083	rum-official-d...	Virtual	10.21.1.151	BROWSER DE...	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec2a73fe-f24...	-
10	868519645	wpm-official-...	Virtual	10.21.1.181	WPM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec24c670-e5...	-
11	1393962412	qa-node-demo	Virtual	10.21.1.241	APM QA	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec230c90-1a...	-
12	1341918756	node-official...	Virtual	10.21.1.98	APM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec240246-05...	-
13	332299567	python-demo...	Virtual	10.21.1.254	APM DEMO	Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec2a29c4-db...	-
14	2078406413	python-demo...	Virtual	10.21.1.43	APM DEMO	Linux/Ubuntu	Ubuntu 22.04	Xen (HYM do...	ec2e113a-c5...	-
15	1187373871	EC2AMAZ-27...	Virtual	10.21.1.227	Windows	Microsoft Wi...	Amazon EC2	ec209420-90...	-	-
16	978189497	ip-10-21-1-200	Virtual	10.21.1.200		Linux/Ubuntu	Ubuntu 22.04	Amazon EC2	ec2b18e8-f8f...	-

- If there is no issue found after verifying the data, click **Apply** to apply the server configuration information from the selected **CSV** file.

- *When uploading a *CSV*, only the changes to **Server alias**, **Default Group** and custom columns (**custom**) are applied.
- When uploading **CSV**, the columns that were not previously registered are not reflected.

Modification

Select the [🔗 Edit](#) button to edit the server inventory list.

- You can configure the settings for the **Server alias** and **Default Group** columns.
- You can set the contents of custom columns (**custom**) added by users.

Server List

Home > Select Project > **Server List** > **Server List**

Select a project in the initial screen of the WhaTap monitoring service, and then select **Server List** > **Server List** under Project Menu. The **Server List** displays the overall status of the collection server and provides the comparison functions between server management functions and between servers.

Server List

|| LIVE 08:25:35

All 11

Inactive 3

Warning 7

Normal 1

Agent name, IP

+ There is no filter applied.

Column settings

Refresh

Transfer

CSV

Compare

Management

Server information		CPU	Memory	Disk		Network	ETC	
Status	Detail	Agent name	CPU used %	Memory used %	Disk used % (max)	Disk I/O % (max)	Network traffic (max)	PRIVATE_IP
		25fbd76273a2	> 32.87%	32.09% (4.9G)	79.05%	0.04%	lo 0	127.0.0.1
		64311a6a8bb1	> 32.8%	32.09% (4.9G)	79.05%	0.04%	lo 0	127.0.0.1
		9ae35ad8c03b	> 32.59%	32.09% (4.9G)	79.05%	0.04%	lo 0	127.0.0.1
		9e6a0322b3a2	> 32.62%	32.09% (4.9G)	79.05%	0.04%	lo 0	127.0.0.1
		d374f2cebf49	> 32.69%	32.09% (4.9G)	79.05%	0.04%	lo 0	127.0.0.1
		qa-compose-demo	> 32.87%	32.09% (4.9G)	79.05%	0.04%	ens5 944K	10.21.11.118 / 1...
		qa-region	> 20.15%	86.7% (26.6G)	85.47%	2.24%	lo 8.4M	10.21.11.187 / 1...
		qa-node-demo	> 11.47%	16.23% (1.2G)	47.59%	0.08%	lo 40.7M	10.21.1.241 / 12...
			> 0%	0%	0%	0%	0	10.21.1.241

1 Data refreshing and server status

Data update

The data on the server list page is automatically updated every 60 seconds.

- **|| Pause** : You can stop automatic update of the data.

-  **Playback** : You can restart the auto update you have stopped.

 For more information on how to use the time selector.

Server Status

It displays the current status of the collection server by status. The status is updated every 1 minute.

-  **Inactive**: The agent cannot collect data from the server.
-  **Fatal**: Any fatal events occur in the collection server.
-  **Warning**: Any warning events occur in the collection server.
-  **Normal**: The agent is running normally in the collection server.
-  **Paused**: The agent temporarily paused collecting data.

1 Searching

You can search for agent names and IPs by using the  **search bar** at the upper right.

2 Filter

You can specify the filter by directly entering a value in the input field or by adding a filter. The filter tags are entered in the following order: **filter key** , **condition** , and **search value** .

Enter the filter key, condition, and search value.

- When you enter the **filter key** , it provides the following options: Agent Name, CPU Used %, Memory Used %, Disk Used %, Disk I/O %, Network Traffic, and private_ip.
- When entering a **condition** , the options such as **>** , **<** , **<=** , **>=** , **==** , and **!=** are provided.
- Upon entry of the **search value** for matched search (**>** , **<** , **<=** , **>=** , **==**), it is highlighted in **blue**, In case of excluded search (**!=**), it is highlighted in **red**.
- Upon entry of the **search value** , you can search using the case-sensitive option.

- If the filter tags exceed 2 lines, you can fold it by selecting the  **Fold** icon.
- After entering a filter tag, you can use the Shift and Tab keys simultaneously in the input field to go to the previous filter tag.
- After entering a filter tag, you can use the Tab key in the input field to go to the next filter tag.

Deleting filter tags

- A tag can be deleted with the Backspace key.
- You can delete a tag by selecting the  icon for it.
- You can delete all tags by selecting the Delete All  icon in the input field.

Filter favorites

You can select the  **Favorites** icon on the right of the input window to **add**, **delete**, and select desired filter search conditions in the favorites.

Exceptional cases for application of filters

- In the tag with a numeric index (**search key** that ends with **.n**), only numbers can be input for the **search value** .
- Duplicate 'search key' and 'search value' cannot be entered.
- Searching is not possible when there is a tag that does not have either **condition key** or **search value** . Invalid tags are displayed in **gray**.

3 Server Status

It displays the status of all servers to which agents are assigned. If an agent is selected in the list, the **Server Detail** page appears. For more information about **Server Detail**, see [the following](#).

File system information

[File system information](#) can be checked by clicking the  icon on the left of the agent name to view in detail.

Column Settings

- Column settings

3 In the menu at the top, if you click **Column settings**, you can select columns to displays in the table as follows.

CPU	Disk	Memory	Network	Server information	ETC
☐ CPU idle	☐ Disk block size	☐ Memory buffers	☐ Network dropped in	☐ CPU core	☐ Elapsed time after agent creation
☐ CPU I/O wait	☐ Disk free %	☐ Memory cached	☐ Network dropped out	☐ OID	☒ IP
☐ CPU IRQ utilization	☐ Disk free space	☐ Memory free	☐ Network error in	☐ Host name	☐ Open file descriptors
☐ CPU load average 1min	☐ Disk inode total	☐ Memory pagefaults	☐ Network error out	☐ OS architecture	☐ Public IP
☐ CPU load average 15mins	☐ Disk inode used	☐ Memory shared	☐ Network IP	☐ OS CPU vendor	☐ Server time
☐ CPU load average 5mins	☒ Disk I/O % (max)	☐ Memory swap total	☐ Network packet in	☐ OS memory	☐ Tags
☐ CPU nice usage rate	☐ Disk queue length	☒ Memory used %	☒ Network traffic (max)	☐ OS Name	☐ Uptime
☐ CPU SOFT IRQ	☐ Disk read bps		☐ Network traffic in	☐ OS release	☐ Agent version
☐ CPU steal utilization	☐ Disk read IOPS		☐ Network traffic out		
☐ CPU system utilization	☐ Disk total space				
☒ CPU used %	☐ Disk used %				
☐ CPU user utilization	☒ Disk used % (max)				
	☐ Disk used space				
	☐ Disk write bps				
	☐ Disk write IOPS				

① The selected **Columns** are stored as the browser cookie values so that the status is maintained after refreshing the page. If an error occurs between settings due to cookie deletion or other reasons, it is initialized to the full selection state.

- Column Width Setting

You can set the column width by dragging the edge of the column title.

- Align Columns

Sort Columns By clicking the column title, you can sort in ascending or descending order.

- ⓘ The settings for **Columns Width** and **Columns Align** are stored in the browser cookies and the status is maintained after refreshing the page. If an error occurs between settings due to cookie deletion or other reasons, it is initialized to the unset state.

Preset management

In **Server List > Server List**, if you click  **Column settings**, you can manage and share the presets.

The selected column combination can be saved as a preset, and shared with other users who access the same project. Recently used presets are stored in the personal browser and automatically loaded when connection is made again. If the recent preset is deleted, it is reset to the default preset (**Default**). For information on how to set presets, see [the following](#).

- The widths and alignments for each column are saved in the user's browser and automatically applied when the menu is re-accessed. These settings are applied only to the user and cannot be shared with other users.

- ⚠ **Only the users with the project edit role can save and delete presets.** For more information about member roles, see [the following](#).

Column settings

Refresh

Filter

Transfer

CSV

Compare

Management

Select all Deselect all Initialize(default)

Default

CPU	Disk	Memory	Network	Server information	ETC
☐ CPU idle	☐ Disk block size	☐ Memory buffers	☐ Network dropped in	☐ CPU core	☐ Elapsed time after agent creation
☐ CPU I/O wait	☐ Disk free %	☐ Memory cached	☐ Network dropped out	☐ OID	☒ IP
☐ CPU IRQ utilization	☐ Disk free space	☐ Memory free	☐ Network error in	☐ Host name	☐ Open file descriptors
☐ CPU load average 1min	☐ Disk inode total	☐ Memory pagefaults	☐ Network error out	☐ OS architecture	☐ Public IP
☐ CPU load average 15mins	☐ Disk inode used	☐ Memory shared	☐ Network IP	☐ OS CPU vendor	☐ Server time
☐ CPU load average 5mins	☐ Disk inode used %	☐ Memory swap total	☐ Network packet in	☐ OS memory	☐ Tags
☐ CPU nice usage rate	☐ Disk I/O %	☐ Memory swap used	☐ Network packet out	☐ OS Name	☐ Uptime
☐ CPU SOFT IRQ	☒ Disk I/O % (max)	☐ Memory total	☐ Network traffic	☐ OS release	☐ Agent version
☐ CPU steal utilization	☐ Disk queue length	☒ Memory used %	☒ Network traffic (max)		
☐ CPU system utilization	☐ Disk read bps		☐ Network traffic in		
☒ CPU used %	☐ Disk read IOPS		☐ Network traffic out		
☐ CPU user utilization	☐ Disk total space				
	☐ Disk used %				
	☒ Disk used % (max)				
	☐ Disk used space				
	☐ Disk write bps				
	☐ Disk write IOPS				

Refresh

If you click  **Refresh**, the data is updated based on the current time.

Agent Transfer

By selecting  **Transfer**, you can transfer the agent to another project.

Agent transfer
✕

i

You can move agents installed in a server project to another.

Select the agent you want to move between project A and B, and click the "Arrow button" in the middle.

✕

Server Monitoring Demo (29763)
 ▼

↻

Selected agent: 0

☐ All

☐ rum-official-demo

☐ demo-k8s-master

☐ demo-k8s-worker-01

☐ mysql-official-demo-02

☐ node-official-demo

☐ virtual-java-agent

☐ go-official-demo

☐ url-official-demo

☐ python-demo-target

☐ python-demo-attack

☐ demo-k8s-worker-03

☐ demo-k8s-worker-04

B. Select project
 🔍

All available agents

All available agents

Demo agent

Server Monitoring (29763)

Spring Web MVC (29763)

Spring Web Services (29763)

Spring Web Services (29763)

Spring Web Services (29763)

Spring Web Services (29763)

Spring Web Services (29763)

>
<

No Data

A crash may occur if you move an agent again within 10 minutes of running. (Asynchronous Processing)

⇌
Save


1. Select a project to transfer on the right screen.
2. Select an agent to transfer on the left screen.
3. By clicking > [Right arrow](#), you can transfer to the target project.
4. If you click **Save**, the transfer is completed.

- !
 - If you move the same agent again within 10 minutes just after execution, it may cause a conflict.
 - The collected data is not transferred. Do not delete the existing projects for a while to retrieve the


past data.

Comparison between servers

If you click  **Compare**, the cross-server comparison chart is provided as follows.



1. Select  **Compare**.

2. Like **6**, check whether the **Compare** column has been added on the left of the server list.
3. In the **Compare** column, select a server for comparison..
4. To enable the comparison function, you can see the chart to compare in the **7** area.

- The comparison servers can be selected up to 5.
- Unique colors are assigned to each server. The unique colors are used to distinguish servers in **7 Compare**.
- If you select **Check period**, the past data can be retrieved.

Monitoring Management

This function manages servers through agent commands. Click **Management** and then select a server to apply the management features from the list. On the right, **Pause**, **Restart**, and **Delete** buttons are activated.

- **|| Pause:** Data collection for the agent is paused. It is displayed as the paused state in the server list.
- **▶ Restart:** The agent is restarted.
- **🗑 Exclude:** The agent operation is stopped. It is excluded from the server list and the server status cannot be checked.

Inactive Agent Retention Period

The inactive agent retention period is the period before an agent that has stopped collecting data is automatically deleted from the server list. It is set next to the status summary (Total, Inactive, Critical, Normal) at the top of the server list screen. If the inactive state lasts for the configured period, the agent is automatically deleted from the server list. Setting a shorter retention period reduces the time that unused servers remain in the list.

This setting is applied per project, and the available retention periods are as follows.

Option	Description
1 day	Automatically deleted when the inactive state lasts for 1 day


Option	Description
3 days	Automatically deleted when the inactive state lasts for 3 days
7 days	Automatically deleted when the inactive state lasts for 7 days
1 month	Automatically deleted when the inactive state lasts for 1 month
Do not delete	Keeps inactive agents in the list without automatic deletion (default)

- Only users with project **Edit** permission can set the inactive agent retention period.
- If a value not in the options above is already set, that value is shown as a temporary option.

Server Details

Home > Select Project > **Server List** > **Server Detail**

Select a project in the initial screen of the WhaTap monitoring service, and then select **Server Detail** under **Server List**. **Server Detail** provides detailed information on the server resource status and running processes.



Options at the top

In the ① area at the top, you can set the data for [data update time](#) and [server](#). You can see the [Summary Info](#) of the server, along with the [Data Update Time](#) and [Server](#) options.

Data update time

The data on the server detail page is automatically updated every 5 seconds.

-  **Pause** : You can stop automatic update of the data.
-  **Playback** : You can restart the auto update you have stopped.

 For more information on how to use the time selector.

Select Server

The **Select server** option allows you to set the server to view. It also set to perform **Displays inactive servers** during a view. If you select a server to view from the list, you can see the server resources and detailed information of running processes in the **2** area.

2 Detail

Server resources

In the **2** area, you can see the resources of the server through the [CPU Usage](#), [Memory Usage](#), [Network](#), and [Disk I/O](#) charts.

-  : You can specify the maximum value of the Y-axis in the chart configuration.
-  : You can save the chart records in a CSV file.
-  : The chart data is saved in an image data.
-  : You can see the charts for CPU utilization, memory utilization, network, and disk I/O.
 - CPU Utilization
[CPU Usage](#), [CPU Idle](#), [CPU Nice](#), [CPU I/O Wait](#), [CPU Steal](#), [CPU IRQ](#), [CPU Soft IRQ](#), [CPU Load](#)
 - Memory utilization
[Memory Usage](#), [Memory Available](#), [Memory SReclaimable](#), [Memory SUNreclaim](#), [Memory Slab](#), [Memory Swap Used](#), [Memory Page Faults](#)
 - Network
[Traffic in/out](#), [Packet in/out](#), [Error in/out](#), [Dropped in/out](#)

- Disk I/O

[Disk I/O](#), [IOPS read/write](#), [Disk Bph Read/Write](#), [Used space](#), [Queue length](#), [Inode Used](#), [Free space](#)

Process List

Align

- You can search by setting a sorting in descending or ascending order according to the selected metric at the top of **Process** list.
- According to the metric selected at the top of **Process** list, you can check by setting the top [10](#), [20](#), [30](#), and **All** processes.

Grouping

A process group is a set of multiple processes with the same behaviors in an instance. You can add process groups by selecting the **Group by** button. For more information about process grouping, see [the following](#).

Server information panel

You can see the following in the **Server information** panel. In case of *** Agent Name*** and [Memo](#), you can edit only in **Server information** panel.

-  **Server Alert List**

At the top of **Server information** panel, if you select  **Server Alert List**, the alerts for the current server appear.

- [Agent Name / ID](#)

You can check the agent name (alias for ONAME) and ID (OID) of the selected server. You can modify the agent name by selecting the  button.

 ONAME is set to the host name by default.

- [OS](#)

You can see the OS information of the selected server.

- [IP Address](#)

You can see the IP address of the selected server.

- [Agent](#)

You can see the agent version of the selected server.

- [Memo](#)

By checking the written memo or selecting [✎](#), you can modify or add any memo.

Adding The Process Group

What is the process group?

A process group is a group of multiple processes that perform the same operation in a single instance. Use the process groups to monitor or receive notifications for each group rather than individual processes.


You can see at a glance the resource usages for each divided process in Process Group. There is no need to calculate the resource usages for individual processes.

You can set to receive notifications when the usage sum of resources for the process group exceeds a specific threshold. For more information about the notification settings, see [the following](#).

Adding the process group

ⓘ The project **Edit** or **Alert settings** role is required. For more information about the roles, see [the following](#).

1. Go to Home > Select Project > **Server List** > **Server Detail** and then click **Group by** in the process list.



- If the process group setting window appears, in **Pattern**, enter the process name pattern to be included in the group including asterisks (*) before and after.

e.g. whatap → *whatap*

Process Group

Pattern

ex) *whatap*

×

🔍

Pattern Alias

Apply

ℹ

The set of process pattern name is replaced with the alias.

After entry, select  on the right. You can view the list of processes that corresponds to the pattern. The process name is based on the process name displayed on the screen.

3. Enter the process group name in **Pattern Alias** and then select **Apply** on the right. Once the grouping has been set, grouping is performed for **all servers** of the **set project**.

Process group example

- **** Application of process grouping****

The following screen displays an example where `*tomcat*` is entered in **Pattern**, and `test_group` is set in **Pattern Alias**.



Process Group

Pattern


`*tomcat*` 

Pattern Alias

`test_group` **Apply**

 The set of process pattern name is replaced with the alias.

- **Process list before process grouping**



The following lists the processes to be grouped by [Pattern](#).

- e.g. javatomcat_left/bin/test.jarcom.virtual.App
- e.g. javatomcat_right/bin/test.jarcom.virtual.App

• Checking the process group registration

You can check that a new process group, test_group has been registered in the process list as follows.

Checking the process group registration

- ! • After applying the grouping, you can check the sum of the data corresponding to each process in the process group.
- Check the data for each process to see the data before grouping.

WhaTap GPU Monitoring

WhaTap GPU Monitoring provides integrated analysis of GPU resource status and utilization across server and Kubernetes environments. While GPU is one of many resources in a server, it is a critical resource that directly impacts the performance of AI/ML, LLM, and HPC workloads — and comes with significantly higher costs compared to CPU or memory.

Beyond simply checking **whether a GPU is installed and running**, you need visibility into how much it is being utilized, which workloads are occupying it, whether any anomalies exist, and whether resource allocation is appropriate.

WhaTap GPU Monitoring Scope

! WhaTap GPU Monitoring supported environments

- GPU monitoring in Server environments
- GPU monitoring in Kubernetes environments

Server Environment

- GPU server monitoring periodically collects and stores key metrics such as GPU utilization, memory usage, temperature, power, clock speed, errors, and PCIe/NVLink communication status, enabling comprehensive assessment of GPU device health from a server infrastructure perspective.
- Process-level GPU occupancy data is also provided, so you can see which workloads are using which GPUs and how much. This goes beyond simple device health checks to help you analyze how GPU resources are actually being used within the server.
- Monitoring enables early detection of signs such as overheating, throttling, power limits, abnormal errors, and PCIe/NVLink communication issues — helping you prevent failures and respond quickly. This directly contributes to service stability and improved operational efficiency.

Kubernetes Environment

- In Kubernetes environments, GPUs are primarily allocated to AI training/inference workloads, batch jobs, and high-performance computing applications. Viewing only node-level GPU status is often insufficient for understanding real-world operational conditions.
- WhaTap Kubernetes GPU monitoring provides visibility into GPU resource usage at the cluster, node, and pod/container levels. This helps you understand which workloads are occupying GPUs, whether usage is concentrated on specific nodes, and how efficiently GPUs are being used relative to requests and allocations.
- GPU monitoring from a Kubernetes perspective is also useful for analyzing scheduling efficiency, resource imbalance, and over/under-allocation. Ultimately, it improves the stability of GPU-based workload operations and enhances overall cluster resource utilization.

Enterprise Environment

- Enterprise environments typically run many GPU servers and Kubernetes clusters together. Monitoring individual servers or Kubernetes environments in isolation makes it difficult to get a comprehensive view of overall resource status, so an integrated management approach is required.
- Organizations operating hundreds to thousands of GPUs want continuous visibility into how effectively GPU resources allocated by team or workload are being used. This data is essential for making informed decisions about idle resource reallocation, capacity expansion, and procurement planning.
- By unifying multiple Kubernetes-based GPU environments and bare-metal GPU server environments under a single monitoring framework, you can manage distributed GPU resources from a single perspective. This delivers enterprise-grade GPU monitoring that provides a consolidated view of GPU resources across the organization and connects to Capacity Planning.

Value of WhaTap GPU Monitoring

GPU Resource Visibility

Gain a multi-dimensional view of GPU usage down to the server, cluster, node, pod, and process level.

Early Anomaly Detection and Rapid Response

Quickly identify signs such as temperature spikes, power limits, clock degradation, errors, and communication issues before they escalate into failures.

GPU Utilization Optimization

Identify idle GPUs, uneven usage, and underutilized resources to enable reallocation and operational optimization.

Improved Operational Efficiency

Manage server and Kubernetes environments in an integrated manner — rather than separately — to reduce operational complexity.

Informed Expansion and Investment Decisions

Use real usage data to evaluate the need for GPU expansion, and leverage it as supporting evidence for budget and procurement planning.

Capacity Planning Support

Connect short-term incident response to medium- and long-term resource demand forecasting and operational strategy planning.

GPU Device Anomaly Detection

Methods for detecting anomalies in GPU devices and workloads.

1. Xid-Based Anomaly Detection

NVIDIA GPUs record Xid events in logs when a failure or abnormal condition occurs. Detecting the presence of Xid keywords and codes in system logs or driver logs enables relatively fast identification of GPU device anomalies.

For example, specific Xid events can indicate GPU computation errors, memory access issues, driver/hardware faults, or reset events — making Xid events a valuable failure signal in production environments.

2. Status Metric-Based Anomaly Detection

GPU anomalies are not always reflected in error logs alone. In production environments, monitoring the change patterns of key status metrics — such as GPU utilization, memory usage, temperature, power, clock speed, and PCIe/NVLink communication status — enables earlier anomaly detection.

However, it is important to **interpret these metrics not solely by absolute values, but in conjunction with workload characteristics, task type, time-of-day patterns, and typical baselines**. For example, training, inference, batch processing, and data preprocessing workloads each have different normal GPU usage patterns, so the same value can mean different things depending on the operational context.

From this perspective, the following pattern changes are key areas to monitor:

- GPU utilization is excessively high compared to typical workload patterns, or allocated resources remain persistently underutilized
- Memory occupancy remains abnormally high after a task completes, or memory usage spikes sharply from a certain point
- Temperature continuously rises beyond the normal operating range, or a sustained high-temperature condition persists for an extended period
- Power consumption hovers near the limit for a prolonged period, suggesting performance throttling due to power capping
- Clock speed remains lower than expected relative to the load level, suggesting throttling or abnormal control behavior
- PCIe/NVLink communication volume shows abnormal increases or decreases relative to workload characteristics, or error indicators are observed on InfiniBand

Interpreting status metrics as change patterns within operational context — rather than as individual values — allows you to more accurately identify issues such as performance degradation, overheating, throttling, power limits, communication bottlenecks, and abnormal behavior.

3. Process/Workload-Based Anomaly Detection

Viewing workload information such as processes, pods, and containers occupying the GPU alongside device

metrics helps clarify anomalies that would be difficult to identify from device metrics alone. However, this also requires considering workload type, execution stage, repetition cycle, and typical occupancy patterns — not just whether something is occupying the GPU.

Since training, inference, batch processing, and data preprocessing workloads use GPUs differently, the criteria for normal versus abnormal can vary depending on the type of task — even when the occupancy state appears the same.

From this perspective, the following patterns are key areas to monitor:

- A specific process or workload occupies GPU memory in excess of what its task characteristics would suggest, without releasing it within the expected timeframe
- Memory occupancy is sustained for a long time while GPU utilization is low, suggesting task stagnation, abnormal waiting, or potential memory leaks
- A specific task repeatedly restarts or fails while inefficiently occupying GPU resources
- Workloads are persistently concentrated on a subset of GPUs, causing usage imbalance or localized overheating
- Compared to similar workloads, a specific process shows excessively long execution time, high memory occupancy, or abnormal occupancy patterns
- After a pod or container is rescheduled, GPU occupancy patterns change drastically, suggesting scheduling inefficiency or resource imbalance

These patterns can help identify not only device-level issues but also application or workload operational problems.

GPU Dashboard

The GPU dashboard is an AI infrastructure operations dashboard that provides an integrated view of GPU resource composition, utilization efficiency, and issue history on a single screen. When GPU servers grow to many units, checking each server one by one is difficult. The GPU dashboard collects all GPU resources by model and by server and displays them, helping you quickly grasp the utilization status and anomalies.

! Supported environment

For an overview of GPU monitoring in server environments, see the [WhaTap GPU Monitoring](#) document.

Basic screen guide

The GPU dashboard consists of two tabs.

- **GPU operation status:** Collects and displays the GPU resource inventory, utilization efficiency, power consumption, and issues by widget.
- **GPU resource board:** Provides an overview of the allocation and utilization status of all GPUs as a summary bar and map.

Option bar

You can adjust the query conditions in the option bar at the top of the screen.

- **Time selection:** Specify the time range to query.
- **Filter:** Specify filters to query only the GPUs that match conditions such as a specific model or server.
- **View toggle:** Switch the display format of the widgets.
- **Preset:** Save frequently used combinations of filters and Top 5 metric selections as a preset and load them again.

GPU operation status tab

The GPU operation status tab provides the following widgets.

GPU inventory (by model)

Aggregates the GPUs you have by model and displays them as a donut. When you click a model slice of the donut or a legend item, the [GPU Inventory](#) with the filter applied for that model opens in a new tab.

GPU power consumption trend

Displays the GPU power consumption trend over time.

 The displayed power consumption is an **estimate** calculated based on GPU metrics and may differ from actual power meter measurements.

Server severity distribution by GPU model

Displays, for each GPU model, the ongoing event status and representative severity distribution of the servers equipped with that model, as a stacked bar. Occurred x / Total y servers indicates the number of servers, among all servers equipped with the GPU model, that currently have at least one ongoing event. Each server's severity is classified by the highest severity among its ongoing events, and the server counts are shown per Critical, Warning, Info, and Normal. Total n events indicates the total number of events currently ongoing across the servers where events occurred. When you click an item, you can view the list of ongoing events for the servers equipped with that GPU model in a drawer.

Top 5 Server / Top 5 GPU

Displays the top 5 servers (Top 5 Server) and the top 5 GPUs (Top 5 GPU) by a specified metric in separate panels. The default reference metric is GPU utilization for servers and power limit violation rate for GPUs. When you click an item, a drawer showing the recent trend of that target opens. You can change the reference metric in the option bar or a preset.

Xid

Displays the history of Xid events that occurred on GPUs. For how to detect anomalies using Xid events, see the [WhaTap GPU Monitoring](#) document.

GPU resource board tab

The GPU resource board tab provides an overview of the allocation and utilization status of all GPUs as a summary bar and map. You can see at a glance how much each GPU is allocated and used, and which GPUs are idle or close to running out of memory.

Summary bar

Summarizes the allocation and utilization status of all GPUs by card. Each card shows the number of physical GPUs (Physical) and MIG instances together.

Item	Description
Total GPU	Total number of GPUs
Allocated GPU	Number of allocated GPUs and the allocation rate
Idle GPU	Number of idle GPUs that are allocated but not used
Active GPU	Number of GPUs in use
Effective GPU	Number of GPUs used efficiently

The summary bar card names are displayed in English (Total GPU, Allocated GPU, Idle GPU, Active GPU, Effective GPU). The idle/active/effective classification is based on GPU utilization: less than 1% is Idle, 1% or more and less than 50% is Active, and 50% or more is Effective.

GPU Resource Map

A scatter plot that arranges all GPUs by utilization. It shows the distribution by dividing each GPU into the Idle, Active, and Effective ranges, and separately marks GPUs that are close to running out of memory due to high memory usage (FB usage 95% or more). When you select a specific area on the map, a **GPU Timeline** drawer opens where you can check the recent trend of those GPUs.

Utilization status by GPU model

Aggregates GPUs by model and shows the ratio of the Idle, Active, Effective, and Unallocated ranges. You can compare how the allocation and utilization status is distributed by model.

GPU utilization status by server

Displays the GPUs held by each server as a slot grid. Each slot represents one physical GPU, and a GPU split by MIG is displayed by dividing the inside of the slot into instances.

! GPU utilization status by server is displayed automatically in environments with 20 or fewer GPU servers.

GPU Inventory

The **GPU Inventory** feature allows you to systematically manage the specifications and status of GPUs installed in servers.

! Supported Agent Version

Only supported on Linux. Requires agent version 2.9.3 or higher.

Basic Screen Overview

GPU Inventory monitors a complete list of GPU and MIG (Multi-Instance GPU) resources assigned to the system. You can view detailed information for each GPU instance.

- You can search by keyword using the search bar.
- Click on a column header in the results to toggle between ascending and descending sort order.

🔍 Keyword Search Filter

By using the filter search function, you can quickly check only the GPU instances that match your desired conditions.

Enter the filter in the 🔍 search box and press Enter to search.

⚙️ Column Settings

Click the **Column Settings** button at the top right of the screen to select the columns you want to display. In the Column Settings window, choose the desired columns and optionally add custom columns. Click **Confirm** to save your settings. You can rearrange column order by dragging the items in the list.

- **Column Selection**

You can choose which columns to display in the table.

Click **Reset** to return to the default column selection.

- ⓘ Your selected columns are stored in browser cookies, so they persist even after refreshing the page. If cookies are deleted or a settings error occurs, the selection will be reset to the default columns.

- **Add Column**

By clicking the + **Add Column** button, you can add custom columns (`custom`) beyond the default ones. Use this to track specific information that matters to your workflow and monitor server status and performance in more detail.

CSV Download & Upload

1. Click the  **Download CSV** button to download the current GPU inventory data.
2. Open the downloaded file and modify the column values as needed.
3. Click the  **Upload CSV** button to upload the edited file and update the values in bulk.
 - You can check the number of updated items, errors, and skipped entries.
4. Click the **Apply** button to reflect the changes in the GPU inventory list.

ⓘ **Items that are not applied**

Any column not included in the existing inventory will be ignored.

- To add new columns, use the + **Add Column** button in the inventory list. GPU additions or deletions will not be reflected.

Edit

Click the  **Edit** button to update the GPU inventory list.

- You can edit values in the **Server Alias** and **Default Group** columns.
- You can also set values for custom columns (`custom`) added by the user.

GPU Performance Summary

Monitor and manage the status and performance metrics of GPUs connected to the system in real time. You can view both physical GPUs and MIG instances together, and efficiently sort or filter each item as needed.

 Supported Agent Version

Only supported on Linux. Requires agent version 2.9.3 or higher.

Basic Screen Overview

Click  **Live** to view or pause real-time data. You can search for GPU information using the search box.

Click on any column header in the data list to toggle between ascending and descending sort order.

- **Total:** Total number of GPU instances (including Physical + MIG)
- **Physical:** Number of actual physical GPU devices
- **MIG:** Number of MIG (Multi-Instance GPU) instances

GPU Information

Column	Description
Status	Current status of the GPU instance - N/A : GPU status is displayed only for Physical GPU targets. If the GPU Type is MIG, it is shown as N/A - Allocated : Same as the previous Active status - Unallocated : Same as the previous Inactive status - Active : Allocated and the 5-minute average GPU Utilization is 1% or higher - Idle : Allocated and the 5-minute average GPU Utilization is less than 1% - Effective : Allocated and the 5-minute average GPU Utilization exceeds 50%
HostName	Hostname of the server to which the GPU is connected

Column	Description
GPU Index	Unique index of the physical GPU or MIG instance (MIG is shown in a format like 0/6/0)
Model Name	GPU model name (e.g., NVIDIA A100-SXM4)
GPU Type	Type of the instance: either Physical or MIG

GPU Performance Metrics

For detailed descriptions of each metric, see [GPU Metrics](#).

↓ Download CSV

You can download the data as a CSV file.

☰ Column Settings

Click the ☰ **Column Settings** button at the top right of the screen to open the **Column Settings** window. In the **Column Settings** window, you can select which columns to display in the table.

- ✓ **Select All:** Select all columns
- ↺ **Reset to Default Order:** Reset the column order to default
- ↺ **Reset to Default Selection:** Restore to the default column selection

ⓘ The selected columns are saved in your browser cookies, so the selection persists even after refreshing the page.
If cookies are deleted or an error occurs, the default columns will be restored.

1. In the **Column Settings** window, select the columns to display under **GPU Information** and **GPU Performance Metrics**.
 - You can select columns by group or adjust their order.

 The Status and HostName fields under GPU Information cannot be removed.

2. Click the **Apply** button to save your settings.

Add Filter

You can add filter conditions to quickly find the GPUs you need.

1. Click the  **Filter** search bar to open the **Add Filter** window.
2. In the **Add Filter** window, select the desired GPU condition (filter key and condition), then click the **Apply** button.

Compare GPUs

When you select multiple GPUs to compare, the **Compare** panel appears below the GPU list, allowing you to visually compare metrics using charts.

You can track trends for the last 10 minutes and identify performance anomalies or bottlenecks.

- You can compare up to 20 items at a time.

GPU Detail

The **GPU Detail** page lets you monitor the GPU usage of a specific server with time-series charts and view the processes occupying the GPU in real time. Physical GPU and MIG (Multi-Instance GPU) instance information are available on a single screen, allowing you to quickly identify which workloads are consuming which GPU resources and by how much.

! Supported Agent Version

Supported on Linux only. Agent version 2.9.16 or later is recommended.

Main Screen

In the options bar at the top of the screen, you can select the time range and server to query. The selected server's inline information (IP, OS Type, CPU Cores, Memory Total, OS Version) and a **Feature Activation Guide** button are displayed together. Below the options bar, you can filter the GPUs to display by selecting them.

Usage for each item is as follows:

- **Time Range:** Select the time period to query. Use the left and right arrow buttons to navigate between periods.
- **Server Selection:** Select the server to monitor from the dropdown. The hostname and IP are displayed together.
- **GPU Selection:** Filter the GPUs to display. Click the selection box to open the GPU selection modal. The default is `* (Select All)`, which selects all Physical GPUs and MIG instances connected to the server.
- **Feature Activation Guide:** Click to open a panel with information on the GPU monitoring supported environment and the list of collectible metrics.

GPU indexes are displayed in the following format across chart legends and the process table:

- **Physical GPU:** Integers such as `0`, `1`, `4`
- **MIG Instance:** Slash notation such as `7/1/0`, `6/4/0`, `5/8/0` (`Physical/GI/CI`)

GPU Metric Charts

The left area of the screen displays the GPU usage of the selected server as five time-series chart cards. These consist of 2 line charts (both Physical GPU + MIG instances) and 3 stack/line toggle charts (Physical GPU only).

You can switch to line mode using the toggle button in the upper right of each card.

Chart	Type	Description
GPU Utilization (%)	Line	GPU core utilization
FB Memory Used (%)	Line	GPU frame buffer memory utilization
Power Usage	Stack/Line	GPU power consumption
PCIe TX+RX	Stack/Line	PCIe transmit/receive bandwidth
NVLink TX+RX	Stack/Line	NVLink transmit/receive bandwidth

Using the icons in the upper right of each chart card, you can download the data as a CSV file or save the chart as an image. Clicking the expand icon opens the **Chart Detail Drawer** for that metric category.

! For detailed definitions of each metric, refer to the [GPU Metrics](#) document.

Chart Detail (Drawer)

Clicking the expand icon in the upper right of each chart card opens a drawer on the right side where you can view the detailed metrics for that metric category. Within the drawer, you can change the time range and GPU selection independently from the main screen.

In the **Interpretation Guide** section at the bottom of the drawer, you can check the meaning and monitoring points for each metric.

Drawer	Included Metrics
GPU Utilization Detail	GPU core utilization, SM Active/Occupancy, FP-Tensor-DRAM activity, etc.
GPU Memory Detail	FB Total/Used/Free/Reserved, FB Used (%), memory copy utilization, etc.
Power/Temperature Detail	Power Limit/Usage, power violation rate, GPU temperature, thermal violation rate, etc.
PCIe Traffic Detail	PCIe TX+RX, PCIe TX, PCIe RX, NVLink TX+RX, etc.
NVLink Traffic Detail	NVLink TX+RX, NVLink TX, NVLink RX, PCIe TX+RX, etc.

GPU Process Table

The **GPU Process** table in the upper right of the screen displays the list of processes occupying the GPU on the selected server. You can check the GPU, CPU, and memory usage trends of each process at a glance, and filter by process name using the search box in the upper right of the table.

Each metric cell displays a **time-series mini chart** along with the maximum value (**Max**) for the query period.

Column	Description
Process Info	Process name, execution command, <code>PID xxxx · account</code> (3 lines)
GPU Utilization (%)	GPU Utilization of the entire GPU device to which the process is bound. This is not a directly measured value per process but shows the overall utilization of the connected GPU, so multiple processes using the same GPU may show identical or similar values. Mini chart + <code>Max N%</code>
FB Used (MiB)	GPU frame buffer memory size occupied by the process. Mini chart + <code>Max N MiB</code>
CPU (%)	CPU utilization of the process. Mini chart + <code>Max N%</code>
RSS (byte)	Host system memory usage of the process (Resident Set Size). Mini chart + <code>Max N GiB</code>

Column	Description
Process IO Total (bytes/s)	I/O bandwidth of the process. Mini chart + Max N KiB/s. Displays - if no data is collected
GPU Index	GPU index occupied by the process. Physical is an integer (0 , 1), MIG uses slash notation (5/8/0)

! VLLM::Worker_TP0 , VLLM::EngineCore , python3 app.py , etc. are the names of processes actually running in the user environment. All processes occupying the GPU are displayed in the table.

Server Information · GPU Inventory Panel

In the lower right panel, select the **Server Information** or **GPU Inventory** tab to view detailed information about the server currently being queried. You can adjust the ratio between the GPU process table and this panel using the drag handle between them.

Server Information

Item	Description
Host name (oid)	Hostname and OID
OS	Operating system type and version
GPU model	GPU model name (e.g., A100-40GB)
GPU count	Number of GPUs connected to the server (in the format Physical N / MIG M)
Server memory	Total server memory
Agent	Installed agent version

GPU Inventory

You can view the list of Physical GPUs and MIG instances connected to the server with the following columns. Click the **Detail View** button on the right side of the tab to open a drawer showing all columns.

ⓘ Only the inventory at the current point in time can be queried. To query previous points in time, use the **GPU Inventory** menu.

Column	Description
GPU Index	Physical is an integer (0 , 1), MIG uses slash notation (5/7/0 , 5/8/0)
GPU Type	Physical or MIG
Model Name	Official GPU model name (e.g., NVIDIA A100-SXM4-40GB)
Model Family	GPU model family (e.g., A100)
GPU Memory Size (MiB)	GPU memory size. Physical shows total memory, MIG shows the partitioned fragment size

For the inventory menu itself and how to use it, refer to the [GPU Inventory](#) page.


CUBE

Home > Select Project > **Analysis** > **Cube**

On the initial screen of the WhaTap Monitoring service, select a project and then select **Analysis** > **Cube**.

What is Cube?

We call the performance statistics generated every 5 minutes as **Cube**. Cube analysis uses each 5-minute performance data stored in cubes.



1 Cube selection panel

Select a cube based on the specific time in the cube selection panel in the left. Because the cubes are stored every 5 minutes, select a specific time.

- **Panel interest metric**

The interest metric in the server cube select panel are **CPU** and **Memory**.

- **Cube color**

The cubes are displayed in different colors depending on the number of the selected metrics.

- Dark blue: high figure
- Light blue: low figure

2 Application selector

Through the application selector option in the 2 area at the top, you can select data for specific applications. If you select  on the right, all available applications can be selected as follows:



- **Select all:** All applications are viewed. By default, **Select all** and **Active application** are selected.
- **View only selected:** Among the applications, only the specified applications are queried.
- **Search:** If there are many applications, you can search for them through the input field.
- **Number/Number:** Displays the number of applications specified/number of all applications.

3 Cube data

- If you select the  icon on the upper right of the chart, the bar-type chart appears.
- If you select the  icon on the upper right of the chart, line charts are provided.

In the area of 3, the cube data can be checked through the charts for **CPU**, **Memory**, **Disk**, and **Network**.

- **CPU**

Displays the applications in the project sorted in order of highest CPU usage.

- **Memory**

Displays the applications in the project sorted in order of highest memory usage.

- **Disk**

Displays the applications in the project sorted in order of highest disk usage.

- **Network**

Displays the applications in the project sorted in order of highest network traffic.

ⓘ WhaTap uses the IP2Location LITE database for [IP geolocation](#).

Metric Anomaly Detection

Home > Select Project > **Analysis** > **Metrics Anomaly Detection**

You can compare the patterns of various metrics with the expected patterns learned by the AI. Anomaly detections that deviate from the expected patterns can be viewed through the graph chart. You can check recurring patterns based on the past data and use them to predict future metrics values.

- ❗ The widgets placed on the screen are not saved and reset when you go to another menu.
- It is similar to **Analysis** > **Metrics Chart** except for the pattern and anomaly display.
- The alert, **Anomaly detection** is based on the **Metric anomaly detection**. For more information about the **Anomaly detection** alert setting, see [the following](#).

Checking the widgets



- ❶ The bright color areas of the graph are the estimated patterns analyzed by the AI.
- ❷ The blue graph is the metrics trend of the project.
- If it deviates from the expected pattern analyzed by the AI, it displays dots on the graph in Orange and Red steps. The values that are far from the expected pattern range appear in red.

Placing the widgets



1. In **1 Time**, set a desired time interval. You can set the interval up to 1 month.
2. In **2 Filter**, select a range of the metrics.
3. In the following list, select a metric for monitoring.


Place the selected metric to the right of the screen.

- Select **2 Tag** to check the metrics separately by selecting a monitoring target. A tag selection list appears on the right of the screen. Select desired items, and then add widgets.
- The widgets placed on the screen are not reflected in the chart even if the **Time** or **Tag** values have been changed.
- Select **X** on the upper right of the widget on the screen to delete.
- You can reposition a widget by dragging it while selecting the upper left part of the widget.
- You can resize a widget by selecting the upper right part of the widget and then dragging it.

Deleting a widget

To delete all widgets placed on the screen, select **Delete All Widgets** on the screen.

Checking the related metrics



Drag part of the chart in the widget. The **Similar pattern indicator** window appears on the right of the screen. You can search the related metrics for up to 6 hours based on the start time of the time range you dragged.

Metrics


What is a metric?

A **metric** is time-series data that represents the state of a monitoring target as numerical values. It is the result of processing performance data collected by the WhaTap agent from monitoring targets.

Through metrics, you can view key performance indicators such as memory usage per server and DB connection time at a glance. After identifying a problem, you can perform detailed analysis through logs and traces.

Resource usage statistics help you estimate the required amount of resources, which can also be used for cost optimization.

How WhaTap collects metrics



To collect metrics, install the WhaTap agent on the monitoring target. Metrics sent by the agent are stored on the WhaTap collection server.

For example, to monitor **Server**, first [install the agent](#). See [Server metrics](#) for available metrics.

Metric components

WhaTap **metrics** consist of the following components:


- **Category:** A name that groups related metrics.

- **Tags:** Labels that identify collection targets. Stores rarely-changing unique information such as IP, Oname, and Host. A metric can contain multiple tags.
- **Fields:** Stores actual measured values. A metric can contain multiple fields.
- **Time:** The time when the metric was collected.
- **Oid:** The unique number of the agent that collected the metric.
- **Oname:** The name of the agent that collected the metric.

Viewing and visualizing metrics

WhaTap provides menus to view collected metrics in various ways.

- **Metrics Search:** View the original metric data.
- **Metrics Chart:** View metric data through visualized charts.
- **Metrics Anomaly Detection:** Detect anomalies by comparing with metric patterns learned by AI.



Server Metrics

Home > Select Project > **Analysis** > **Metrics Search**

Metrics available in the **Metrics Search** menu. Each metric contains the following information. Some metrics may not have tag information.

- **Tags:** Information to identify the collection target (agent name, IP, etc.)
- **Fields:** Actual measured values (CPU usage, response time, etc.)

! CategoryV2 metrics are documented in [Server Metrics V2](#).

Category V1

agent_status_summary

- Collects agent status metrics
- Collection interval: 10 seconds

Fields

Field	Type	Unit	Description
inActTime	-	Millisecond(ms)	Amount of time the agent remains inactive
isActive	-	-	Whether the current agent is active or not (true / false)
isRebooted	-	-	Whether the server was recently restarted (true / false)
lastActTime	-	Millisecond(ms)	Last time when the agent was activated (0 : Disabled)
oid	-	-	Unique IDs for each agent in the project

Field	Type	Unit	Description
startTime	-	Millisecond(ms)	Timestamp when the agent was started.
status	-	-	Status of the current agent
uptime	-	Second (sec)	Server's uptime

❗ The values in the `status` field are distinguished as follows:

- `ok` : The server is in normal status where it is active.
- `inactive` : The server is disabled.
- `warning` : A warning event is in progress on the server.
- `fatal` : A critical event is in progress on the server.
- `paused` : The agent has paused data collection.

server_base

- Collects server resource metrics
- Collection interval: 5 seconds
- Statistical data: 5 minutes, 1 hour

Tags

Tag	Description	Remarks
cpu_cores	Number of server cores	-
host_ip	Server IP	-
oname	Server name	-
okindName	Server type name	-

Tag	Description	Remarks
os	OS type (Linux or Windows)	-

Fields

Field	Unit	Description	Remarks
cpu	Percent	All CPU utilization	Common
cpu_c1	Percent	CPU c1 state utilization	Windows
cpu_c2	Percent	CPU c2 state utilization	Windows
cpu_c3	Percent	CPU c3 state utilization	Windows
cpu_ctxt	Prime number	Context switching utilization	Linux
cpu_dpc	Percent	CPU DPC time percentage	Windows
cpu_idle	Percent	CPU idle rate	Common
cpu_interrupt	Percent	cpu interrupt	Linux
cpu_iowait	Percent	CPU iowait rate	Linux
cpu_irq	Percent	CPU IRQ utilization	Linux
cpu_load1	Prime number	Load average for 1 minute	Linux
cpu_load5	Prime number	Load average for 5 minutes	Linux
cpu_load15	Prime number	Load average for 15 minutes	Linux
cpu_new_procs_forked	Positive number	OS process fork call	Linux
cpu_procs_blocked	Positive number	Process blocked state count	Linux

Field	Unit	Description	Remarks
cpu_procs_running	Positive number	Process running state count	Linux
cpu_softirq	Percent	CPU softirq utilization	Linux
cpu_steal	Percent	CPU steal utilization	Linux
cpu_sys	Percent	CPU system utilization	Common
cpu_usr	Percent	CPU user utilization	Common
epochtime	Positive number	Epoch time in which the metric was collected	Common
memory_available	Byte	Available memory size	Common
memory_buffers	Byte	Memory buffer size	Linux
memory_cached	Byte	Memory cache size	Common
memory_free	Byte	Free memory size	Common
memory_pageFaults	Positive number	Memory page fault count	Common
memory_pavailble	Percent	Available memory utilization	Common
memory_pused	Percent	Memory utilization	Common
memory_shared	Byte	Shared memory size	Linux
memory_slab	Byte	Slab memory size	Linux
memory_sreclaimable	Byte	Slab reclaimable memory size	Linux
memory_sunreclaim	Byte	Slab claimable memory size	Linux
memory_swappused	Percent	Swap utilization	Common

Field	Unit	Description	Remarks
memory_swaptotal	Byte	Total swap size	Common
memory_swapused	Byte	Swap usage	Common
memory_total	Byte	Total memory size	Common
memory_used	Byte	Memory usage	Common
processorQueueLength	Positive number	Process queue length	Windows
uptime	Positive number	Server uptime after a boot (TimeStamp)	Common

server_cpu_core

- Collects metrics for each server CPU core
- Collection interval: 5 seconds
- Statistical data: -

Tags

Tag	Description	Remarks
cpu	Server core number	-
ip	Server IP	-
oname	Server name	-
os	OS type (Linux or Windows)	-

Fields

Field	Unit	Description	Remarks
idle	Percent	Core idle rate	Common
iowait	Percent	Core iowait rate	Windows
irq	Percent	Core IRQ utilization	Linux
load1	Prime number	Core load average for 1 minute	Linux
load5	Prime number	Core load average for 5 minutes	Linux
load15	Prime number	Core load average for 15 minutes	Linux
nice	Percent	Core nice utilization	Common
processorQueueLength	Positive number	Core process queue length	Windows
softirq	Percent	Core softirq utilization	Linux
steal	Percent	Core steal utilization	Linux
sys	Percent	Core sys utilization	Linux
usr	Percent	Core usr utilization	Linux

server_disk

- Collects disk metrics for each file system
- Collection interval: 5 seconds
- Statistical data: 5 minutes, 1 hour

Tags

Tag	Description	Remarks
deviceId	Device ID	-
deviceIdHash	Device ID hash value	-
fileSystem	File system type	-
mountPoint	File system's mount point	-
mountPointHash	File system's mount point hash	-
oname	Server name	-
os	OS type (Linux or Windows)	-

Fields

Field	Unit	Description	Remarks
blksize	Byte	Block size	Linux
count	Positive number	WhaTap only	Common
freePercent	Percent	File system's free space rate	Linux
freeSpace	Byte	File system's free space amount	Linux
ioPercent	Percent	IO utilization In some disk types, inaccurate values can be measured.	Linux
ioutilOrder	Positive number	WhaTap only	Linux

Field	Unit	Description	Remarks
queueLength	Mistake	Queue length	Common
totalInode	Positive number	Total number of Inodes	Linux
totalSpace	Byte	Total size of the file system	Linux
usedPercent	Percent	File system utilization	Linux
usedInode	Positive number	Number of used inodes	Linux
usedInodePercent	Percent	Inode utilization	Linux
userSpaceOrder	Positive number	WhaTap only	Common
readBps	Byte	Bytes read per second	Common
writeBps	Byte	Bytes written per second	Common
readIops	Prime number	Read count per second	Common
writeIops	Prime number	Written count per second	Common

server_network

- Collects network interface metrics
- Collection interval: 5 seconds
- Statistical data: 5 minutes, 1 hour

Tags

Tag	Description	Remarks
desc	Network interface name	-
descHash	Network interface's hash value	-
hwAddr	MAC address	-
ip	Interface IP address	-
oname	Server name	-
os	OS type (Linux or Windows)	-

Fields

Field	Unit	Description	Remarks
bandwidthOrder	Positive number	WhaTap only	-
count	Positive number	WhaTap only	Common
droppedIn	Positive number	Inbound drop	-
droppedOut	Positive number	Outbound drop	-
errorIn	Positive number	Inbound error	-
errorOut	Positive number	Outbound error	-
packetIn	pps	Inbound packet	-
packetOut	pps	Outbound packet	-

Field	Unit	Description	Remarks
trafficIn	bps	Inbound traffic	-
trafficInAcct	Positive number	WhaTap only	-
trafficOut	bps	Outbound traffic	-
trafficOutAcct	Positive number	WhaTap only	-

server_netstat

- Collects network protocol statistics
- Collection interval: 5 seconds
- Statistical data: 5 minutes, 1 hour

Tags

Tag	Description	Remarks
oname	Server name	-
os_arch	OS architecture name	-
os_cpucore	Number of server CPU cores	-
os_cpusocket	Number of server CPU sockets	-
os_cpuvendor	Server CPU's vendor name	-
os_ip	Server IP	-
os_memory	Server memory size	-
OS name type (Linux or Windows)	-	

Tag	Description	Remarks
os_public_ip	Public server IP	-
os_release	Server OS release name	-
whatap_build	WhaTap agent's build hash value	-
whatap_home	Path of the WhaTap agent Home	-
whatap_product	WhaTap product type	-
whatap_version	WhaTap agent version	-

Fields

Field	Unit	Description	Remarks
netstat_Close_Wait	Positive number	Number of CLOSE_WAIT state sockets	Common
netstat_Closing	Positive number	Number of closed sockets	Common
netstat_Established	Positive number	Number of established sockets	Common
netstat_Fin_Wait1	Positive number	Number of FIN_WAIT1 sockets	Common
netstat_Fin_Wait2	Positive number	Number of FIN_WAIT2 sockets	Common
netstat_Last_Ack	Positive number	Number of LASK_ASK sockets	Common
netstat_Listen	Positive number	Number of LISTEN sockets	Common
netstat_Syn_Recv	Positive number	Number of SYN_RECV sockets	Common
netstat_Syn_Sent	Positive number	Number of SYN_SENT sockets	Common
netstat_Time_Wait	Positive number	Number of TIME_WAIT sockets	Common

Field	Unit	Description	Remarks
netstat_Unknown	Positive number	Number of unknown sockets	Common
Server_Filedescriptors	Positive number	Number of OS file descriptors	Common

server_process

- Collects server process information
- Collection interval: 20 seconds
- Statistical data: 5 minutes

Tags

Tag	Description	Remarks
alias	Server alias	-
cmd1	Process execution command 1	-
cmd2	Process execution parameter	-
hash	Process hash value	-
name	Process name	-
oname	Server name	-
user	Linux user name for the process	-

Fields

Field	Unit	Description	Remarks
count	Positive number	Number of processes	Common

Field	Unit	Description	Remarks
cpu	Percent	CPU Utilization	Common
cpuAccumulated	Percent	WhaTap only	Common
cpuOrder	Order	WhaTap only	Common
memOrder	Order	WhaTap only	Common
memory	Percent	Memory utilization	Common
memoryAccumulated	Percent	WhaTap only	Common
openFileDescriptors	Positive number	Number of open file descriptors	Common
rbps	Byte	Bytes of the disk read per second	Common
rbpsAccmulated	Byte	WhaTap only	Common
riops	Positive number	Count read per second in the disk	Common
riopsAccmulated	Positive number	WhaTap only	Common
wbps	Byte	Bytes written per second in the disk	Common
wbpsAccmulated	Byte	WhaTap only	Common
wiops	Positive number	Count written per second in the disk	Common
wiopsAccmulated	Positive number	WhaTap only	Common
rss	Byte	RSS memory size	Common
rssAccmulated	Byte	WhaTap only	Common

References

- [Windows metrics material](#)
- [Linux metrics material](#)

Server Metrics V2

Home > Select Project > **Analysis** > **Metrics Search**

Metrics available in the **Metrics Search** menu. Each metric contains the following information. Some metrics may not have tag information.

- **Tags:** Information to identify the collection target (agent name, IP, etc.)
- **Fields:** Actual measured values (CPU usage, response time, etc.)

The CategoryV2 metrics of WhaTap server monitoring are a metric system that standardizes server resources in operational environments by domain. From core system resources such as CPU, memory, Swap, and kernel to major infrastructure indicators including disk device I/O, filesystem, network traffic, NIC configuration, NTP synchronization, processes, Windows services, InfiniBand, ZFS, and Veritas LV, they are provided in a consistent structure across on-premises and cloud server environments based on Linux, Windows, and Unix (HP-UX, AIX, Solaris).

CategoryV2 provides both categories that compare resource usage across all servers and categories that analyze each server's components at the detailed instance level. With these, operators can quickly grasp the resource distribution and anomalies of the entire infrastructure, and progressively trace down to the detailed components of a server where a problem occurred.

Whereas the existing CategoryV1 grouped indicators from multiple domains into a few categories, CategoryV2 improves the analysis structure and collection efficiency by separating categories per domain and differentiating collection intervals according to indicator characteristics. The collected CategoryV2 metrics integrate with the Metrics Explorer, so users can directly query the indicators they want and perform detailed analysis under various conditions.

! CategoryV2 can be used by enabling the `categoryv2.enabled=true` option in the agent settings.

CategoryV2 metric details

infra_cpu

A category that collects the usage state of the server's overall CPU resources. It provides indicators for checking CPU usage and load level across the board, such as user mode, kernel mode, I/O wait, idle, interrupt, steal, and load average. Use it to identify the server's overall CPU bottlenecks and system load state.

- Category: CPU
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
cpu	CPU	%	It refers to the overall CPU utilization excluding the idle percentage at 100%.	HP-UX, AIX, Solaris, Linux, Windows
user	User	%	It refers to the ratio of total CPU time spent executing general applications/processes in user mode.	HP-UX, AIX, Solaris, Linux, Windows
sys	Sys	%	The ratio of total CPU time spent in kernel mode (System mode) for handling system calls, kernel tasks, and driver processing.	HP-UX, AIX, Solaris, Linux, Windows
nice	Nice	%	It refers to the percentage of total CPU time used for executing low-priority (niced) user processes. The utilization of user processes with normal priority is typically included in user (%).	HP-UX, AIX, Solaris, Linux

Field	Title	Unit	Description	OS
wait	Wait	%	The ratio of total CPU time during which the CPU was idle waiting for I/O completion (I/O wait). This is interpreted as the time spent waiting for responses from external resources such as disk/network, rather than the time the CPU actually performed computations.	HP-UX, AIX, Solaris, Linux
total	User+Sys	%	The sum of the ratios used in user mode and kernel mode out of the total CPU utilization. Generally interpreted as the pure CPU execution utilization, excluding I/O wait, irq, softirq, and steal.	HP-UX, AIX, Solaris, Linux, Windows
total_wait	User+Sys+Wait	%	The sum of user, system, and wait time out of the total CPU time. This is a usage metric that includes user processing, kernel processing, and I/O waiting.	HP-UX, AIX, Solaris, Linux
idle	Idle	%	The ratio of total CPU time during which the CPU was idle and not performing any tasks. A higher value indicates a greater availability of CPU resources.	HP-UX, AIX, Solaris, Linux, Windows
lCores	Logical Cores	count	This refers to the number of logical CPUs (logical cores) recognized by the OS. It may differ from the number of physical cores, and when Hyper-Threading/SMT is enabled, the number of logical cores may appear larger.	HP-UX, AIX, Solaris, Linux, Windows
irq	Irq	%	The ratio of total CPU time spent handling hardware interrupts (IRQ). It may increase when the load of processing network cards, disks, and other hardware events is high.	Linux
softirq	Softirq	%	It refers to the ratio of total CPU time spent handling Software interrupts (SoftIRQ). This can increase when	Linux

Field	Title	Unit	Description	OS
			there are many deferred processing tasks in the kernel, such as network packet processing and block I/O post-processing.	
steal	Steal	%	In a virtualized environment, the ratio of time that the current virtual CPU is unable to execute due to the hypervisor prioritizing other virtual CPUs is referred to as CPU steal time. A high value may indicate insufficient CPU resources allocated to the VM or a contention state.	Linux
procUtil	Processor Util	%	A Windows-specific metric that refers to the overall CPU utilization of the entire system's processors. It is the same as the total CPU utilization seen in the Windows Task Manager.	Windows
loadavg_1m	LoadAvg 1m	count	The average system load (load average) over the last 1 minute. It reflects tasks that are currently running or waiting to run, as well as tasks in an uninterruptible wait state, and is a different concept from CPU utilization.	HP-UX, AIX, Solaris, Linux
loadavg_5m	LoadAvg 5m	count	The average system load over the last 5 minutes. It is used to observe a trend that is slightly smoother than short-term fluctuations.	HP-UX, AIX, Solaris, Linux
loadavg_15m	LoadAvg 15m	count	The average system load over the last 15 minutes. It is used to check the long-term trend of system load.	HP-UX, AIX, Solaris, Linux

infra_mem

A category that collects the server's overall memory usage state. It provides indicators for checking memory usage and pressure across the board, such as total memory capacity, usage, usage rate, cache, buffer, available memory, slab, page fault, and swap in/out. Use it to identify memory shortage, reclaim, paging, and cache

utilization state.

- Category: MEMORY
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
pUsed	Used	%	This refers to the ratio of currently used memory to the total memory. Linux typically interprets used as total - available based on MemAvailable, considering page cache and reclaimable memory as free memory. Windows generally interprets usage based on the amount excluding Available. AIX, HP-UX, and Solaris do not have a concept equivalent to Linux's MemAvailable, so they usually calculate based on the used value provided by the OS.	AIX, HP-UX, Solaris, Linux, Windows
used	Used	byte	This refers to the current memory usage. Linux typically calculates this based on the value of MemTotal - MemAvailable, converted to bytes. Windows generally interprets this based on usage excluding Available. AIX, HP-UX, and Solaris do not have a MemAvailable criterion like Linux, so they derive this from the used value provided by the OS.	AIX, HP-UX, Solaris, Linux, Windows
pCache	Cache	%	The ratio of the cache area to the total memory. In Linux, it is generally based on the Cached value from /proc/meminfo, referring to memory of the page cache nature, but does not include Buffers or SReclaimable. AIX, HP-UX, and Solaris do not have a concept equivalent to Linux's MemAvailable, so the cache value provided by the OS is typically used for calculation.	HP-UX, AIX, Solaris, Linux
cache	Cache	byte	The size of memory used as cache. In Linux, it is	HP-UX,

Field	Title	Unit	Description	OS
			generally based on the Cached value from /proc/meminfo, referring to memory of the page cache nature, but does not include Buffers or SReclaimable. AIX, HP-UX, and Solaris do not have a concept equivalent to Linux's MemAvailable, so the cache value provided by the OS is typically used for calculation.	AIX, Solaris, Linux
buffer	Buffer	byte	This value represents the amount of memory used as I/O buffers, expressed in bytes. It is a Linux-specific metric based on the Buffers value from /proc/meminfo.	Linux
pUsed_pCache	Used+Cache	%	This is an indicator that represents the current usage rate (pUsed) and cache ratio (pCache) together. It serves as a reference metric to view the proportion of memory actually in use by users/kernel and the cache area, and is more of a composite indicator for operational interpretation rather than a standard memory metric.	HP-UX, AIX, Solaris, Linux
shared	Shared Size	byte	This is the value representing the size of shared memory used across the entire server in bytes. It is a Linux-specific metric and may include tmpfs, shared memory, and some shared mapping areas.	Linux
size	Total Size	byte	The value representing the total physical memory size allocated to the system in bytes.	HP-UX, AIX, Solaris, Linux, Windows
available	Available	byte	This value represents the currently available memory size in bytes. In Linux, it refers to	Linux

Field	Title	Unit	Description	OS
			MemAvailable from /proc/meminfo, which is closer to an 'immediate usable estimate' that reflects not only free memory but also reclaimable cache. While instantaneous values are important, a continuous decrease over time, along with an increase in pageMajFaults, swapIn, and swapOut, may indicate that actual memory pressure is increasing.	
sreclaimable	Sreclaimable	byte	This value represents the size of the reclaimable area in the kernel slab cache in bytes. It may be reclaimed under memory pressure. A larger value does not necessarily indicate a problem; it also reflects the current level of cache utilization. However, if it excessively increases over time or if there is a decrease in available memory without proper reclamation, further investigation may be necessary.	Linux
sunreclaim	Sunreclaim	byte	This is the value representing the size of the area in the kernel slab cache that cannot be reclaimed, expressed in bytes. Since this memory must be retained by the kernel, it may have a more direct implication in terms of memory pressure compared to SReclaimable. A pattern of continuous increase in this value can help interpret the expansion of kernel memory usage or the possibility of abnormal accumulation.	Linux
slab	Slab	byte	This value represents the total size of slab memory used by the kernel for object caching, expressed in bytes. It generally refers to the Slab in /proc/meminfo and is interpreted as the sum of SReclaimable and SUNreclaim. While the current value is important, a continuous	Linux

Field	Title	Unit	Description	OS
			increase over time may indicate an expansion of kernel memory usage. Particularly, if there is an accompanying increase in SUnreclaim, it is necessary to consider the possibility of actual kernel memory pressure rather than just an increase in reclaimable cache.	
pageFaults	Page Faults	counts/s	This metric represents the total number of page faults that occurred during memory access across the entire system, expressed as a per-second value. It includes both minor faults and major faults, and does not necessarily involve disk I/O. While a momentary increase alone may not indicate a problem, a sustained rise in value, especially when accompanied by increases in pageMajFaults, scanned, and swapIn/Out, can help interpret the possibility of memory pressure or an increase in page cache misses.	Linux
pageMajFaults	Page Major Faults	counts/s	This metric represents the number of major page faults per second, which indicates the instances where the system had to actually read the required pages from disk/storage. The source of the reads can be file-backed pages from the filesystem or the swap area. If the value remains consistently high or shows an increasing trend, it suggests potential memory pressure that involves actual I/O, rather than just an increase in memory access, and can be directly linked to response latency. It is important to observe this alongside swapIn, fsIn, and available metrics.	Linux
scanned	Page Scan	pages/s	This refers to the number of pages scanned by the system to reclaim memory. A higher value	HP-UX, AIX,

Field	Title	Unit	Description	OS
			indicates more active reclaim activities for memory, and consistently high levels may suggest increasing memory pressure. It is important to interpret the duration of sustained increases rather than one-time spikes, along with the relationship to freed, steal, and swapIn/Out.	Solaris, Linux
freed	Page Free	pages/s	The number of free pages secured through page reclamation, page release, etc. This is an indicator that shows the results of the system's reclaim or memory return activities. It is important to observe the relationship with scanned as well as the continuous increase of the value itself. If scanned is high but freed does not keep up sufficiently, it may indicate low memory reclamation efficiency.	HP-UX, AIX, Solaris, Linux
steal	Page Steal	pages/s	It refers to the actual number of pages reclaimed during the page reclaim process in Linux. This is associated with the process of securing page cache or inactive pages in situations of memory shortage. A temporary increase in this value may indicate normal reclaim behavior; however, if it remains consistently high, it may suggest that the system is continuously reclaiming pages under memory pressure.	Linux
stealRatio	Page Steal Ratio	%	This refers to the ratio of actual pages retrieved from the pages scanned in Linux. It is generally calculated as $\text{steal} / \text{scanned} * 100$, where a higher value indicates better retrieval efficiency compared to scanning, and a lower value may suggest a higher proportion of active pages, indicating a situation where retrieval is not	Linux

Field	Title	Unit	Description	OS
			effective. (0 has no meaning) However, the interpretation may vary depending on the workload characteristics, so it is important to consider it alongside scanned, freed, and available.	
swapIn	Swap In	pages/s	It refers to the number of pages read back into memory per second from the swap (paging space) area of the disk. If the value consistently occurs or shows an increasing trend, it may indicate a higher dependency on swap due to memory shortages, which could be directly linked to response delays. It is advisable to prioritize the persistence of occurrences over temporary spikes.	HP-UX, AIX, Solaris, Linux, Windows
swapOut	Swap Out	pages/s	It refers to the number of pages exported per second from memory to the disk's swap (paging space) area. A continuous occurrence of this value indicates that the system is pushing pages to swap to free up memory. A consistently high pattern observed alongside swapIn may suggest a memory pressure state with a high potential for performance degradation.	HP-UX, AIX, Solaris, Linux, Windows
fsIn	Filesystem In	pages/s	This refers to the number of pages that have been brought into memory by file system read operations. Unlike Swap In, it specifically refers to pages loaded by file system I/O. While the current value is significant, the increasing pattern along with pageMajFaults, the process's rdiskIo, and available can help interpret that actual file-based page-in activity is increasing.	HP-UX, AIX, Solaris, Linux, Windows
fsOut	Filesystem	pages/s	This refers to the number of pages written to	HP-UX,

Field	Title	Unit	Description	OS
	Out		disk as a result of filesystem write operations. Unlike Swap Out, it is more appropriate to interpret this as the volume of page writes of the filesystem writeback/flush nature. Continuous increases or periodic spikes in the value can be correlated with the characteristics of operations such as flush/writeback/checkpoint.	AIX, Solaris, Linux, Windows
<code>poolPaged</code>	Pool Paged	byte	A Windows-specific metric that represents the size in bytes of the paged pool. It refers to a kernel memory area that can be paged to disk if necessary.	Windows
<code>poolNonPaged</code>	Pool Nonpaged	byte	A Windows-specific metric that indicates the byte size of the nonpaged pool. It refers to a kernel memory area that must always reside in physical memory.	Windows
<code>systemCacheResident</code>	System Cache Resident	byte	A Windows-specific metric that indicates the size in bytes of the system cache memory currently residing in physical memory. It reflects the kernel memory usage related to the file system cache.	Windows

infra_swap

A category that collects the server's overall swap or paging space usage state. With total swap size, usage, and usage rate, you can grasp the level of disk-based virtual memory use. Use it to check whether swap is being used due to memory shortage and the overall virtual memory pressure state.

- Category: SWAP
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
size	Total Size	byte	This is the value representing the total size of the swap area configured in the system, expressed in bytes. In Unix/Linux systems, it refers to the total size of the swap/paging space, while in Windows, it corresponds to the total size of the page file-based virtual memory area.	HP-UX, AIX, Solaris, Linux, Windows
pUsed	Used	%	This refers to the current percentage of swap used compared to the total swap space. It is calculated as $\text{used} / \text{size} * 100$. A higher value may indicate an increased reliance on disk-based virtual memory due to memory shortages.	HP-UX, AIX, Solaris, Linux, Windows
used	Used	byte	This is the value representing the size of the currently used swap space in bytes. In Unix/Linux systems, it refers to the actual occupied space in the swap/paging space, while in Windows, it corresponds to the page file usage.	HP-UX, AIX, Solaris, Linux, Windows

infra_kernel

A category that collects the OS kernel and system behavior state. It provides operational indicators from the kernel and scheduler perspective, such as run queue, block queue, context switch, interrupt, syscall, fork, and exec. Use it to identify CPU contention, I/O wait, scheduling load, and kernel entry frequency.

- Category: System Kernel
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
runq	RunQueue	count	It refers to the number of processes or threads that are ready to run but have not yet been allocated CPU time, sitting in the run queue. A higher value may indicate CPU contention or that there are more runnable tasks than available CPU resources. It is important to view the run queue alongside CPU utilization. If CPU utilization is high	HP-UX, AIX, Solaris, Linux, Windows

Field	Title	Unit	Description	OS
			and the run queue is also high, there is a significant possibility of a CPU bottleneck. Conversely, if CPU utilization is low but the run queue is high, issues such as locks, scheduling, or priority problems may be suspected.	
blockq	BlockQueue	count	The number of processes or threads that are in a waiting state and cannot be executed immediately due to I/O, paging, or resource contention. This number can generally increase when there is a high volume of disk I/O, memory waits, or resource contention. A high value suggests a potential bottleneck related to I/O, paging, or resource waits rather than a CPU shortage. It is advisable to monitor disk performance metrics alongside swap metrics.	HP-UX, AIX, Solaris, Linux
pendingq	PendingQueue	count	This is a metric specific to AIX, generally referring to the number of threads waiting for raw device I/O requests to complete. It may increase during periods of storage response delays or high raw I/O contention.	AIX
waitq	WaitQueue	count	In the context of HP-UX / Solaris systems, this refers to the number of processes or threads that have been swapped out or are temporarily removed from scheduling target and are waiting. Unlike the simple run queue, this is distinct from being immediately runnable state that can receive CPU.	HP-UX, Solaris
fork	Fork	counts/s	The number of fork() calls per second. Fork is a system call that creates a new child process by duplicating the current process. A high value may indicate an environment where short-lived processes are frequently created. It can naturally be high in environments that frequently generate processes, such as shell scripts, batch jobs, and CGI/process-oriented applications, so it is more important to determine if there has been a sudden increase compared to usual levels.	HP-UX, AIX, Solaris, Linux

Field	Title	Unit	Description	OS
exec	Exec	counts/s	The number of exec() calls per second. exec is a system call that replaces the memory image and execution context of the current process with a new program while maintaining the process's PID. Unlike fork, it does not create a new process but substitutes the existing process with another program. In environments where program replacement occurs frequently, such as shell scripts, batch jobs, and CGI/process-based applications, it can naturally be high, so it is more important to determine if there has been a sudden increase compared to usual levels.	HP-UX, AIX, Solaris
interrupt	Interrupt	counts/s	The number of hardware interrupts processed by the CPU per second. This can increase with a higher number of network, disk, timer, and other device events. A high value indicates that it is advisable to analyze the increase in network packet throughput, storage I/O, and whether there is a concentration of interrupts from specific devices.	HP-UX, AIX, Solaris, Linux, Windows
syscall	SystemCall	counts/s	It refers to the number of system calls processed by the kernel per second. This number increases when user processes request kernel services such as file I/O, networking, and process control. A high value does not necessarily indicate an anomaly, but it may reflect changes in workload characteristics or an increase in kernel entry frequency. It can be high in applications with significant file/socket I/O.	HP-UX, AIX, Solaris, Windows
ctxswitch	ContextSwitch	counts/s	It refers to the number of context switches per second. It increases when the CPU switches from one running process or thread to another task. A higher value may indicate more multitasking, thread contention, lock contention, and frequent short-duration tasks. If the context switch count is very high, it may suggest that the CPU is spending more time on scheduling overhead than	HP-UX, AIX, Solaris, Linux, Windows

Field	Title	Unit	Description	OS
			on actual work, so it is advisable to analyze it alongside run queue, system calls, interrupts, and lock contention situations.	

infra_disk

A category that collects the aggregate performance of all physical disks connected to the server. It provides indicators for checking the server's overall disk I/O scale, such as read/write throughput (Byte/s) and IOPS. Use it to identify the server's overall disk load level.

- Category: DISK
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
read	Read	bytes/s	The sum of the data read per second from all physical disks of the server, expressed in Byte/s. In other words, it is the aggregated value of the read throughput of all physical disks in the server.	HP-UX, AIX, Solaris, Linux, Windows
write	Write	bytes/s	The sum of the data written per second across all physical disks of the server, expressed in Byte/s. In other words, it is the aggregated value of the write throughput for all physical disks of the server.	HP-UX, AIX, Solaris, Linux, Windows
total	IO Total	bytes/s	The total data throughput, expressed in Byte/s, is the sum of the read and write throughput from all physical disks of the server. This is the aggregated value of the server's overall disk I/O throughput corresponding to read + write.	HP-UX, AIX, Solaris, Linux, Windows
readIops	Read IOPS	counts/s	This value represents the sum of the number of read I/O operations processed per second across all physical disks of the	HP-UX, AIX,

Field	Title	Unit	Description	OS
			server, expressed in IOPS. In other words, it is the aggregated read IOPS of all physical disks in the server. IOPS is a metric that indicates the number of I/O requests processed by the disk per second, without reflecting the size of the data itself. Therefore, in environments with a high number of small block I/O, IOPS may be high while throughput (Byte/s) may be relatively low, whereas in environments with a high number of large block I/O, IOPS may be low but throughput (Byte/s) may appear high.	Solaris, Linux, Windows
writeIops	Write IOPS	counts/s	This value represents the sum of the number of write I/O operations processed per second across all physical disks of the server, expressed in IOPS. In other words, it is the aggregated write IOPS of all physical disks in the server. IOPS is a metric that indicates the number of I/O requests processed by the disk per second, without reflecting the size of the data itself. Therefore, in environments with a high number of small block I/O, IOPS may be high while throughput (Byte/s) may be relatively low, whereas in environments with a high number of large block I/O, IOPS may be low but throughput (Byte/s) may appear high.	HP-UX, AIX, Solaris, Linux, Windows
totalIops	Total IOPS	counts/s	The total number of I/O operations, expressed in IOPS, is the sum of the read and write I/O operations processed per second across all physical disks of the server. It corresponds to the aggregate value of the server's total disk I/O operations, which is readIops + writeIops. IOPS is a metric that indicates the number of I/O requests processed by the disk per second, without reflecting the size of the data itself. Therefore, in environments with a high number of small block I/O operations, IOPS may be high while throughput (Byte/s) may be relatively low, whereas in environments with a high number of large block I/O operations, IOPS may be low but throughput (Byte/s) may appear high.	HP-UX, AIX, Solaris, Linux, Windows
busyMax	Busy Max	%	Among all physical disks connected to the server, this indicates the utilization of the disk with the highest busy (%) value during the collection interval. In other words, it represents the busy level of the disk that operated the most actively among all disks on the	HP-UX, AIX, Solaris, Linux,

Field	Title	Unit	Description	OS
			server. A higher value may indicate that I/O load is concentrated on a specific disk or that there is a potential bottleneck. However, this value is collected at a default interval of 2 seconds, while the busy values for individual disks are collected at a default interval of 5 seconds, so individual disk busy values may not match exactly, even on charts at the same timestamp. Particularly during low busy intervals or moments of sudden load increases, the gap can look larger because of the mismatch between the collection interval and the sampling point. Therefore, this value can be used as a representative indicator to quickly check for signs of disk load at the server level, and it is advisable to examine detailed metrics for individual disks to confirm actual bottleneck disks.	Windows
diskName	Busy Disk	string	This is the name of the disk with the highest busy (%) value among all physical disks connected to the server.	HP-UX, AIX, Solaris, Linux, Windows

infra_network

A category that aggregates and provides the server's overall network usage and per-interface usage level. It collects Tx/Rx bps, Tx/Rx pps, total traffic, maximum NIC usage rate, and the corresponding NIC name for NICs excluding loopback. With these, you can check the trend of the server's overall network load and identify the interface with the highest network usage rate.

- Category: NETWORK
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
rxbps	Net Rx	bps	This value represents the total amount of data received (Rx) per second through all NICs of the server (excluding loopback) in bps (bits per second). In other words, it is the aggregated	HP-UX, AIX, Solaris,

Field	Title	Unit	Description	OS
			network input bandwidth of the server.	Linux, Windows
txbps	Net Tx	bps	This value represents the total amount of data transmitted (Tx) per second through all NICs of the server (excluding loopback), expressed in bps (bits per second). In other words, it is the aggregated network transmission bandwidth of the entire server.	HP-UX, AIX, Solaris, Linux, Windows
totalbps	Net Total	bps	This value represents the total data transfer amount received (Rx) and transmitted (Tx) per second through all NICs of the server (excluding loopback), expressed in bps (bits per second). It corresponds to the aggregate value of the server's overall network bandwidth, which is the sum of rxbps and txbps.	HP-UX, AIX, Solaris, Linux, Windows
rxpps	Rx PPS	counts/s	This is the sum of the packets received (Rx) per second through all NICs of the server (excluding loopback), expressed in pps (packets per second) or count/s. In other words, it is the aggregate value of the server's total network received packet throughput. Since pps is based on the number of packets regardless of the size of the transmitted data, a high number of small packets can result in a high pps while the bps may be relatively low.	HP-UX, AIX, Solaris, Linux, Windows
txpps	Tx PPS	counts/s	This value represents the sum of packets transmitted (Tx) per second through all NICs of the server (excluding loopback) in pps (packets per second) or count/s units. In other words, it is the aggregate value of the server's overall network transmission packet throughput. Since pps is based on the number of packets regardless of the size of the transmission, a high number of small packets can result in a high pps while the bps may be relatively low.	HP-UX, AIX, Solaris, Linux, Windows
totalpps	Total PPS	counts/s	The total number of packets received (Rx) and transmitted (Tx) per second through all NICs of the server (excluding loopback)	HP-UX, AIX,

Field	Title	Unit	Description	OS
			is represented in pps (packets per second) or count/s units. This is the aggregate value corresponding to rxpps + txpps for the entire network packet throughput of the server. Since pps is based on the number of packets regardless of the size of the transmitted data, a high number of small packets can result in a high pps while the bps may be relatively low.	Solaris, Linux, Windows
linkUsageMax	NIC Max Usage	%	This is the percentage value of the utilization of the NIC with the highest link usage (excluding loopback) among all NICs present on the server. It serves as an indicator to quickly identify which NIC among the server's network interfaces is under the highest load. Generally, it corresponds to the maximum value of the link usage of each NIC, and NICs for which link speed information cannot be obtained may be excluded from the calculation or may not reflect a value.	HP-UX, AIX, Solaris, Linux, Windows
nicName	Max Usage NIC Name	string	This value indicates the name of the NIC with the highest link utilization (excluding loopback) among all NICs present on the server. It serves as an identifier for which NIC the linkUsageMax value occurred on. For example, it may display names such as eth0, ens192, bond0 on Linux, or adapter names on Windows. If there are no calculable NICs, the value may be empty.	HP-UX, AIX, Solaris, Linux, Windows

infra_netstat

Information aggregating the server's TCP network connection states by state.

- Category: Network State Summary
- Default collection interval: 60 sec
- Minimum supported agent: Linux 2.9.18+, Windows 2.9.18+, AIX 1.4.9+, Solaris 1.4.9+, HP-UX 1.4.9+

Field	Title	Unit	Description	OS
netstat_Listen	Listen	count	The number of TCP sockets in LISTEN state waiting to receive external connection requests on the current server. This is primarily used to verify whether the service port is properly opened.	HP-UX, AIX, Solaris, Linux
netstat_Established	Established	count	The number of TCP connections in the ESTABLISHED state that are currently successfully established on the server and capable of data transmission. This indicates the actual number of active sessions.	HP-UX, AIX, Solaris, Linux
netstat_Close_Wait	Close_Wait	count	The number of TCP connections in the CLOSE_WAIT state, where a remote request to terminate the connection has been made, but the local application has not yet fully closed the socket. If this value remains consistently high, it is necessary to check for potential socket cleanup delays or leaks in the application.	HP-UX, AIX, Solaris, Linux
netstat_Closing	Closing	count	The number of TCP connections in the CLOSING state, where both sides are in the process of terminating the connection. Typically represented briefly, and if it persists or increases excessively, it is necessary to check for abnormal termination handling.	HP-UX, AIX, Solaris, Linux
netstat_Fin_Wait1	Fin_Wait1	count	The number of TCP connections in the FIN_WAIT1 state, where the local side has initiated the connection termination and sent a FIN, waiting for a response from the remote side. This corresponds to the initial stage of the connection termination procedure.	HP-UX, AIX, Solaris, Linux
netstat_Fin_Wait2	Fin_Wait2	count	The number of TCP connections in the FIN_WAIT2 state, which are waiting for the final termination signal from the remote side after receiving a connection termination request from the local side. If this state persists for a long time, it may indicate a termination delay or issues on the peer side.	HP-UX, AIX, Solaris, Linux

Field	Title	Unit	Description	OS
<code>netstat_Time_Wait</code>	Time_Wait	count	The number of TCP connections in the TIME_WAIT state, which is maintained for a certain period to prevent socket reuse conflicts similar to delayed packet processing after connection termination. In environments with a very high number of short connections, this value can increase significantly.	HP-UX, AIX, Solaris, Linux
<code>netstat_Syn_Recv</code>	Syn_Recv	count	The number of TCP connections in the SYN_RECV state that have received a connection request (SYN) from the remote side, responded to it, and are now waiting for final confirmation. If the value increases sharply, it is necessary to check for a surge in connection requests or the possibility of a SYN Flood attack.	HP-UX, AIX, Solaris, Linux
<code>netstat_Syn_Sent</code>	Syn_Sent	count	The number of TCP connections in the SYN_SENT state after the local side has sent a connection request (SYN) and is waiting for a response from the remote side. This can increase due to remote server response delays, network path issues, or firewall problems.	HP-UX, AIX, Solaris, Linux
<code>netstat_Last_Ack</code>	Last_Ack	count	The number of TCP connections in the LAST_ACK state after receiving a termination request from the remote side, where the local side also sends a termination response and waits for the final confirmation. This state is generally short-lived, and if it persists at a high level, it is necessary to check for any delays in the termination process.	HP-UX, AIX, Solaris, Linux
<code>netstat_Unknown</code>	Unknown	count	The number of connections that are not classified as defined major TCP states or whose states could not be clearly identified at the time of collection. Generally, this value should not be large, and if it continues to increase, it is necessary to check the collection logic or the state mapping criteria specific to the OS.	HP-UX, AIX, Solaris, Linux

infra_ntp_info

A category that collects the server's time synchronization state. It provides NTP usage, synchronization daemon, synchronization state, the NTP server in use, offset, offset direction, time zone, etc. Use it to check how much the system time differs from the reference time and whether NTP synchronization is normal.

- Category: NTP Information
- Default collection interval: 60 sec
- Minimum supported agent: Linux 2.9.18+, Windows 2.9.18+, AIX 1.4.9+, Solaris 1.4.9+, HP-UX 1.4.9+

Field	Title	Unit	Description	OS
checkTime	Check Time	string	This refers to the time when NTP information collection or inspection was performed. It indicates the local system time of the server where the Agent is installed at the moment the result was generated.	HP-UX, AIX, Solaris, Linux, Windows
ntpInUse	NTP In Use	string	Indicates whether the NTP-based time synchronization feature is in use on the system. 'Yes' means that NTP or a time synchronization daemon is configured and operational, while 'No' indicates that it is not in use or the related daemon is not running.	HP-UX, AIX, Solaris, Linux, Windows
ntpDaemon	NTP Daemon	string	This indicates the name of the time synchronization daemon or service used in the system. For example, in Linux/Unix systems, it may display chrony, ntpd, xntpd, and in Windows, it may show w32time.	HP-UX, AIX, Solaris, Linux, Windows
timeSynced	Time Synced	string	Indicates whether the current system is properly synchronized with the NTP server. 'Yes' means that it is currently synchronized with a valid NTP server, while 'No' indicates that NTP is in use but has not yet synchronized or has not secured a valid synchronization target.	HP-UX, AIX, Solaris, Linux, Windows
ntpServer	NTP	string	Indicates the hostname or IP address of the NTP server	HP-UX,

Field	Title	Unit	Description	OS
	Server		currently used for time synchronization. If synchronization is not functioning properly, the value may be empty.	AIX, Solaris, Linux, Windows
offset	Offset (ms)	ms	This value represents the absolute difference in milliseconds (ms) between the current system time and the reference NTP server time. The direction information regarding whether it is fast or slow is distinguished in the offsetDirection field.	HP-UX, AIX, Solaris, Linux, Windows
offsetDirection	Offset Direction	string	Indicates whether the current system time is faster or slower than the reference NTP server time. For example, FAST (+) means the local system time is ahead, while SLOW (-) means the local system time is behind.	HP-UX, AIX, Solaris, Linux, Windows
timeZone	Time Zone	string	This indicates the time zone information set in the system. Depending on the operating system, it may be displayed in different formats such as Asia/Seoul, KST, or Korea Standard Time.	HP-UX, AIX, Solaris, Linux, Windows

infra_process

A category that aggregates and shows the server's overall process state. You can check the overall process/thread state, such as total process count, thread count, zombie count, runnable/sleeping counts, and file descriptor usage. Use it to identify system operational state such as process count spikes, zombie accumulation, and FD exhaustion.

- Category: Process State Summary
- Default collection interval: 2 sec

Field	Title	Unit	Description	OS
procTotal	Process Total	count	It refers to the total number of processes currently existing in the OS. A high process count does not necessarily indicate an issue. In environments with batch processing, web server prefork structures, and many short job units, it can naturally be high.	HP-UX, AIX, Solaris, Linux, Windows
threadTotal	Thread Total	count	This refers to the total number of threads currently existing in the OS. It can increase with the number of multithreaded applications and may reflect system load more directly than the number of processes. The thread count can better represent the actual level of concurrency than the process count. In Java, WAS, DB, and multithreaded application environments, the trend of threadTotal is important.	HP-UX, AIX, Solaris, Linux, Windows
handleTotal	Handle Total	count	The total number of handles currently open or in use in the Windows OS. This includes references to various kernel objects such as files, registry keys, processes, threads, events, and mutexes.	Windows
runnableTotal	Runnable Total	count	The number of processes with a state of R in Linux refers to those that are currently running or are runnable but waiting for CPU in the run queue.	Linux
sleepingTotal	Sleeping Total	count	It refers to the number of processes in the S state on Linux. This state is interruptible sleep, which includes processes that can wake up due to signals while waiting for events or resource responses. It is the most common range for processes in a waiting state.	Linux
idleTotal	Idle Total	count	It refers to the number of processes in the I state on Linux. From the perspective of the kernel scheduler, it indicates the number of tasks marked as idle, which may be observed more in kernel internal workers or idle-type tasks rather than in general user processes.	Linux
disksleepTotal	Disk	count	It refers to the number of processes in the D state on Linux.	Linux

Field	Title	Unit	Description	OS
	Sleep Total		This state is known as uninterruptible sleep, primarily including processes waiting for disk I/O, filesystem, devices, or kernel resources. A high and persistent value may indicate a storage or kernel wait bottleneck.	
zombieTotal	Zombie Total	count	It refers to the number of processes that have terminated but remain in a zombie (Z) state because they have not yet been collected by the parent process. A continuous increase may indicate issues with the parent process's wait handling or delays in process collection.	HP-UX, AIX, Solaris, Linux
stoppedTotal	Stopped Total	count	It refers to the number of processes in the T state on Linux. Generally, it indicates processes that have been stopped by job control signals (SIGSTOP, SIGTSTP, etc.).	Linux
tracingstopTotal	Tracing Stop Total	count	The number of processes in the process state t in Linux. Generally, it indicates processes that are stopped in a tracing stop state by a debugger (ptrace).	Linux
deadTotal	Dead Total	count	It refers to the number of processes in state X on Linux. This indicates that the processes are in the process of termination or have completely terminated, but it is rarely observed in typical operating environments.	Linux
wakekillTotal	Wake Kill Total	count	The number of processes in the K state in Linux. This may be classified as an internal kernel state and is rarely observed in typical operating environments, or its meaning may vary depending on the kernel version.	Linux
wakingTotal	Waking Total	count	It refers to the number of processes in the W state on Linux. This can be viewed as a transitional state during waking processing, and in the latest kernels, it may be rarely observed or only environment-dependent.	Linux
otherTotal	Other Total	count	It refers to the number of processes with other process states than the explicitly separated R/S/I/D/Z/T/t/X/K/W in	Linux

Field	Title	Unit	Description	OS
			Linux. The inclusion criteria may vary depending on kernel version differences, transient states at the moment of observation, and implementation differences.	
openfdCurrent	Open fd Used	count	The number of currently used file descriptors across the Linux OS. This includes the number of open objects managed by the kernel as file descriptors, such as files, sockets, pipes, and devices.	Linux
openfdMax	Open fd Max	count	It refers to the maximum number of file descriptors set across the entire Linux OS. Generally, it is interpreted based on the system-wide file descriptor limit (/proc/sys/fs/file-max).	Linux
openfdPUsed	Open fd Used	%	This refers to the overall file descriptor usage of the Linux OS. It is generally calculated as $\text{openfdCurrent} / \text{openfdMax} * 100$, and a higher value indicates an increased risk of global FD exhaustion in the system.	Linux

infra_nic_info

A category that collects the configuration information of each network interface in the server. It provides static or semi-static information such as NIC name, IP, MAC, link, state, duplex, speed, mode, and driver. Use it for monitoring NIC configuration changes, checking redundancy structures, and identifying interfaces.

- Category: NIC Information
- Default collection interval: 60 sec
- Minimum supported agent: Linux 2.9.18+, Windows 2.9.18+, AIX 1.4.9+, Solaris 1.4.9+, HP-UX 1.4.9+

Field	Title	Unit	Description	OS
nicName	NIC Name	string	The name of the network interface recognized by the operating system, which can display not only physical NICs but also alias interfaces, bond/team/bridge, EtherChannel,	HP-UX, AIX, Solaris,

Field	Title	Unit	Description	OS
			APA, AGGR, IPMP, and their member NICs together.	Linux, Windows
ipAddress	IP Address	string	This is the IPv4 address configured on the NIC, which may be displayed together with multiple addresses separated by commas (,), and if there is no IP, it may be indicated as none.	HP-UX, AIX, Solaris, Linux, Windows
macAddress	MAC Address	string	The MAC address assigned to the NIC, and the logical/ aggregate interface may display the representative MAC reported by the operating system, and if it cannot be verified, it may be displayed as none.	HP-UX, AIX, Solaris, Linux, Windows
link	Link	string	The link connection status determined by the operating system or driver is indicated, where 'yes' means connected and 'no' means disconnected or inactive, which may not fully align with the actual cable status.	HP-UX, AIX, Solaris, Linux, Windows
state	State	string	The current interface status of the NIC is indicated, where UP means available, DOWN means inactive or unavailable, and Unknown/Unknown! means status cannot be determined.	HP-UX, AIX, Solaris, Linux, Windows
duplex	Duplex	string	The Duplex mode of the NIC is indicated, where Full means full duplex, Half means half duplex, and unknown refers to a state that cannot be determined due to the characteristics of the operating system, driver, or device.	HP-UX, AIX, Solaris, Linux, Windows
speed	Speed	string	The current link speed of the NIC or the operational speed reported by the operating system, which refers to the speed currently negotiated or recognized rather than the maximum	HP-UX, AIX, Solaris,

Field	Title	Unit	Description	OS
			supported speed, and some aggregated interfaces may be displayed as a sum.	Linux, Windows
mode	Mode	string	Additional information regarding the NIC aggregation, redundancy, or teaming modes is indicated, for example, EtherChannel:<mode> in AIX, APA in HP-UX, Teaming Mode/ Load Balancing in Windows, or AGGR or IPMP in Solaris.	HP-UX, AIX, Solaris, Linux, Windows
windowsDriver	Windows Driver	string	This is the driver or adapter description information used by the NIC in the Windows environment.	Windows
Redundancy configuration display				

The `nicName` and `mode` fields are displayed as follows depending on the redundancy configuration.

Item	Description
Linux bonding	Member interfaces belonging to a bond are displayed as <code>bond0_eth1</code> , in the <code>bondname_membername</code> form, and <code>mode</code> shows the mode of that bond
Windows team interface	<code>mode</code> is displayed in the <code>LBFO/<team mode>/<load balancing method></code> form (for example, <code>LBFO/SwitchIndependent/TransportPorts</code>)
Windows team member NIC	<code>mode</code> is displayed in the <code>LBFO_MEMBER/<configured role>/<operational status></code> form (for example, <code>LBFO_MEMBER/Active/Active</code>)

A Windows team member NIC is a hidden adapter without its own IP address, so it did not appear in the list before; now member NICs are collected as well. On Linux, the issue where the control entry of the bonding driver (`bonding_masters`) was incorrectly displayed as a network interface has been fixed.

infra_process_group

A category that targets all processes collected from the server without exclusion and, by default, groups processes run with the same cmd by the same User even if their run arguments differ, providing resource usage by group. It sums and displays usage such as CPU, memory, I/O, and page fault per group, so users can grasp at a glance how a given program family affects server resources. In addition, through the user-defined process group management feature, you can directly define groups by specific server/owner/argsPattern conditions and manage them with a user-specified process group name.

- Category: ProcessGroup
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
cmd	Command	string	This is the execution command (cmd) of the processes belonging to the Process group.	HP-UX, AIX, Solaris, Linux, Windows
name	Name	string	By default, it displays cmd, but if it is a Process Group defined by the user through the process group management menu, it displays the name of the defined Process Group.	HP-UX, AIX, Solaris, Linux, Windows
owner	Owner	string	The execution account (user) of the processes belonging to the Process group. Since they are grouped by the same user, the processes within the group have the same owner.	HP-UX, AIX, Solaris, Linux, Windows
pidList	Process ID List	string	Only the PIDs collected from the Process TopN are provided in the list. Therefore, the number of displayed PIDs may differ from the Process count.	HP-UX, AIX, Solaris, Linux, Windows
custom	Custom	string	Indicates whether the user-defined process group is True/False.	HP-UX, AIX, Solaris, Linux,

Field	Title	Unit	Description	OS
				Windows
cpu	CPU	%	The total CPU utilization used by the processes belonging to the Process Group. This value includes both user mode and system mode utilization, and observing sustained increases over time, prolonged high utilization, and periodic spikes can be useful for interpreting the operational load characteristics of the group or identifying abnormal CPU occupancy.	HP-UX, AIX, Solaris, Linux, Windows
cpuUser	CPU User	%	The CPU utilization of processes belonging to the Process Group in user mode indicates the proportion of CPU used for executing application logic. This metric helps to interpret whether the observed values, including sustained increases over time, repetitive spikes, and correlations with procCnt/threadCnt changes, reflect normal load due to increased throughput or excessive user-space operations.	HP-UX, AIX, Solaris, Linux, Windows
cpuSys	CPU System	%	The CPU utilization of processes belonging to the Process Group in kernel/system mode refers to the CPU usage associated with system calls, I/O processing, and kernel resource usage. A consistently high value or an excessively large proportion of system mode compared to user mode can indicate increased I/O processing, a rise in kernel calls, or lock contention.	HP-UX, AIX, Solaris, Linux, Windows
rss	RSS	byte	The sum of the Resident Set Size (RSS) of the processes belonging to the Process Group. This refers to the size of the memory loaded in actual physical memory. Since this value may include shared memory, in servers with many processes that use shared memory extensively, the sum of the RSS of the processes may exceed the total memory size. Observing not only the current occupancy but also the sustained increase over time and any sharp fluctuation patterns can help determine	HP-UX, AIX, Solaris, Linux

Field	Title	Unit	Description	OS
			whether there is a memory leak or abnormal occupancy.	
pRss	RSS	%	The total RSS of the process group represents the proportion it occupies in the total physical memory of the system. It is calculated as $\text{rss} / \text{total physical memory} * 100$. The total RSS of the process group may include shared memory redundantly, which can make it appear to exceed 100% of the total physical memory of the system. Observing not only the absolute utilization but also the long-term upward trend and peak maintenance can help assess the impact of memory pressure from the process group.	HP-UX, AIX, Solaris, Linux
vsz	VSZ	byte	The sum of the total virtual memory address space size acquired by the processes belonging to the Process Group. This refers to the size of the allocated or reserved address space, not the actual physical memory usage, and may include file mappings, shared libraries, and swap reservations. Therefore, it is important to observe the long-term trend of the current value itself, along with the differences from other actual memory metrics (RSS/PSS/USS).	HP-UX, AIX, Solaris, Linux, Windows
pss	PSS	byte	The sum of the PSS (Proportional Set Size) of processes belonging to a Process Group in Linux. This reflects the physical memory usage based on the proportion divided by the number of processes sharing shared memory, making it useful for understanding the actual memory burden closer to the occupancy level than RSS. Observing the usage trends of shared memory across the entire Process Group along with the increase over time helps interpret changes in memory usage structure. In Linux, collection and aggregation are performed only when the <code>/proc/[pid]/smaps_rollup</code> file is available for each PID, considering the collection cost. From Agent 2.9.20 onward, calculating PSS for a process	Linux: smaps_rollup supported environment

Field	Title	Unit	Description	OS
			whose RSS exceeds a certain size (4 GiB by default) can put load on the target process, so the agent skips PSS collection by default and the value may be empty. Therefore, when multiple processes are aggregated into a single group, processes with no collected PSS are excluded from the group total, so the group's total PSS may appear smaller than the actual memory usage. You can change the threshold in the settings.	
uss	USS	byte	The sum of the Unique Set Size (USS) of processes belonging to a Process Group in Linux. This represents the total amount of non-shared physical memory used exclusively by each process, which is useful for understanding the scale of memory that is likely to be returned upon the termination of the Process Group. Observing the trend of pure occupied memory, excluding shared memory, helps in assessing the potential for memory leaks within a specific service group. In Linux, data is collected and aggregated only when the /proc/[pid]/smaps_rollup file is available for each PID, considering the cost of collection. From Agent 2.9.20 onward, calculating USS for a process whose RSS exceeds a certain size (4 GiB by default) can put load on the target process, so the agent skips USS collection by default and the value may be empty. Therefore, when multiple processes are aggregated into a single group, processes with no collected USS are excluded from the group total, so the group's total USS may appear smaller than the actual memory usage. You can change the threshold in the settings.	Linux: smaps_rollup supported environment
minFlt	Page Minor Fault	counts/s	The sum of the minor page faults per second for the processes belonging to the Process Group. This refers to faults that occurred but were resolved internally in memory without disk I/O. While the value itself is significant, a gradual increase in conjunction with RSS	Linux

Field	Title	Unit	Description	OS
			growth or workload increase may indicate normal memory access expansion, whereas a sudden spike may suggest a change in access patterns or increased memory reclamation.	
majFlt	Page Major Fault	counts/s	The sum of the major page faults per second for the processes belonging to the Process Group. This refers to faults that required reading the necessary pages from disk or swap space. Typically, this includes reads from file-backed pages or the swap area. A consistently high or increasing value serves as an important indicator for interpreting memory pressure, increased page-ins, and potential response latency.	Linux
workingSet	Working Set	byte	This is a Windows-specific metric that represents the sum of the Working Sets of processes belonging to the Process Group. It indicates the size of the memory currently resident in physical memory. While instantaneous values are important, observing long-term trends, recovery after sudden spikes, can help identify memory pressure or abnormal occupancy patterns.	Windows
private	Private	byte	This is a Windows-specific metric that represents the total private memory exclusively used by processes belonging to the Process Group. It indicates the usage of non-shared private memory and serves as a metric to observe the memory growth unique to the process. A pattern of continuous increase over time that is not reclaimed can be useful for interpreting signs of memory leaks.	Windows
poolPaged	Pool Paged	byte	This is a Windows-specific metric that represents the total paged pool memory usage associated with processes belonging to the Process Group. It refers to the kernel memory area that can be paged to disk if	Windows

Field	Title	Unit	Description	OS
			necessary. Observing not only the current value but also the long-term trend of increase and recovery can help identify increases in kernel resource usage or abnormal occupancy patterns.	
totalCharIo	Process IO Total	bytes/s	The total character-based I/O (char I/O) throughput generated per second by the processes belonging to the Process Group. This may include not only disk but also pipes, terminals, caches, and some buffer-based I/O, so there may be differences from actual physical disk I/O. While the current throughput is important, observing the sustained increase over time, spikes, and repetitive work patterns can help interpret the overall I/O activity characteristics of the group.	HP-UX, AIX, Solaris, Linux, Windows
rcharIo	Read Character IO	bytes/s	The total throughput of read character I/O operations generated per second by the processes belonging to the Process Group. This includes not only disk reads but also read requests through buffers, caches, terminals, etc. Therefore, it is important to observe not only instantaneous values but also repetitive patterns associated with the workload cycle, peak moments, and their relationship with CPU and Disk I/O metrics.	Linux, Windows
wcharIo	Write Character IO	bytes/s	The total write character I/O throughput generated per second by the processes belonging to the Process Group. This includes write requests not only to the disk but also through buffers, caches, terminals, etc. A sustained increase or recurring spike pattern in values can help interpret the correlation with log output increases, buffer flushes, and data processing tasks.	Linux, Windows
rdiskIo	Read Disk IO	bytes/s	This is a Linux-specific metric that represents the total throughput of read I/O operations actually performed by processes belonging to a Process Group per second from the disk. Since reads processed from the filesystem	Linux

Field	Title	Unit	Description	OS
			cache are excluded, this value is closer to the actual storage device reads than rcharIo. Therefore, in environments with a high cache hit rate, rcharIo may be high while rdiskIo may be low, making it appropriate to interpret this metric as an indicator of the actual disk-dependent read load.	
wdiskIo	Write Disk IO	bytes/s	This is a Linux-specific metric that represents the total write I/O throughput to the disk generated per second by processes belonging to a Process Group. This value may differ from the total amount of write requests generated by the processes (wcharIo) and can be reflected with a time lag due to the effects of page cache/buffer cache and writeback operations. Therefore, it is appropriate to interpret this metric as an indicator of the write load actually recorded on the storage.	Linux
threadCnt	Thread	count	The current total number of threads belonging to the Process group. This is the sum of the thread counts of all processes within the group.	HP-UX, AIX, Solaris, Linux, Windows
procCnt	Process	count	This refers to the total number of processes belonging to the current Process group. It can be interpreted as the number of processes executed by the same user with the same command.	HP-UX, AIX, Solaris, Linux, Windows
fdCnt	File Descriptor	count	A Linux-specific metric that represents the total number of file descriptors currently opened by processes belonging to a Process Group. This includes open file descriptors for files, sockets, pipes, devices, etc.	Linux
handleCnt	Handle	count	A Windows-specific metric that represents the total number of handles currently held by the processes belonging to a Process group. It includes references to various kernel objects such as files, processes, threads,	Windows

Field	Title	Unit	Description	OS
			events, and registry keys.	

infra_process_topn

A category that selects and stores only the individual processes with high resource usage on the server at collection time. By default, it preferentially collects the top 30 processes with individual CPU usage of 0.1% or higher, and separately from this condition, additionally collects up to the top 10 processes with RSS of 20 MB or more even if CPU usage is 0, storing up to about 40 processes. Through this category, users can check, from a detailed analysis perspective, which PIDs are actually using a lot of CPU or memory. In other words, while Process Group is for viewing the broad flow of resource usage, Process TopN is for concretely tracing the individual processes that are the cause.

- Category: ProcessTopN
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
cmd	Command	string	This refers to the value extracted from the execution command of the process. It is parsed based on cmdline or executable file name.	HP-UX, AIX, Solaris, Linux, Windows
owner	Owner	string	This value represents the user account that executed the process.	HP-UX, AIX, Solaris, Linux, Windows
args	Arguments	string	This value represents the execution arguments (command line parameters) of the process.	HP-UX, AIX, Solaris, Linux, Windows
stime	StartTime	string	This value represents the start time of the process in epoch time.	HP-UX, AIX, Solaris,

Field	Title	Unit	Description	OS
				Linux, Windows
pid	PID	string	This value represents the Process ID (PID) of the corresponding process.	HP-UX, AIX, Solaris, Linux, Windows
ppid	PPID	string	This value represents the PPID (Parent Process ID) of the process, which is the PID of the parent process that created this process.	HP-UX, AIX, Solaris, Linux, Windows
procInfo	Process Information	string	This is an informative field used for process identification and display. It consists of a combination of owner, pid, ppid, stime, cmd(args), etc.	HP-UX, AIX, Solaris, Linux, Windows
cpu	CPU	%	This refers to the total CPU utilization used by the process. It includes the values for user mode and system mode utilization, and in a multi-core environment, it is adjusted based on the total number of logical cores, displayed within a maximum range of 100%. In other words, it can be interpreted as the usage ratio of this process compared to the overall CPU resources of the system. While instantaneous utilization is important, observing whether there is a sustained high utilization or periodic spikes can be useful for interpreting the CPU resource occupancy characteristics.	HP-UX, AIX, Solaris, Linux, Windows
cpuUser	CPU User	%	This indicates the CPU utilization used by the process in user mode. It serves as a metric to observe the proportion of CPU used for application logic execution, and analyzing sustained increases or repetitive spikes over time can help identify the characteristics of user-	HP-UX, AIX, Solaris, Linux, Windows

Field	Title	Unit	Description	OS
			space operational load.	
cpuSys	CPU System	%	This refers to the CPU utilization used by the process in kernel/system mode. It is utilized to interpret the CPU share associated with system calls, I/O processing, and kernel resource usage. A pattern where the value is consistently high or the system share is excessively large compared to user share can help suspect influences such as kernel calls, I/O, and lock contention.	HP-UX, AIX, Solaris, Linux, Windows
rss	RSS	byte	The RSS (Resident Set Size) value of the process indicates the size of memory currently loaded in physical memory. This value includes shared memory, so in servers with processes that utilize a lot of shared memory, the sum of the RSS of the processes may exceed the total memory size. It is a representative metric showing the current memory occupancy, and if it continuously increases over time, it is useful for determining the possibility of memory leaks or abnormal occupancy.	HP-UX, AIX, Solaris, Linux
pRss	RSS	%	The ratio of the RSS of the process to the total physical memory of the system is indicated. Rather than focusing solely on the absolute value, observing the long-term upward trend and whether it maintains a peak can help interpret how much a specific process affects system memory pressure.	HP-UX, AIX, Solaris, Linux
vsz	VSZ	byte	This refers to the total size of the virtual memory address space secured by the process. It is based on the size of the mapped or reserved virtual address space, not the actual physical memory usage. Therefore, it is important to look at the differences between this value and RSS/PSS/USS, as well as the long-term trends. There may be cases where VSZ is	HP-UX, AIX, Solaris, Linux, Windows

Field	Title	Unit	Description	OS
			large, but the actual occupied memory is not significant.	
pss	PSS	byte	This is the PSS (Proportional Set Size) value of the process. It represents the size of physical memory based on the proportion divided by the number of processes sharing shared memory, making it useful for understanding memory usage that is closer to the actual occupancy level than RSS. It is particularly useful when comparing memory burdens per process in environments with heavy use of shared libraries or shared memory. In Linux, it is collected only in environments where the /proc/[pid]/smaps_rollup file is available, considering the collection cost. From Agent 2.9.20 onward, calculating PSS for a process whose RSS exceeds a certain size (4 GiB by default) can put load on the target process, so the agent skips PSS collection by default and the value may be empty. You can change the threshold in the settings.	Linux: smaps_rollup supported environment
uss	USS	byte	This is the USS (Unique Set Size) value of the process. It indicates the size of the non-shared physical memory that this process uses exclusively, and is useful for understanding the scale of memory that is likely to be returned upon process termination. It serves as an indicator for viewing the pure occupied memory, excluding shared memory, and a pattern of continuous increase and non-recovery over time is particularly useful for assessing the potential memory leak of the process itself. In Linux, it is collected only in environments where the /proc/[pid]/smaps_rollup file is available, considering the collection cost. From Agent 2.9.20 onward, calculating USS for a process whose RSS exceeds a certain size (4 GiB by default) can put load on the target process, so the agent skips USS collection by default and the value may be empty. You	Linux: smaps_rollup supported environment

Field	Title	Unit	Description	OS
			can change the threshold in the settings.	
minFlt	Page Minor Fault	counts/s	The number of minor page faults per second for the process. This refers to faults that occurred but were resolved internally in memory without disk I/O. While the value itself is significant, a gradual increase in conjunction with RSS growth or workload increase may indicate normal memory access expansion, whereas a sudden spike may suggest a change in access patterns or increased memory recycling.	Linux
majFlt	Page Major Fault	counts/s	This refers to the number of major page faults per second for the process. It indicates faults where the required pages had to be read from disk/swap space. Typically, this may include reads from file-backed pages or swap space. If the value remains consistently high or shows an increasing trend, it becomes an important indicator for interpreting memory pressure, increased page-ins, and potential response delays.	Linux
workingSet	Working Set	byte	This is a Windows-specific metric that indicates the size of memory currently being used by the process in physical memory. It serves as a representative metric for observing the actual memory usage of a process. While the current value is important, analyzing long-term trends, recovery after sudden spikes, and other patterns can help identify memory pressure or abnormal occupancy patterns.	Windows
private	Private	byte	A Windows-specific metric that indicates the size of private memory exclusively used by the process. It represents the amount of dedicated memory that is not shared with other processes, serving as an indicator for observing the memory increase unique to the process. A pattern of continuous increase over time that is not reclaimed can be useful for	Windows

Field	Title	Unit	Description	OS
			interpreting signs of memory leaks.	
poolPaged	Pool Paged	byte	This is a Windows-specific metric that indicates the paged pool memory usage associated with the process. It relates to the usage of kernel memory areas that can be paged to disk if necessary. Observing not only the current value but also the long-term trend of increase and recovery can help identify increases in kernel resources or abnormal memory usage patterns.	Windows
totalCharIo	Process IO Total	bytes/s	This refers to the total character-based I/O (char I/O) throughput generated by the process per second. It may include I/O through not only disks but also pipes, terminals, buffers, caches, and file system layers, so there may be differences from actual physical disk I/O. Observing both the current throughput and the spikes over time, as well as sustained high usage patterns, can help interpret the I/O activity characteristics of the process.	HP-UX, AIX, Solaris, Linux, Windows
rcharIo	Read Character IO	bytes/s	This refers to the read character I/O throughput generated by the process per second. It may include read requests not only from disk reads but also from buffers, caches, terminals, etc. Therefore, it is important to consider not only instantaneous values but also the repetitive patterns associated with the workload cycle and the points of spikes.	Linux, Windows
wcharIo	Write Character IO	bytes/s	This refers to the write character I/O throughput generated by the process per second. It may include write requests through disk writes, as well as buffers, caches, and terminals. A sustained increase in values or repetitive spike patterns can help interpret the correlation with log outputs, buffer flushes, and data processing tasks.	Linux, Windows

Field	Title	Unit	Description	OS
<code>rdiskIo</code>	Read Disk IO	bytes/s	This is a Linux-specific metric that represents the total throughput of read I/O operations actually performed on the disk by the process per second. Since reads processed from the filesystem cache are excluded, this value is closer to the actual storage device reads than <code>rcharIo</code> . Therefore, in environments with a high cache hit rate, <code>rcharIo</code> may be high while <code>rdiskIo</code> may be low, making it appropriate to interpret this metric as an indicator of the actual disk-dependent read load.	Linux
<code>wdiskIo</code>	Write Disk IO	bytes/s	This is a Linux-specific metric that represents the total write I/O throughput actually generated by the process to the disk per second. This value may differ from the total amount of write requests generated by the process (<code>wcharIo</code>) and can be reflected with a time lag due to the effects of page cache/buffer cache and writeback operations. Therefore, it is appropriate to interpret this as a metric for observing the write load recorded on the actual storage.	Linux
<code>threadCnt</code>	Thread	count	The number of threads currently held by the process. It is useful for understanding the concurrency level of multi-threaded applications.	HP-UX, AIX, Solaris, Linux, Windows
<code>fdCnt</code>	File Descriptor	count	A Linux-specific metric that indicates the number of file descriptors currently opened by the process. This includes open file descriptors for files, sockets, pipes, and devices.	Linux
<code>handleCnt</code>	Handle	count	A Windows-specific metric that indicates the number of handles currently held by the process. It includes references to various kernel objects such as files, processes, threads, events, and registry keys.	Windows

Kubernetes metadata fields

On Linux servers where a container runtime is running, the Pod information the process belongs to is collected as well. After finding a process that uses a lot of resources on the server, you can immediately check which container of which Pod that process is.

Field	Title	Unit	Description	OS
<code>namespace</code>	Namespace	string	Kubernetes namespace of the container the process belongs to. Empty for a non-container process	Linux
<code>podName</code>	Pod Name	string	Name of the Pod the process belongs to	Linux
<code>containerName</code>	Container Name	string	Name of the container the process belongs to	Linux

The three fields are collected by default. If you set `process.topn.k8s.enabled=false`, collection stops and the three fields are not sent. On servers without a container runtime, the values are empty.

The supported agent version is Linux 2.9.23 or later.

infra_win_service

A category that collects state information per Windows service. It provides service name, PID, state, status, start mode, exit code, etc. Use it to check the run state, startup configuration, and abnormal termination of Windows services.

- Category: Windows Service
- Default collection interval: 60 sec
- Minimum supported agent: Windows 2.9.19+

Field	Title	Unit	Description	OS
<code>svcName</code>	Service Name	string	This value represents the unique service name of the Windows service. It corresponds to the internal identifier in the Service Management tool (services.msc) or the system service name, and may differ from the Display Name.	Windows
<code>pid</code>	PID	string	This value represents the PID (Process ID) of the process currently	Windows

Field	Title	Unit	Description	OS
			running the service. If the service is not running, it may be displayed as empty or 0. In cases where multiple services operate within a single process, such as svchost.exe, they may share the same PID.	
state	State	string	This value represents the current service status according to the Service Control Manager (SCM). For example, states such as Running, Stopped, Start Pending, Stop Pending, Paused, Pause Pending, and Continue Pending may be displayed.	Windows
startMode	Start Mode	string	This value indicates how this service is configured to start during system boot or service initiation. It can display as Automatic, Manual, or Disabled, and may be differentiated as Automatic (Delayed Start) depending on the environment.	Windows
exitCode	Exit Code	string	This value represents the exit code returned during the last termination or start/stop operation of the service. Generally, 0 is interpreted as a normal termination or no error, while a non-zero value can be referenced when tracing causes such as service start failure, execution errors, or abnormal termination.	Windows

infra_nic_perf

A category that collects performance indicators per individual NIC. It provides Tx/Rx bandwidth, packet count, error, drop, collision, etc. per interface. Use it to check a specific NIC's bottleneck, packet loss, and concentration of interrupts/errors.

- Category: by NIC
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
name	NIC Name	string	This value represents the name of the network interface. For example, it displays eth0, ens192, bond0 in Linux, and adapter names in Windows.	HP-UX, AIX, Solaris,

Field	Title	Unit	Description	OS
				Linux, Windows
totalbps	NIC Total	bps	This is the total amount of data transmitted (Tx) and received (Rx) through the specified NIC, expressed in bps (bits per second).	HP-UX, AIX, Solaris, Linux, Windows
rxbps	NIC RX	bps	This value represents the amount of data received (Rx) per second through the specified NIC, measured in bps (bits per second).	HP-UX, AIX, Solaris, Linux, Windows
txbps	NIC Tx	bps	The value representing the amount of data transmitted per second (Tx) through the specified NIC, measured in bps (bits per second).	HP-UX, AIX, Solaris, Linux, Windows
totalpps	NIC Total PPS	counts/s	The total number of packets received (Rx) and transmitted (Tx) through the specified NIC per second is represented in pps (packets per second) or count/s. This corresponds to rxpps + txpps.	HP-UX, AIX, Solaris, Linux, Windows
rxpps	NIC Rx PPS	counts/s	The number of packets received per second (Rx) through the specified NIC is represented in pps (packets per second) or count/s.	HP-UX, AIX, Solaris, Linux, Windows
txpps	NIC Tx PPS	counts/s	The number of packets transmitted (Tx) per second through the specified NIC is represented in pps (packets per second) or count/s.	HP-UX, AIX, Solaris,

Field	Title	Unit	Description	OS
				Linux, Windows
collision	NIC Collision	counts/s	This is the value converted to an average rate per second of the number of packet collisions detected during the transmission (Tx) process of the corresponding NIC. It is primarily significant in half-duplex Ethernet environments, while in modern full-duplex environments, it is generally normal for this value to be close to 0.	HP-UX, AIX, Solaris, Linux
errIn	NIC Error In	counts/s	This is the value converted to the average rate of change per second of the number of packets received (Rx) through the specified NIC that encountered errors. It may include received errors such as CRC errors, frame errors, and overruns, and the scope of inclusion may vary depending on the OS/driver.	HP-UX, AIX, Solaris, Linux, Windows
errOut	NIC Error Out	counts/s	This is the value converted to the average rate of packets with errors that were transmitted (Tx) through the corresponding NIC per second. It may include transmission errors such as carrier errors, aborts, and transmit queue issues, and the scope of inclusion may vary depending on the OS/driver.	HP-UX, AIX, Solaris, Linux, Windows
dropIn	NIC Drop In	counts/s	This is the average rate of packets dropped that were received (Rx) through the specified NIC and not processed at the kernel or NIC layer. It can increase due to buffer shortages, queue overflow, driver processing delays, etc.	HP-UX, AIX, Solaris, Linux, Windows
dropOut	NIC Drop Out	counts/s	This is the average rate of packets dropped before actual transmission among the packets transmitted (Tx) through the specified NIC, expressed as a value per second. It can increase due to factors such as transmission queue overflow, driver issues, or interface status anomalies.	HP-UX, AIX, Solaris, Linux, Windows
linkSpeed	NIC Link Speed	bps	This is the value representing the link speed of the NIC in Mbps. It refers to the transmission bandwidth capacity of the	HP-UX, AIX,

Field	Title	Unit	Description	OS
			currently negotiated or configured interface, not the actual usage. However, if the link speed cannot be determined due to the OS, driver, or type of virtual interface, it may be collected as unknown or null, and in such cases, the value may not be displayed.	Solaris, Linux, Windows
linkUsage	NIC Utilization	%	This value represents the total transmission and reception throughput of the NIC (NIC Total) expressed as a percentage of the link speed (NIC Link Speed). It indicates the actual traffic usage level compared to the available link bandwidth of the NIC, with higher values signifying that the link is being utilized more heavily. However, if the link speed (NIC Link Speed) is unknown or null, the utilization cannot be calculated, and thus no value may be displayed.	HP-UX, AIX, Solaris, Linux, Windows

infra_infiniband

A category that collects state and performance indicators per InfiniBand or HCA port. It provides port state, link layer, speed, LID, Tx/Rx throughput, packet count, symbol error, link recovery, constraint error, etc. Use it to check IB fabric connection state, link quality, and per-port bottlenecks/errors.

- Category: by Infiniband
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
interfaceName	Interface Name	string	This is an identifier that represents the HCA device name along with the port number. For example, mlx5_0:1 refers to port 1 of the mlx5_0 device.	Linux
linkLayer	Link Layer	string	Indicates the link layer type of the port. It displays either InfiniBand or Ethernet and is used to check which link mode the port is operating in.	Linux

Field	Title	Unit	Description	OS
state	State	string	Indicates the logical state of the port. For example, Down, Init, Active, etc. can be displayed, which serves as a criterion for determining whether the port is participating normally in the IB fabric.	Linux
physState	Physical State	string	Indicates the physical link status of the port. For example, Polling, LinkUp, Disabled, etc. are displayed and are used to determine the cable connection and physical layer negotiation status.	Linux
rate	Link Speed	string	Indicates the current link speed of the port. For example, values such as 100Gb/s and 200Gb/s are displayed, which are used to verify the results of link negotiation or the maximum operational speed of the port.	Linux
lid	Local Identifier	string	This represents the LID (Local Identifier) assigned to the corresponding port in the InfiniBand environment. It serves as the logical address of the port within the IB fabric, and may be meaningless or empty in Ethernet mode.	Linux
portXmitData	Transmit Data	bytes/s	This value represents the data throughput transmitted (Tx) per second on the specified port, measured in Bytes/s. It indicates the transmission usage based on the port.	Linux
portRcvData	Receive Data	bytes/s	This value represents the throughput of data received (Rx) per second on the specified port, measured in Bytes/s. It indicates the port-based receive usage.	Linux
portXmitPackets	Transmit	counts/s	This value represents the total number of	Linux

Field	Title	Unit	Description	OS
	Packets		packets transmitted per second on the specified port. It is a packet rate metric independent of packet size.	
portRcvPackets	Receive Packets	counts/s	This value represents the total number of packets received per second on the specified port. It is a packet rate metric that is independent of packet size.	Linux
unicastRcvPackets	Unicast Receive Packets	counts/s	This value indicates the number of unicast packets received per second on the specified port. It is useful for monitoring the general traffic flow with a single destination.	Linux
unicastXmitPackets	Unicast Transmit Packets	counts/s	The value representing the number of unicast packets transmitted per second on the specified port.	Linux
multicastRcvPackets	Multicast Receive Packets	counts/s	This value represents the number of multicast packets received per second on the specified port. It can be referenced when analyzing the impact of multicast-based communication or IB fabric control traffic.	Linux
multicastXmitPackets	Multicast Transmit Packets	counts/s	The value representing the number of multicast packets transmitted per second on the specified port.	Linux
symbolError	Symbol Errors	counts/s	This value represents the rate of symbol errors occurring per second at the physical link layer. Generally, it may indicate bit errors, cable quality issues, or instability in optical modules/ports.	Linux
portRcvErrors	Port Receive	counts/s	This value represents the per-second change rate of the total error counter recorded	Linux

Field	Title	Unit	Description	OS
	Errors		during the reception process for the port. It may include received frame anomalies, integrity issues, protocol errors, etc.	
portRcvRemotePhysicalErrors	Remote Physical Errors	counts/s	This value indicates the rate of change per second of the error counter received by this port due to a physical layer issue occurring at the opposite port. It may suggest a problem with the remote port or the cable/module status on the other side of the link.	Linux
portRcvSwitchRelayErrors	Switch Relay Errors	counts/s	This value represents the rate of change per second of the error counter recorded when the switch fails to properly relay packets. It may indicate issues with the switch's forwarding path or anomalies in the fabric.	Linux
linkErrorRecovery	Link Error Recovery	counts/s	This value represents the rate of change per second of link error recovery events. It may indicate that link reinitialization or error recovery is occurring frequently. If the value continues to increase, link instability may be suspected.	Linux
portXmitConstraintErrors	Transmit Constraint Errors	counts/s	This value represents the per-second change rate of errors recorded due to violations of transmission-side constraints. For example, inconsistencies in fabric policies such as SL, VL, QoS, and partition key can be the cause.	Linux
portRcvConstraintErrors	Receive Constraint Errors	counts/s	This value represents the rate of change per second of errors recorded due to violations of receiving-side constraints. It may increase in cases of mismatches in fabric policies, partitions, and QoS between the transmitting and receiving sides.	Linux

Field	Title	Unit	Description	OS
localLinkIntegrityErrors	Local Link Integrity Errors	counts/s	This value represents the per-second change rate of link integrity errors based on the local port. It may be associated with CRC/bit errors, signal quality degradation, or issues with cable or port hardware.	Linux
excessiveBufferOverflowErrors	Buffer Overflow Errors	counts/s	This value indicates the rate of errors that could not be processed normally due to a receive buffer overrun, which may increase when burst traffic is high relative to port processing capacity or when there are delays in system/driver/firmware processing.	Linux
linkDowned	Link Down Events	counts/s	This value represents the rate of change per second for events where the link transitions to a Down state. It may increase in situations such as cable disconnection, disabling the relative port, link instability, and equipment rebooting.	Linux
portXmitDiscards	Transmit Discards	counts/s	This value represents the rate of change per second of packets discarded before actual transmission in the transmission queue. Causes may include buffer shortages, congestion, queue constraints, and port status anomalies.	Linux
vl15Dropped	VL15 Dropped Packets	counts/s	This value indicates the rate of change per second of the number of discarded packets in the VL15 (Virtual Lane dedicated to Subnet Management) queue. It may suggest congestion in Subnet Manager-related control traffic or anomalies in the fabric management path.	Linux
Field	Title	Unit	Description	OS

Field	Title	Unit	Description	OS
---	---	:---	---	---
<code>rdmaNetDevice</code>	RDMA Net Device	string	Name of the operating system network interface connected to that RDMA device and port. Multiple names are separated by (for example, <code>ib0</code> , <code>enp175s0f0</code>). Empty when there is no mapping information	Linux

You can view the RDMA port together with the operating system network interface, so use it when checking the error metrics of an InfiniBand port and the traffic metrics of the corresponding interface at the same time. When port-level mapping is not provided, such as for a pure InfiniBand port, device-level mapping is displayed instead.

The supported agent version is Linux 2.9.23 or later.

infra_disk_perf

A category that collects performance indicators per individual physical disk. It provides detailed I/O performance such as per-disk throughput, IOPS, busy, service time, and wait time. Use it to check a specific disk's bottleneck, latency, and queue growth.

- Category: by Disk
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
<code>name</code>	Disk Name	string	Value displaying the name of the physical disk.	HP-UX, AIX, Solaris, Linux, Windows
<code>busy</code>	Disk Busy	%	This value represents the ratio of time the physical disk was busy during the observation period.	HP-UX, AIX, Solaris, Linux,

Field	Title	Unit	Description	OS
				Windows
total	Disk IO Total	bytes/s	This value represents the total throughput of the physical disk per second, which is the sum of the read and write byte throughput.	HP-UX, AIX, Solaris, Linux, Windows
read	Disk Read	bytes/s	This value indicates the read throughput of the physical disk in bytes per second. It is interpreted based on Bytes/sec.	HP-UX, AIX, Solaris, Linux, Windows
write	Disk Write	bytes/s	This value represents the write throughput of the physical disk in bytes per second. It is interpreted based on Bytes/sec.	HP-UX, AIX, Solaris, Linux, Windows
totalIops	Disk Total IOPS	counts/s	This value represents the total I/O operations per second for the physical disk. It is the sum of the read IOPS and write IOPS.	HP-UX, AIX, Solaris, Linux, Windows
readIops	Disk Read IOPS	counts/s	This value represents the number of read I/O operations per second for the physical disk.	HP-UX, AIX, Solaris, Linux, Windows
writeIops	Disk Write IOPS	counts/s	This value represents the number of write I/O operations per second for the physical disk.	HP-UX, AIX, Solaris, Linux,

Field	Title	Unit	Description	OS
				Windows
readSvcTime	Disk Read Service Time	ms/io	The value displayed in milliseconds represents the time taken by the physical disk to process a single Read I/O request.	AIX, Solaris, Linux, Windows
writeSvcTime	Disk Write Service Time	ms/io	The value displayed in milliseconds represents the time taken by the physical disk to process a single Write IO request.	AIX, Solaris, Linux, Windows
avgWaitTime	Disk Average Wait Time	ms	This is a metric exclusive to HP-UX that indicates the average time, in milliseconds, that IO processing requests have spent in an idle state in the disk queue.	HP-UX
avgSvcTime	Disk Average Service Time	ms	This is a value that indicates the average time in milliseconds taken to process each transfer request to the disk (including seek time, rotational latency, and data transfer time) for HP-UX.	HP-UX
qLength	Disk Q-Length	count	This metric indicates the queue length of I/O requests that are either pending or being processed for the specified disk. A low value suggests that requests are being handled smoothly, while a consistently high value may indicate that the disk is unable to process requests in a timely manner. However, the absolute benchmark may vary depending on the type of disk (HDD/SSD/NVMe), system architecture, and workload characteristics, so it should be interpreted comprehensively along with busy (%), wait time, and service time.	Linux, Windows

infra_filesystem

A category that collects capacity and performance information per filesystem or mount point. It provides total capacity, usage, free space, usage rate, inode usage rate, read/write throughput, IOPS, queue length, etc. Use it to check a specific filesystem's space shortage, inode exhaustion, and I/O load.

- Category: by Filesystem
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
mountName	Mount Name	string	This value represents the path or identifier name where the file system is mounted. For example, in Linux/Unix, mount points such as /, /var, and /data may be displayed, while in Windows, drive letters or volume paths may be shown.	HP-UX, AIX, Solaris, Linux, Windows
fsType	FileSystem Type	string	This value indicates the file system type. For example, ext4, xfs, zfs, vxfs, ntfs, etc.	HP-UX, AIX, Solaris, Linux, Windows
devName	Device Name	string	This value represents the name of the device or the device identification path to which the file system is actually connected. In Linux, it is displayed based on the KNAME column of the lsblk command, while in Unix-like systems, it is based on the FILESYSTEM column of the df or bdf command results. In Windows, it is displayed based on the device identification path to which the file system is connected, rather than the drive letter.	HP-UX, AIX, Solaris, Linux, Windows
pDevName	Parent Device Name	string	This value indicates the name of the parent device to which the device belongs. In Linux, it is displayed based on the PKNAME value of lsblk, and can be used to identify which higher disk hierarchy a partition or logical device is connected to. For example, it may be displayed as	Linux

Field	Title	Unit	Description	OS
			nvme0n1 for nvme0n1p1, and as sda for sda1. In Unix-like systems and Windows, it is difficult to provide information based on the same criteria as Linux's PKNAME in a consistent manner, so it may be empty or excluded from the collection target.	
totalSize	FS Total Size	GiB	The total capacity of the file system expressed in bytes.	HP-UX, AIX, Solaris, Linux, Windows
used	FS Used	GiB	The value representing the currently used capacity in bytes for the file system.	HP-UX, AIX, Solaris, Linux, Windows
free	FS Free	GiB	The value representing the currently available remaining capacity in the file system, in bytes.	HP-UX, AIX, Solaris, Linux, Windows
pUsed	FS Used	%	This value represents the ratio of used capacity to the total capacity of the file system.	HP-UX, AIX, Solaris, Linux, Windows
pFree	FS Free	%	This value represents the percentage of available remaining capacity compared to the total capacity of the file system.	HP-UX, AIX, Solaris, Linux, Windows

Field	Title	Unit	Description	OS
pInodeUsed	FS i-node Used	%	This value represents the ratio of used inodes to the total number of inodes in the file system. When the number of files is very high, it may not be possible to create new files due to inode exhaustion, even if there is available capacity, so it is necessary to check this as well.	HP-UX, AIX, Solaris, Linux
read	FS Read	bytes/s	This value represents the data throughput read per second from the file system, expressed in Byte/s. In other words, it is the read throughput based on the file system.	HHP-UX, AIX, Solaris, Linux, Windows
write	FS Write	bytes/s	This value represents the data throughput recorded per second on the file system, expressed in Byte/s. In other words, it is the write throughput based on the file system.	HP-UX, AIX, Solaris, Linux, Windows
ioTotal	FS IO Total	bytes/s	This value represents the total data throughput processed per second in the file system, measured in Byte/s. It is the sum of the read and write data throughput, indicating the total I/O throughput of the file system.	HP-UX, AIX, Solaris, Linux, Windows
readIops	FS Read IOPS	counts/s	This value represents the number of read I/O operations processed per second on the file system. It is interpreted as the file system's read IOPS.	HP-UX, AIX, Solaris, Linux, Windows
writeIops	FS Write IOPS	counts/s	This value represents the number of write I/O operations processed per second on the file system. It is interpreted as file system write IOPS.	HP-UX, AIX, Solaris, Linux, Windows

Field	Title	Unit	Description	OS
IopsTotal	FS Total IOPS	counts/s	This value represents the total number of I/O operations processed per second on the file system. It is the sum of read IOPS and write IOPS, indicating the total IOPS for the file system.	HP-UX, AIX, Solaris, Linux, Windows
iopsTotal	FS Total IOPS	counts/s	This value represents the total number of I/O operations processed per second on the file system. It is the sum of read IOPS and write IOPS, indicating the total IOPS for the file system.	HP-UX, AIX, Solaris, Linux, Windows
qLength	FS Q-Length	count	This metric indicates the queue length level of I/O processing requests associated with the file system. It can be interpreted as the average cumulative level of requests that are either waiting or being processed. A temporary spike in the value may be normal due to a momentary increase in load; however, if a high state persists, it may suggest potential processing delays in the file system or underlying storage. It is important to note that absolute benchmark values can vary depending on the type of file system, storage architecture, and workload characteristics, so it should be interpreted comprehensively alongside metrics such as read/write throughput, IOPS, and response time series.	Linux, Windows

infra_zpool_perf

A category that collects performance/state information per ZFS zpool and vdev. It provides zpool capacity, read/write throughput, IOPS, queue and disk wait family indicators, and wait times related to scrub/trim/rebuild. Use it to check the overall state of the ZFS storage pool and bottleneck sections of a specific vdev or device.

- Category: by ZFS
- Default collection interval: 5 sec
- Minimum supported agent: Linux 2.9.15+, Solaris 1.4.9+

Field	Title	Unit	Description	OS
name	Zpool Display Name	string	This is the display name of the device. It is an identifier value that combines the pool name and the device name, used as a display name to distinguish between zpool or vdev on the screen.	Linux, Solaris
poolName	Zpool Name	string	This value represents the name of the ZFS zpool to which the device belongs. It is used when interpreting multiple vdevs or leaf devices grouped within the same pool.	Linux, Solaris
devName	Zpool Device Name	string	This value represents the original name of the actual vdev or leaf device. It is used to identify devices where performance bottlenecks have occurred.	Linux, Solaris
devType	Zpool Device Type	string	This value indicates the ZFS device type. For example, it can be represented as zpool, data, special, log, cache, spare, etc., and is used to distinguish the roles of the devices.	Linux, Solaris
totalSize	Total Size	byte	This is the total size of the zpool or device (Allocated + Free) expressed in bytes.	Linux, Solaris
allocated	Allocated	byte	This value represents the currently allocated capacity for the zpool or device in bytes. It is used to understand the scale of space that has been logically assigned.	Linux, Solaris
free	Free	byte	This value represents the currently available remaining capacity in bytes for the specified zpool or device. It is used to check for insufficient available space.	Linux, Solaris
pctUsed	Used	%	This value represents the percentage of used capacity relative to the total capacity of the zpool or device. A higher value indicates less available space.	Linux, Solaris

Field	Title	Unit	Description	OS
<code>readBytes</code>	Read	bytes/s	This value represents the read data throughput of the specified zpool or device in bytes/s. It indicates the ZFS read throughput metric.	Linux, Solaris
<code>writeBytes</code>	Write	bytes/s	This value represents the write data throughput of the specified zpool or device in bytes/s. It indicates the ZFS standard write throughput metric.	Linux, Solaris
<code>totalBytes</code>	Read+Write	bytes/s	This value represents the total data throughput of the specified zpool or device in bytes/s. It is the total throughput metric obtained by summing the read and write throughput.	Linux, Solaris
<code>readOperation</code>	Read IOPS	counts/s	This value represents the number of read I/O operations processed per second on the specified zpool or device. It refers to the ZFS standard read IOPS.	Linux, Solaris
<code>writeOperation</code>	Write IOPS	counts/s	This value represents the number of write I/O operations processed per second on the specified zpool or device. It refers to the ZFS standard write IOPS.	Linux, Solaris
<code>totalOperation</code>	Read+Write IOPS	counts/s	This value represents the total number of I/O operations processed per second on the specified zpool or device. It is the total IOPS metric obtained by summing the read IOPS and write IOPS.	Linux, Solaris
<code>asyncqReadWait_ns</code>	Asyncq Read Wait	ns	The average time in nanoseconds that read I/O requests have waited in the ZFS internal asynchronous queue. A higher value indicates that asynchronous read requests are being delayed in the queue.	Linux
<code>asyncqWriteWait_ns</code>	Asyncq Write Wait	ns	The average time in nanoseconds that write I/O requests wait in the ZFS internal asynchronous	Linux

Field	Title	Unit	Description	OS
			queue. A higher value indicates that asynchronous write requests are being delayed in the queue.	
totalAsyncqWait_ns	Asyncq Total Wait	ns	The total average time in nanoseconds that read/write I/O requests waited in the ZFS internal asynchronous queue. This is used to comprehensively check the bottleneck level of the async queue.	Linux
syncqReadWait_ns	Syncq Read Wait	ns	The average time in nanoseconds that read I/O requests waited in the ZFS internal synchronous queue. This is used to assess the latency level of synchronous read request processing.	Linux
syncqWriteWait_ns	Syncq Write Wait	ns	The average time in nanoseconds that write I/O requests waited in the ZFS internal synchronous queue. This is used to assess the level of latency in processing synchronous write requests.	Linux
totalSyncqWait_ns	Syncq Total Wait	ns	The total average time in nanoseconds that read/write I/O requests have waited in the ZFS internal synchronous queue. This is used to comprehensively check for sync queue bottlenecks.	Linux
diskReadWait_ns	Disk Read Wait	ns	The average time taken for read I/O requests to be processed at the actual disk layer, expressed in nanoseconds. This is used to understand the actual device processing delay level after queue waiting.	Linux
diskWriteWait_ns	Disk Write Wait	ns	The value represents the average time taken for write I/O requests to be processed at the actual disk layer, measured in nanoseconds. It is used to understand the actual device processing latency level after queue waiting.	Linux
totalDiskWait_ns	Disk Total	ns	The value represents the total average time in	Linux

Field	Title	Unit	Description	OS
	Wait		nanoseconds for read/write I/O requests to be processed at the actual disk layer. It is used to comprehensively assess the processing latency level of the device itself.	
totalReadWait_ns	Total Read Wait	ns	The total average latency of read I/O requests expressed in nanoseconds. It includes both the internal queue wait time of ZFS and the actual disk processing time.	Linux
totalWriteWait_ns	Total Write Wait	ns	The value representing the overall average latency of write I/O requests in nanoseconds. It includes both the internal queue wait time of ZFS and the actual disk processing time.	Linux
totalWait_ns	Total Wait	ns	This value represents the average latency of all read/write I/O requests in nanoseconds. It is a consolidated wait metric that includes both queue waiting time and actual disk processing time. A higher value indicates a potential for I/O latency for the corresponding zpool or device.	Linux
scrub_ns	Scrub	ns	The average time (in ns) that scrub operations wait in the ZFS internal queue. Scrub is the process of sequentially reading blocks within the pool to verify checksums and check for data anomalies.	Linux
trim_ns	Trim	ns	The average time that trim operations wait in the ZFS internal queue is represented in nanoseconds. Trim is the operation that returns unused blocks from SSDs or thin provisioning devices.	Linux
rebuild_ns	Rebuild	ns	The average time (in ns) that the rebuild operation waits in the ZFS internal queue. It indicates the level of delay in the process of reconstructing the necessary data during device replacement or	Linux

Field	Title	Unit	Description	OS
			recovery.	

infra_vxlv_perf

Data collecting I/O performance information per Logical Volume (vxlv) in a Veritas Volume Manager environment. It provides indicators such as read/write IOPS, throughput, and average latency per disk group and volume, and is used to analyze per-volume performance trends, bottlenecks, and read/write skew.

- Category: by VxLV
- Default collection interval: 5 sec
- Minimum supported agent: Linux 2.9.15+

Field	Title	Unit	Description	OS
name	VxLV Display Name	string	This is the display name. It can be used as an identifier combining the Disk Group name and the VxLvol name.	Linux
diskGroup	VxLV Disk Group	string	The name of the Veritas Disk Group to which the volume belongs. It is a management unit that logically groups multiple volumes and is used to identify or aggregate volumes within the same Disk Group.	Linux
vxLvolName	VxLV Lvol Name	string	The name of the logical volume managed by Veritas Volume Manager. It is the volume identifier that is the subject of performance metrics measurement.	Linux
blockSize	VxLV Block Size	string	This value represents the I/O block size of the volume. It is expressed in bytes and is used to convert block count-based metrics into byte-based throughput. It is based on the output of the command 'vxprint -F %iosize', and in case of a retrieval failure or abnormal value, it falls back to 512 bytes.	Linux
readIops	Read	counts/s	The average number of read operations per second is calculated	Linux

Field	Title	Unit	Description	OS
	IOPS		by dividing the number of read operations that occurred during the collection interval by the duration of the interval. This indicates the frequency of read request processing for the volume.	
writeIops	Write IOPS	counts/s	The average number of write operations per second is calculated by dividing the number of write operations that occurred during the collection interval by the duration of the interval. This indicates the frequency of write requests to the volume.	Linux
IopsTotal	Total IOPS	counts/s	The average number of operations per second is calculated by dividing the total number of I/O operations that occurred during the collection interval by the duration of the interval. It is interpreted as the sum of Read IOPS and Write IOPS.	Linux
read	Read	bytes/s	The average read throughput per second is calculated by multiplying the number of read blocks processed during the collection interval by the I/O block size and then dividing by the interval time. The unit is bytes/s.	Linux
write	Write	bytes/s	The average write throughput per second is calculated by multiplying the number of write blocks processed during the collection interval by the I/O block size and then dividing by the interval duration. The unit is bytes/s.	Linux
total	IO Total	bytes/s	The average throughput processed during the collection interval, calculated by dividing the total I/O data amount by the interval time. It refers to the sum of read throughput and write throughput.	Linux
readAvgMs	Read Average Latency	ms	Average latency of read operations during the collection interval. This is the average time taken for each read request to complete, measured in milliseconds (ms).	Linux
writeAvgMs	Write Average	ms	The average latency of write operations that occurred during the collection interval. It is the average time taken for each write	Linux

Field	Title	Unit	Description	OS
	Latency		request to complete, measured in milliseconds (ms).	

infra_cpu_core

A category that collects the usage state per logical CPU core. It provides user, sys, wait, idle, irq, softirq, steal, core CPU usage, etc. by coreId. Use it to check core skew, IRQ concentration, and single-thread bottlenecks.

- Category: by CPU Core
- Default collection interval: 5 sec

Field	Title	Unit	Description	OS
coreId	Core ID	string	This is the ID value used to identify CPU cores. It refers to the logical core number recognized by the OS and may not correspond 1:1 with physical cores. When Hyper-Threading/SMT is enabled, multiple logical cores can be mapped to a single physical core.	HP-UX, AIX, Solaris, Linux, Windows
coreCpu	CPU	%	This refers to the overall CPU utilization of the core. It is calculated as 100 minus idle and interpreted as the usage ratio that includes all non-idle times such as user mode, kernel mode, nice, I/O wait, irq, softirq, and steal. The included items may vary depending on the operating system.	HP-UX, AIX, Solaris, Linux, Windows
user	User	%	It refers to the ratio of the total CPU time of the core that is used to execute general applications/processes in user mode.	HP-UX, AIX, Solaris, Linux, Windows
sys	Sys	%	It refers to the ratio of the total CPU time of the core that is used for system call processing, kernel tasks, and driver handling in kernel mode (System mode).	HP-UX, AIX, Solaris, Linux,

Field	Title	Unit	Description	OS
				Windows
nice	Nice	%	It refers to the ratio of the total CPU time of the core that is used for executing low-priority (niced) user processes. The utilization of user processes with normal priority is typically included in user(%).	HP-UX
wait	Wait	%	The ratio of the total CPU time of the core during which the CPU was idle waiting for I/O completion (I/O wait). This is interpreted as the time spent waiting for responses from external resources such as disk/network, rather than the time the CPU actually performed computations.	HP-UX, AIX, Solaris, Linux
total	User+Sys	%	It refers to the sum of the ratios of CPU usage in user mode and kernel mode out of the total CPU utilization of the respective core. This is interpreted as the pure CPU execution utilization, excluding I/O wait, irq, softirq, and steal.	HP-UX, AIX, Solaris, Linux, Windows
total_wait	User+Sys+Wait	%	The sum of user, system, and wait time of the total CPU time for the corresponding core. In other words, it is a usage metric that includes user processing, kernel processing, and I/O waiting. There may be differences in the actual total CPU utilization depending on the operating system.	HP-UX, AIX, Solaris, Linux
idle	Idle	%	This refers to the ratio of the total CPU time of the core that was idle without performing any tasks. A higher value indicates that the core has more availability.	HP-UX, AIX, Solaris, Linux, Windows
irq	Irq	%	This refers to the ratio of total CPU time of the core that is used to handle hardware interrupts (IRQ). It can increase when network cards, disks, and other hardware events are concentrated on a specific core.	Linux

Field	Title	Unit	Description	OS
softirq	Softirq	%	This refers to the ratio of total CPU time of the core used for handling Software interrupts (SoftIRQ). It can increase when delayed processing tasks of the kernel, such as network packet processing and block I/O post-processing, are concentrated on a specific core.	Linux
steal	Steal	%	In a virtualized environment, the ratio of time that a core is unable to execute due to the hypervisor prioritizing other virtual CPUs is referred to as CPU steal utilization. A high value may indicate insufficient CPU resources allocated to the VM or a contention state.	Linux

Field	Title	Unit	Description	OS
numaNodeId	NUMA Node ID	string	NUMA node number the logical core belongs to. -1 for cores whose NUMA connection information cannot be determined	Linux

The CPU Core chart legend of Metrics Explorer uses this field to display the `Core cpu1 (Node 0)` form. For cores whose value is `-1` and in Windows and Unix environments where the field is not collected, the `Core cpu1` form is kept.

infra_os_parameter

A category that collects the kernel and system parameter values of the operating system. It provides the main parameters of the network, file system, and kernel areas along with their current values. Use it to compare configuration differences between servers or to check kernel setting values that can cause performance problems. The collected values can be queried on the [OS Parameter information](#) screen in Labs.

- Category: OS Parameter
- Default collection interval: 1 hour

Field	Title	Unit	Description	OS
<code>parameter</code>	Parameter	string	Parameter name (for example, <code>net.ipv4.tcp_keepalive_intvl</code>)	AIX, HP-UX, Solaris, Linux
<code>category</code>	Category	string	Area the parameter belongs to (for example, <code>net</code> , <code>fs</code> , <code>kernel</code>)	AIX, HP-UX, Solaris, Linux
<code>os</code>	OS	string	Type of the operating system the parameter was collected from	AIX, HP-UX, Solaris, Linux
<code>value</code>	Value	string	Current value of the parameter	AIX, HP-UX, Solaris, Linux
<code>custom</code>	Custom	string	Whether the parameter was added by the user with the <code>os.parameter.custom</code> option instead of being a default collection target	AIX, HP-UX, Solaris, Linux

To add parameters other than the default collection targets, enter the parameter names separated by commas (,) in the `os.parameter.custom` option. Adjust the collection interval with the `infra_os_parameter.interval` option. The default value is 3600 seconds.

Windows does not collect this category.

References

- [Windows reference](#)
- [Linux reference](#)

Server Inventory Collection Items

Apart from the performance metrics, the server monitoring agent periodically collects the hardware configuration of the server. Performance metrics alone do not show which device is connected where, so even when a metric looks abnormal, it is difficult to identify the device causing it. Viewing the inventory items together shows which NUMA node the GPU and the network card are attached to, and how the GPUs are connected to each other. This document describes the collected fields of each inventory item, for operators who run AI and HPC workloads and need to check the device configuration as well.

Inventory items are sent only when a value changes. The agent checks the configuration every hour and sends only the items that were added, changed, or deleted compared with the previous collection.

Fields marked with a version and **New** in the **Remarks** column of each table were added in that version. Fields without a mark have been provided since the category was added.

server_rdma_inventory

Configuration information of RDMA (InfiniBand, RoCE) devices and ports. One row is collected per port of an RDMA device, and it provides everything from the device model, firmware, and link speed to the connected operating system network interface.

- Collection target: Linux
- Default collection interval: 1 hour

Field	Type	Unit	Description
interfaceName	string	-	RDMA device name and port number (for example, <code>mlx5_0:1</code>)
rdmaDevice	string	-	RDMA device name (for example, <code>mlx5_0</code>)
rdmaNetdev	string	-	Operating system network interface connected to the device. Multiple names are separated by <code> </code>
rdmaNetdevIp	string	-	IP address per network interface (CIDR notation). Multiple addresses are separated by <code> </code>

Field	Type	Unit	Description
rdmaDeviceType	string	-	ASIC type of the device (for example, MT4123)
rdmaDeviceModel	string	-	Device model name (for example, ConnectX-6)
rdmaLinkLayer	string	-	Link layer. InfiniBand or Ethernet
rdmaNumaNode	string	-	NUMA node number the device belongs to. -1 when it cannot be determined
rdmaPciBusId	string	-	PCI address (for example, 0000:81:00.0)
rdmaPort	string	-	Port number
rdmaRate	string	-	Port link speed (for example, 100 Gb/sec (4X EDR))
rdmaMtu	string	-	Active MTU of the port
rdmaNetdevMtu	string	-	MTU per network interface. Multiple values are separated by
rdmaDriverName	string	-	Kernel driver name (for example, mlx5_core)
rdmaDriverVersion	string	-	Driver version
rdmaFirmwareVersion	string	-	Device firmware version
rdmaBondMaster	string	-	Name of the bond the network interface belongs to
rdmaNodeGuid	string	-	Node GUID of the device
rdmaPortGuid	string	-	Port GUID of the port
rdmaSysImageGuid	string	-	System Image GUID of the device
rdmaSubnetPrefix	string	-	Subnet prefix the port belongs to
rdmaVlanId	string	-	VLAN ID per network interface

Field	Type	Unit	Description
<code>storageMount</code>	<code>string</code>	-	Remote storage mount point connected through that network interface

Adjust whether to collect and the interval with the following options.

whatap.conf

```
rdma.inventory.enabled=true # Whether to collect (Default: true)
rdma.inventory.interval=3600 # Collection interval in seconds (Default: 3600)
```

server_gpu_topology

The connection relationship between two GPUs. One row is collected per physical GPU pair. Use it to check by which path two GPUs are connected and whether they can exchange data directly.

- Collection target: Linux
- Default collection interval: 1 hour

Field	Type	Unit	Description
<code>baseGpuIndex</code>	<code>string</code>	-	Base GPU number
<code>targetGpuIndex</code>	<code>string</code>	-	Peer GPU number. Always larger than the base GPU number
<code>topologyType</code>	<code>string</code>	-	Connection level between the two GPUs. One of <code>PIX</code> , <code>PXB</code> , <code>PHB</code> , <code>NODE</code> , <code>SYS</code> , or <code>NVn</code> . <code>INCONSISTENT</code> when the values are not consistent
<code>topologyConsistent</code>	<code>string</code>	-	Whether the results queried in both directions match. <code>true</code> or <code>false</code>
<code>interconnectType</code>	<code>string</code>	-	Connection method. <code>NVSwitch</code> , <code>NVLink</code> , <code>PCIe</code> , or <code>Unknown</code>
<code>nvLinkCount</code>	-	<code>count</code>	Number of NVLink connections. <code>0</code> when it is not an NVLink connection

Field	Type	Unit	Description
baseGpuMigMode	string	-	MIG configuration status of the base GPU. <code>Enabled</code> , <code>Disabled</code> , <code>NotSupported</code> , or <code>Unknown</code>
targetGpuMigMode	string	-	MIG configuration status of the peer GPU
peerToPeerSupported	string	-	Whether direct communication (P2P) between the two GPUs is supported. <code>true</code> , <code>false</code> , or <code>unknown</code>
peerToPeerStatus	string	-	P2P read status code (for example, <code>OK</code> , <code>NS</code> , <code>CNS</code>)

The closer the connection level is to `PIX`, the physically closer the two GPUs are; the closer it is to `SYS`, the farther apart they are. Grouping GPUs that handle the same job into close pairs reduces communication latency.

whatap.conf

```
gpu.topology.inventory.enabled=true # Whether to collect (Default: true)
gpu.topology.inventory.interval=3600 # Collection interval in seconds (Default: 3600)
```

❗ If GPU collection is disabled or the GPU driver cannot be accessed, this item is not collected. On a server with a single GPU, no pair is formed, so the collection result is empty.

server_system_topology

Configuration information that links CPU sockets, NUMA nodes, and PCIe devices in parent-child relationships. One connection relationship becomes one row. Use it to view the device layout of the whole server as a tree.

- Collection target: Linux, Windows
- Default collection interval: 1 hour

Field	Type	Unit	Description
parentDevice	string	-	Name of the parent device (for example, <code>SYSTEM</code> , <code>NUMA_NODE:0</code> ,

Field	Type	Unit	Description
			PCI_BRIDGE:0000:00:01.1)
childDevice	string	-	Name of the child device (for example, CPU_SOCKET:0, GPU0, ens192)
deviceType	string	-	Type of the child device. CPU_SOCKET, NUMA_NODE, CPU_CORES, MEMORY, PCI_ROOT, GPU, NVSWITCH, RDMA_NIC, NIC, NVME, PCI_BRIDGE, HOST_BRIDGE, or PCI_DEVICE
pciBusId	string	-	PCI address for devices in the PCI layer, and an identifier in the socket:0 or node:0 form for devices in the CPU layer
pcieDepth	string	-	Depth of the PCIe path. 0 for the CPU, NUMA, and memory layers

The name of the top of the tree is `SYSTEM`. The collection uses only the information provided by the kernel (sysfs on Linux, device information on Windows), so it does not run any separate command.

whatap.conf

```
system.topology.enabled=true # Whether to collect (Default: true)
system.topology.interval=3600 # Collection interval in seconds (Default: 3600)
```

server_numa_inventory

Configuration information per NUMA node. One NUMA node becomes one row, and it provides the CPUs and memory that belong to the node, together with the GPU, RDMA, and NVMe devices attached to the same node.

- Collection target: Linux
- Default collection interval: 1 hour

Field	Type	Unit	Description	Remarks
numaNodeId	string	-	NUMA node number	
cpuList	string	-	List of CPUs that belong to the node (for example,	

Field	Type	Unit	Description	Remarks
			0-7,16-23)	
logicalCores	-	count	Number of logical CPUs that belong to the node	
memorySize	-	byte	Memory size that belongs to the node	
gpuIndexes	string	-	List of NVIDIA GPU indexes attached to the node. Separated by a comma and a space	
rdmaDevices	string	-	List of RDMA device names attached to the node (for example, mlx5_0, mlx5_1). Separated by a comma and a space	
rdmaNetDevs	string	-	List of network interfaces bound to those RDMA devices. Separated by a comma and a space	
nvmeDevices	string	-	List of NVMe device names connected to that NUMA node (for example, nvme0, nvme1). Separated by a comma and a space	New in 2.9.22
gpuRdmaAffinity	string	-	The closest RDMA device for each GPU of that NUMA node, displayed in the GPU0:mlx5_0(PIX) form. Multiple entries are separated by a comma and a space	New in 2.9.22

The value in parentheses in `gpuRdmaAffinity` is the distance between the GPU and the RDMA device. From closest to farthest, the order is `PIX`, `PXB`, `PHB`, `NODE`, and `SYS`, and only the device at the closest level is displayed for each GPU. Use it to check whether the two devices are attached to the same node in configurations where GPU communication goes through an RDMA device.

Adjust whether to collect and the interval with the following options.

whatap.conf

```

numa.inventory.enabled=true # Whether to collect (Default: true)
numa.inventory.interval=3600 # Collection interval in seconds (Default: 3600)

```

VMware Collection Items

The server monitoring agent collects the status of ESXi hosts and virtual machines together through the VMware API. The resource usage seen from inside a virtual machine alone makes it difficult to identify the cause of performance degradation. A broken fan on the host, or virtual machines packed onto one host so that vCPUs wait for each other's turn, hardly shows up in the guest metrics. This document describes the items collected through the VMware integration, for operators who run virtualization-based infrastructure and need to trace causes down to the host level.

The VMware data integration can collect data by connecting directly to an ESXi host, even without vCenter. However, the `vm_vcenter_inventory` category is collected only when the agent is connected to vCenter.

The supported agent versions are Linux and Windows 2.9.23 or later. Fields marked `New in 2.9.23` in the **Remarks** column of each table were added in 2.9.23. Fields without a mark have been provided since earlier versions.

infra_vm_vhost_hw

A category that collects the hardware status of an ESXi host by part. One host becomes one row, and it provides the status per part, such as power, cooling, memory, and processor, along with the highest temperature. Use it to check hardware failures before they lead to performance degradation or an outage. This category was added in 2.9.23, so every field in the table is new.

- Collection target: VMware ESXi

Status values are standardized into the following four. When a part has multiple sensors, the most severe status is displayed as the representative value, and a part with no sensor data is displayed as `unknown`.

Item	Description
<code>green</code>	Normal
<code>yellow</code>	Warning
<code>red</code>	Failure

Item	Description
unknown	The status cannot be determined

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each hardware status came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
hardwareHealth	string	-	Overall status representing the most severe state among all collectable hardware statuses
powerSupplyStatus	string	-	Status of the power supply units
powerSupplyRedundancyStatus	string	-	Redundancy status of the power supply units. Used to check whether redundancy is degraded by a partial failure
networkDeviceStatus	string	-	Hardware status of the physical network devices and network adapters. Distinct from the network link up and down state
batteryStatus	string	-	Status of the internal battery or the hardware management controller battery. Includes CMOS batteries and RAID cache batteries
processorsStatus	string	-	Hardware status of the physical processors and CPU sockets. A physical hardware status unrelated to CPU usage
memoryModulesStatus	string	-	Hardware status of the physical memory modules. Includes DIMM errors and module failure or warning

Field	Type	Unit	Description
			states
coolingUnitStatus	string	-	Status of the cooling units such as fans and blowers
coolingUnitRedundancyStatus	string	-	Redundancy status of the cooling units. Used to check whether cooling is maintained even when some fans fail
temperatureStatus	string	-	Overall status of the temperature sensors such as the CPU, system board, memory, and air intake
temperatureMax	-	°C	Highest temperature among the sensor readings at collection time. Calculated by applying the unit modifier provided by VMware to the raw sensor reading. Empty when no temperature sensor is available
hardwareHealthDataAvailable	boolean	-	Whether the host reported at least one hardware status or sensor data
hardwareHealthSupportStatus	string	-	Whether hardware status collection is supported. <code>supported</code> means sensor data is provided, <code>unsupported</code> means the API is not supported or the sensors are not exposed
hardwareHealthUnavailableReason	string	-	Reason why hardware status data cannot be collected. Records causes such as an unsupported Hardware Health API, unexposed sensors, or a collection error. Empty on normal collection
temperatureMaxSensor	string	-	Name of the temperature sensor where <code>temperatureMax</code> was measured. Identifies which of the CPU, system board, memory, or air intake sensors actually recorded the highest temperature
temperatureUnit	string	-	Unit applied to <code>temperatureMax</code> . Usually in the

Field	Type	Unit	Description
			Degrees C form, and the expression can vary depending on the sensor provider

infra_vm_vhost_resource_capacity

A category that collects the physical resources of an ESXi host together with the resources allocated to virtual machines. One host becomes one row, and it provides the ratio of allocation to physical capacity (overcommit). Use it to check whether too many virtual machines are placed on a specific host. This category was added in 2.9.23, so every field in the table is new.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each resource metric came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
cpuCores	-	count	Total number of physical CPU cores on the host. The basis for allocatable virtual machine resources and CPU capacity planning
totalMemory	-	byte	Total memory capacity installed on the host. The basis for the capacity available for virtual machine placement and memory resource planning
activeVCpu	-	count	Sum of the vCPUs allocated to powered-on virtual machines
allocatedVCpu	-	count	Sum of the vCPUs configured for all virtual machines placed on the host. Includes powered-off and suspended virtual machines

Field	Type	Unit	Description
activeOverCommitCore	-	-	activeVCpu divided by the number of physical CPU cores. Expressed as a multiple showing how many vCPUs are allocated per physical core on average, based on running virtual machines
allocatedOverCommitCore	-	-	allocatedVCpu divided by the number of physical CPU cores. The vCPU overcommit multiple on a configuration basis, including powered-off and suspended virtual machines
activeMemory	-	byte	Sum of the memory configured for powered-on virtual machines
allocatedMemory	-	byte	Sum of the memory configured for all virtual machines placed on the host. Includes powered-off and suspended virtual machines
activeOverCommitMemory	-	-	activeMemory divided by the physical memory capacity. Expressed as a multiple showing how much memory is allocated relative to the physical memory, based on running virtual machines
allocatedOverCommitMemory	-	-	allocatedMemory divided by the physical memory capacity. The memory overcommit multiple on a configuration basis, including powered-off and suspended virtual machines
vmGuestCount	-	count	Number of virtual machine objects connected to the host. Used to identify the placement and density per host
poweredOffVMCount	-	count	Number of powered-off virtual machines. Included in allocatedVCpu but not in activeVCpu
suspendedVMCount	-	count	Number of suspended virtual machines. Because vCPUs are configured, they are included in allocatedVCpu, but they are not included in activeVCpu since they are not running

Field	Type	Unit	Description
unknownPowerStateVMCount	-	count	Number of virtual machines whose power state cannot be determined or is not identified as one of <code>poweredOn</code> , <code>poweredOff</code> , or <code>suspended</code> . Can occur due to vCenter response delays, permissions, or a temporary state mismatch
vcpuSkippedVMCount	-	count	Number of virtual machines excluded from the vCPU sum because the vCPU setting could not be read. Can be caused by a failure to access the configuration, insufficient permissions, an abnormal configuration, or a temporary vCenter response error
memorySkippedVMCount	-	count	Number of virtual machines excluded from the memory sum because the memory setting could not be read. The causes are the same as for <code>vcpuSkippedVMCount</code>

The overcommit ratio is displayed to two decimal places. A value greater than `1` means more resources are allocated to virtual machines than the physical resources. Since a certain level of overcommit is common in virtualized environments, it is better to look at the deviation between hosts in the same cluster and the increasing trend over time rather than the absolute value.

Adjust whether to collect with the following option.

whatap.conf

```
vm.vhost.perf.resource_capacity.enabled=true # Whether to collect (Default: true)
```

infra_vm_vhost_cpu

A category that aggregates the CPU usage and wait time of an ESXi host. One host becomes one row, and it provides the CPU usage and contention of all virtual machines on the host.

- Collection target: VMware ESXi

Field	Type	Unit	Description	Remarks
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each CPU metric came from	
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes	
usage	-	percent	Ratio of the host's total CPU resources actually used. The basic metric for checking the CPU load level of the host	
usageMhz	-	MHz	Amount of CPU resources actually in use on the host. Check it together with <code>usage</code> when comparing hosts with different CPU capacities	
totalCapacity	-	MHz	Total CPU capacity the host can provide. The reference value for interpreting usage, demand, and reserved capacity	
demand	-	MHz	Amount of CPU resources the virtual machines of the host actually demanded	
readiness	-	percent	Ratio of the time a virtual machine was ready to run but was not immediately given a physical CPU. The key metric for judging CPU contention on the host	
ready	-	ms/s	Time spent waiting for a physical CPU while ready to run	
latency	-	percent	Level of delay between requesting CPU resources and actually running	
costop	-	ms/s	Time a multi-vCPU virtual machine was stopped while waiting to be scheduled simultaneously	
wait	-	ms/s	Time the CPU spent waiting without performing actual computation. Includes idle time, I/O waits, and swap waits	

Field	Type	Unit	Description	Remarks
swapwait	-	ms/s	Time CPU execution was delayed by memory swap processing	
idle	-	ms/s	Time the CPU spent idle	
reservedCapacity	-	MHz	CPU capacity secured by reservations. When it is large relative to the total capacity, the actually available CPU headroom decreases	
demandPct	-	percent	demand divided by totalCapacity	New in 2.9.23
costopPct	-	percent	costop converted to a ratio (1,000 ms/s = 100%)	New in 2.9.23

Do not judge CPU contention from a single metric. When `usage` is high and `readiness` and `ready` are high together, CPU contention is likely. Conversely, when `usage` is low but `readiness` is high, also check for over-allocated vCPUs, resource limits, and the scheduling policy. When `demand` stays higher than `usageMhz`, the virtual machines may not be getting as much CPU as they demand.

`wait` includes idle time, I/O waits, and swap waits, so a high value does not by itself mean the CPU is insufficient. Check it separately against `idle` and `swapwait`. `swapwait` is displayed as a CPU metric, but the cause can be insufficient memory or memory contention, so when it increases, also check the memory usage, ballooning, and swap activity of the host and the virtual machines.

Metrics in `ms/s` are converted to per-second values so that they can be compared on the same basis even when the collection interval changes.

`demandPct` and `costopPct` are the sum over all virtual machine vCPUs on the host, so they can exceed 100%, and no separate upper limit is applied. If `costopPct` stays high, the number of vCPUs allocated to the virtual machines may be too large for the host capacity.

infra_vm_vhost_mem

A category that aggregates the memory usage and reclamation activity of an ESXi host. One host becomes one

row, and it provides everything from the actual occupancy and working set to reclamation metrics such as ballooning, compression, and swapping.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each memory metric came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
consumed	-	byte	Physical memory the host is actually consuming. Includes the memory occupied by virtual machines, VMkernel, and the ESXi system area
active	-	byte	Memory the virtual machines on the host have actively used recently. Close to the effective working set
activeConsumedRatio	-	percent	<code>active</code> divided by <code>consumed</code> . The share of the consumed memory that has actually been actively used
granted	-	byte	Memory the host granted to the virtual machines
latency	-	percent	Level of delay occurring during memory access or allocation
vmmemctl	-	byte	Memory the balloon driver reclaimed from the virtual machines
swapused	-	byte	Memory that was pushed out to swap on the host
swapInRate	-	bytes/s	Rate at which memory is read back from the swap area
swapOutRate	-	bytes/s	Rate at which memory pages are pushed out to the swap area
compressed	-	byte	Memory compressed through memory compression
compressionRate	-	bytes/s	Rate at which memory compression occurs

Field	Type	Unit	Description
shared	-	byte	Memory shared between virtual machines or between memory pages
overhead	-	byte	Memory overhead ESXi uses additionally to run the virtual machines
sysUsage	-	byte	Memory in use by the ESXi system area
heapFree	-	byte	Free memory available in the ESXi heap area

When `consumed` is high but `active` is low, many virtual machines have a large allocation and occupancy but little actual usage. `activeConsumedRatio` shows this relationship directly.

Do not judge a memory shortage from a single metric. As memory pressure grows, ESXi escalates reclamation through ballooning, then compression, then swapping. When `compressed` and `compressionRate` rise after `vmmemctl` has increased, and `swapped` and `swapOutRate` also appear, memory contention has progressed considerably. When `latency` is high at the same time, virtual machine performance is likely affected.

A high `swapInRate` means the pages a virtual machine needs are in swap and are being read back, which can lead to performance delays.

`shared`, `overhead`, `sysUsage`, and `heapFree` do not indicate a problem on their own. `shared` is a supplementary metric for the memory savings effect, and `overhead` grows with the number of virtual machines and the complexity of their virtual device configurations. If `sysUsage` or `heapFree` changes abnormally, check the ESXi system state separately.

infra_vm_vhost_disk

A category that aggregates the disk I/O of an ESXi host. One host becomes one row, and it provides the throughput summed across all disks connected to the host along with the highest latency.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each disk metric came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
usage	-	bytes/s	Total disk I/O throughput on the host. The sum of reads and writes
read	-	bytes/s	Disk read throughput on the host
write	-	bytes/s	Disk write throughput on the host
maxTotalLatency	-	ms	Highest disk I/O latency among the disks and devices connected to the host

When `usage` stays high, a large amount of disk I/O is occurring on the virtual machines of that host. Check it together with `maxTotalLatency` to distinguish a simple throughput increase from a storage bottleneck.

When `read` is high and the latency increases together, check for read I/O contention, cache misses, and backend storage response delays. When `write` is high and the latency increases, check for write I/O contention, storage controller delays, and RAID, cache, or synchronization delays.

`maxTotalLatency` rises when a delay occurs on a specific disk, datastore, LUN, or storage path. Even when the overall throughput of the host is not high, a high value is worth checking for I/O delays on a particular storage path or virtual machine.

infra_vm_vhost_disk_perf

Performance per disk device connected to an ESXi host. One disk becomes one row, and it provides the throughput, the request counts, and the latency per layer. Use it to narrow down which disk is the bottleneck.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each disk metric came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
diskName	string	-	Name of the disk or device the metrics were collected from. The basis for distinguishing individual disks, LUNs, local disks, and storage devices connected to the host
usage	-	bytes/s	Total I/O throughput of the disk. The sum of read and write
read	-	bytes/s	Read throughput of the disk
write	-	bytes/s	Write throughput of the disk
numberTotal	-	counts/s	Total number of I/O requests. The sum of numberRead and numberWrite
numberRead	-	counts/s	Number of read I/O requests
numberWrite	-	counts/s	Number of write I/O requests
totalLatency	-	ms	Average latency for processing all I/O requests on the disk. The representative latency metric, including the effects of kernel, device, and queue latency
totalReadLatency	-	ms	Average latency for processing read I/O requests
totalWriteLatency	-	ms	Average latency for processing write I/O requests
deviceLatency	-	ms	Average latency for processing I/O on the physical disk or backend storage device
queueLatency	-	ms	Average time an I/O request waited in the queue before being processed

Field	Type	Unit	Description
kernelLatency	-	ms	Average latency incurred while the ESXi VMkernel layer processed the I/O
maxQueueDepth	-	count	Maximum queue depth available or observed on the disk
busResets	-	counts/s	Number of bus resets that occurred on the disk or storage bus

Looking at the throughput (`usage` , `read` , `write`) together with the request counts (`numberTotal` , `numberRead` , `numberWrite`) gives a sense of the I/O block size. When the request count is high but the throughput is low, there are many small-block I/Os.

Splitting the latency metrics narrows down where the delay originates.

High metric	Where to look
deviceLatency	The backend storage layer such as the storage device, LUN, array, SAN, or NAS
kernelLatency	Host-internal I/O processing, the driver, the HBA or storage adapter, and the VMkernel storage stack
queueLatency	Requests are queuing because they are not processed fast enough. Check it with <code>maxQueueDepth</code> and the request counts

When `busResets` keeps occurring with a non-zero value, also check the storage path stability, the HBA and controller logs, and the ESXi events.

The values summed across the whole host are available in `infra_vm_vhost_disk` .

infra_vm_vhost_storageadapter_perf

Performance per storage adapter (HBA) of an ESXi host. One adapter becomes one row, and it provides the summed performance of the disks and LUNs connected to that adapter. Use it to see which HBA carries the load.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each adapter metric came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
memberCount	-	count	Number of canonical disk or LUN instances included in the adapter aggregation. A larger value means the performance of more disks and LUNs is summed into one adapter
memberPathNames	string	-	List of the disk and LUN paths mapped to the adapter. Used to check which paths the adapter values were aggregated from
adapterName	string	-	Name of the storage adapter the values were aggregated for. Usually an HBA or storage controller name such as vmhba0
total	-	bytes/s	Total I/O throughput of the adapter. The sum of read and write
read	-	bytes/s	Data throughput read through the adapter. The sum of the read throughput of the underlying disks and LUNs
write	-	bytes/s	Data throughput written through the adapter. The sum of the write throughput of the underlying disks and LUNs
numberTotalAveraged	-	counts/s	Total average number of I/Os on the adapter. The sum of numberReadAveraged and numberWriteAveraged
numberReadAveraged	-	counts/s	Average number of read I/Os on the adapter
numberWriteAveraged	-	counts/s	Average number of write I/Os on the adapter
commandsAveraged	-	counts/s	Total number of I/O commands issued on the disks and LUNs connected to the adapter
totalReadLatency	-	ms	Average read I/O latency of the adapter. The read latency of the underlying disks and LUNs, weighted by the read request count

Field	Type	Unit	Description
totalWriteLatency	-	ms	Average write I/O latency of the adapter. The write latency of the underlying disks and LUNs, weighted by the write request count

The values in this category are aggregated from the disks and LUNs under the adapter. Throughput and request counts are simple sums, while latency is an average weighted by the request count. A large `memberCount` means that many disks and LUNs are folded into one value, so a single slow disk may not stand out in the adapter average.

When the adapter latency looks high, check `memberPathNames` to see which paths were aggregated, then narrow down per path in `infra_vm_vhost_storagepath_perf` and per disk in `infra_vm_vhost_disk_perf`.

infra_vm_vhost_storagepath_perf

Performance per storage path of an ESXi host. One path becomes one row, and it provides the per-path throughput and latency along with the adapter, target, and LUN information. Use it to find a single slow path in a multipath configuration.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each path metric came from
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
pathName	string	-	Display name of the storage path. A path name combining the adapter, target, LUN, and canonical disk information
canonicalDiskName	string	-	Canonical name of the SCSI LUN as recognized by the host (for example, <code>naa.*</code> , <code>t10.*</code> , <code>mpx.*</code>). The basis for mapping to the performance counter instance

Field	Type	Unit	Description
adapterName	string	-	Name of the storage adapter the path is connected to. Usually an HBA or storage controller name such as vmhba0
targetName	string	-	Name combining the adapter and the target (for example, vmhba0:T0). Used to identify the target path under an adapter
targetId	string	-	Target number the path is connected to in the ESXi SCSI topology. Used to distinguish multiple targets on the same adapter
lunId	string	-	LUN number the path is connected to
adapterKey	string	-	Key that identifies the storage adapter inside vSphere. Used for API mapping and long-term data consistency checks
total	-	bytes/s	Total I/O throughput of the path. The sum of read and write
read	-	bytes/s	Data throughput read through the path
write	-	bytes/s	Data throughput written through the path
numberTotalAveraged	-	counts/s	Total average number of I/Os on the path. The sum of numberReadAveraged and numberWriteAveraged
numberReadAveraged	-	counts/s	Average number of read I/Os on the path
numberWriteAveraged	-	counts/s	Average number of write I/Os on the path
commandsAveraged	-	counts/s	Total number of I/O commands issued on the path. The storage command load including reads and writes
totalReadLatency	-	ms	Average read I/O latency of the path
totalWriteLatency	-	ms	Average write I/O latency of the path

There are several path identification fields because each narrows the problem at a different level. adapterName

is per HBA, `targetId` is per target within the same adapter, and `lunId` is per LUN. When one LUN has multiple paths, comparing which path alone is slow separates a path failure from a storage device problem.

When `totalReadLatency` or `totalWriteLatency` stays high, check the HBA, target, LUN, and SAN segment of that path. If it is high across every path using the same `canonicalDiskName`, the problem is more likely on the backend storage than on the path.

infra_vm_vhost_datastore_perf

Performance and capacity per datastore accessed by an ESXi host. One datastore becomes one row, and it provides the capacity usage along with the I/O throughput and latency.

- Collection target: VMware ESXi

Field	Type	Unit	Description
<code>vhostName</code>	<code>string</code>	-	Name of the host being collected. The basis for identifying which ESXi host each datastore metric came from
<code>uuid</code>	<code>string</code>	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes
<code>datastoreName</code>	<code>string</code>	-	Name of the datastore the performance and capacity data was collected from. The display name vSphere recognizes, such as VMFS, NFS, or vSAN
<code>datastoreInstanceId</code>	<code>string</code>	-	Datastore identifier mapped to the performance counter instance. A value such as the VMFS volume UUID extracted from the datastore URL, used to link data reliably where names can change or be duplicated
<code>datastoreType</code>	<code>string</code>	-	Type of the datastore (for example, <code>VMFS</code> , <code>NFS</code> , <code>vSAN</code>). Used to distinguish the storage backend
<code>accessible</code>	<code>boolean</code>	-	Whether the datastore is accessible from the host
<code>datastoreUrl</code>	<code>string</code>	-	Internal vSphere URL of the datastore. Includes the VMFS

Field	Type	Unit	Description
			volume UUID or the access path
total	-	bytes/s	Total I/O throughput of the datastore. The sum of read and write
read	-	bytes/s	Data throughput read from the datastore
write	-	bytes/s	Data throughput written to the datastore
numberTotalAveraged	-	counts/s	Total average number of I/Os on the datastore. The sum of numberReadAveraged and numberWriteAveraged
numberReadAveraged	-	counts/s	Average number of read I/Os on the datastore
numberWriteAveraged	-	counts/s	Average number of write I/Os on the datastore
totalReadLatency	-	ms	Average read I/O latency of the datastore
totalWriteLatency	-	ms	Average write I/O latency of the datastore
capacity	-	byte	Total capacity of the datastore
allocated	-	byte	Allocated space. The sum of used and uncommitted
used	-	byte	Actual usage. The total capacity minus the free space
uncommitted	-	byte	Space that is committed but not yet actually used in a thin provisioning environment, and can become actual usage later
free	-	byte	Free space of the datastore
usedPct	-	percent	used divided by capacity

The capacity fields relate as follows.

- $\text{capacity} = \text{used} + \text{free}$

- `allocated` = `used` + `uncommitted`

When `allocated` exceeds `capacity`, thin provisioning has promised more space than physically exists, an overcommit. Since `uncommitted` turns into actual usage as virtual machine disks grow, check the relationship between `allocated` and `capacity` rather than looking at `usedPct` alone.

When `accessible` is `false`, check the connection between the host and the datastore, the storage path, the SAN or NAS network, the permissions, and whether the datastore has failed.

When the latency stays high, check for storage response delays, path problems, backend storage load, and storage network delays. When several hosts share the same datastore, comparing whether the other hosts see the same latency separates a host problem from a backend problem.

infra_vm_vhost_network

A category that aggregates the network usage of an ESXi host. One host becomes one row, and it provides the sum across all physical network interfaces of the host.

- Collection target: VMware ESXi

Field	Type	Unit	Description	Remarks
<code>vhostName</code>	<code>string</code>	-	Name of the host being collected. The basis for identifying which ESXi host each network metric came from	
<code>uuid</code>	<code>string</code>	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes	
<code>usage</code>	-	<code>bytes/s</code>	Total network usage of the host. The sum of received and transmitted traffic	
<code>bytesRx</code>	-	<code>bytes/s</code>	Data throughput the host received over the network. Includes virtual machines, VMkernel, and the management network	
<code>bytesTx</code>	-	<code>bytes/s</code>	Data throughput the host transmitted over the network. Includes virtual machines, VMkernel, and the management	

Field	Type	Unit	Description	Remarks
			network	
packetsTotal	-	counts/s	Sum of received packets (packetsRx) and transmitted packets (packetsTx)	
packetsRx	-	counts/s	Number of network packets the host received	
packetsTx	-	counts/s	Number of network packets the host transmitted	
errorsRx	-	counts/s	Number of errors that occurred during reception	
errorsTx	-	counts/s	Number of errors that occurred during transmission	
droppedRx	-	counts/s	Number of packets dropped during reception	
droppedTx	-	counts/s	Number of packets dropped during transmission	
errorsRxTx	-	counts/s	Sum of receive errors (errorsRx) and transmit errors (errorsTx)	New in 2.9.23
droppedRxTx	-	counts/s	Sum of receive drops (droppedRx) and transmit drops (droppedTx)	New in 2.9.23

When errorsRx keeps increasing, suspect problems with the physical NIC, cable, switch port, driver, or MTU settings. When errorsTx increases, check the physical NIC, switch port, link status, driver, and network congestion.

droppedRx can be caused by an insufficient receive queue, network flooding, a NIC or driver problem, or vSwitch processing delays. Use droppedTx to check for an insufficient transmit queue, network congestion, or processing delays on the physical NIC or vSwitch.

Use errorsRxTx and droppedRxTx when you want to check only the host-level totals without separating receive and transmit.

infra_vm_vhost_nic_perf

Usage per physical network interface of an ESXi host. One interface becomes one row. Use it to check which interface produces errors or drops.

- Collection target: VMware ESXi

Field	Type	Unit	Description	Remarks
vhostName	string	-	Name of the host being collected. The basis for identifying which ESXi host each NIC metric came from	
uuid	string	-	Hardware UUID of the host. Supplementary basis for identifying the same host even when its name or IP changes	
nicName	string	-	Name of the network instance the metrics were collected from. Usually a physical NIC name such as <code>vmnic0</code> or <code>vmnic1</code> , though another name can appear depending on the environment	
usage	-	bytes/s	Total network throughput of the NIC. The sum of reception and transmission	
bytesRx	-	bytes/s	Data throughput received through the NIC	
bytesTx	-	bytes/s	Data throughput transmitted through the NIC	
packetsTotal	-	counts/s	Sum of received packets (<code>packetsRx</code>) and transmitted packets (<code>packetsTx</code>)	
packetsRx	-	counts/s	Number of network packets the NIC received	
packetsTx	-	counts/s	Number of network packets the NIC transmitted	
errorsRx	-	counts/s	Number of errors that occurred during reception	

Field	Type	Unit	Description	Remarks
errorsTx	-	counts/s	Number of errors that occurred during transmission	
droppedRx	-	counts/s	Number of packets dropped during reception	
droppedTx	-	counts/s	Number of packets dropped during transmission	
errorsRxTx	-	counts/s	Sum of receive errors (errorsRx) and transmit errors (errorsTx)	New in 2.9.23
droppedRxTx	-	counts/s	Sum of receive drops (droppedRx) and transmit drops (droppedTx)	New in 2.9.23

When errors or drops increase on only one NIC, first check that NIC's cable, switch port, driver, and MTU settings. The totals summed across the whole host are available in `infra_vm_vhost_network` .

Use `errorsRxTx` and `droppedRxTx` when you want to check only the per-interface totals without separating receive and transmit.

infra_vm_vhost_vmquest_perf

Performance metrics of the virtual machines running on an ESXi host. One powered-on virtual machine becomes one row, and it provides the CPU, memory, disk, and network usage of the virtual machine as seen from the host.

- Collection target: VMware ESXi

Field	Type	Unit	Description	Remarks
vHostName	string	-	Name of the host on which the virtual machine is running	
uuid	string	-	Hardware UUID of the ESXi host on which the virtual machine is running. Supplementary basis for identifying the same host even when its name or IP changes	

Field	Type	Unit	Description	Remarks
vmName	string	-	Name of the virtual machine the performance data was collected from. The name shown in the vCenter inventory	
vmInstanceUuid	string	-	UUID that identifies the virtual machine instance within vCenter. The main identifier for tracking the same virtual machine regardless of a name change or vMotion	
cpuUsage	-	percent	CPU usage of the virtual machine. The vSphere <code>cpu.usage.average</code> value converted to a percentage	
cpuUsagemhz	-	MHz	CPU resources the virtual machine is actually using. Check it with the usage rate to see the actual consumption	
cpuCostop	-	ms	Time the vCPUs of an SMP virtual machine waited because they could not be scheduled simultaneously	
cpuCostopPct	-	percent	<code>cpuCostop</code> divided by the sampling interval and converted to a ratio. Matches <code>%CSTP</code> as shown per VM group in esxtop	New in 2.9.23
cpuCostopPctPerVcpu	-	percent	<code>cpuCostopPct</code> divided by the number of vCPUs, giving the average ratio per vCPU. Always between 0 and 100	New in 2.9.23
cpuReadiness	-	ms	Time the virtual machine waited without being scheduled after it was ready to use the CPU, because of insufficient physical CPU resources	
cpuReadinessPct	-	percent	<code>cpuReadiness</code> divided by the sampling interval and converted to a ratio. Matches <code>%RDY</code> as shown per VM group in esxtop	New in 2.9.23
cpuReadinessPerVcpu	-	percent	Average waiting ratio per vCPU. Provided directly by	New in

Field	Type	Unit	Description	Remarks
			the vSphere <code>cpu.readiness.average</code> counter and always between 0 and 100	2.9.23
<code>cpuMaxlimited</code>	-	percent	Ratio of the CPU time that was available but restricted by a CPU limit or a similar setting	
<code>memConsumed</code>	-	MB	Memory the virtual machine actually consumed from the host memory	
<code>memVmmemctl</code>	-	KB	Virtual machine memory reclaimed by the balloon driver (vmmemctl)	
<code>memActive</code>	-	KB	Memory judged to have been actively used recently. Used to estimate the effective working set	
<code>memGranted</code>	-	KB	Memory the host granted to the virtual machine	
<code>memSwapped</code>	-	KB	Amount of the virtual machine memory the host swapped out	
<code>memSwapoutRate</code>	-	KB/s	Rate at which the virtual machine memory is swapped out to the host swap	
<code>memCompressed</code>	-	KB	Amount of the virtual machine memory the host is keeping compressed	
<code>memCompressionRate</code>	-	KB/s	Rate at which the virtual machine memory is compressed	
<code>diskUsage</code>	-	KB/s	Total disk usage from the guest perspective. The sum of reads and writes	
<code>networkUsage</code>	-	KB/s	Total network usage from the guest perspective. The sum of reception and transmission	

`cpuCostopPct` and `cpuReadinessPct` grow with the number of vCPUs on a multi-vCPU virtual machine, so they

can exceed 100%. When judging thresholds, take the vCPU count into account, or use `cpuCostopPctPerVcpu` and `cpuReadinessPerVcpu`, which are converted to a per-vCPU basis. The per-vCPU values are always between 0 and 100, so virtual machines with different vCPU counts can be compared on the same basis.

The judgment criteria in the official VMware guide are as follows.

Metric	Normal	Caution	Critical
<code>cpuCostopPctPerVcpu</code>	3% or less	3–10%	10% or more
<code>cpuReadinessPerVcpu</code>	5% or less	5–10%	10% or more

When `cpuCostopPctPerVcpu` is at the critical level, consider reducing the vCPUs allocated to the virtual machine or moving it to another host. When `cpuReadinessPerVcpu` is at the critical level, consider insufficient host CPU resources or relocating the virtual machine.

When `memVmmemctl`, `memSwapped`, and `memCompressed` increase, suspect memory pressure on the host. Compression usually occurs before swapping, so when `memSwapped` also increases after `memCompressed` has risen, the memory shortage is worsening.

Counters that vCenter does not provide are collected as empty values.

vm_vcenter_inventory

Configuration information of the vCenter server itself. It provides the number of managed objects, the number of active alarms, and the certificate and license expiration information. Use it to check items that need advance preparation, such as vCenter certificate expiration. This category was added in 2.9.23, so every field in the table is new.

- Collection target: VMware vCenter

Field	Type	Unit	Description
<code>vCenterHostname</code>	string	-	Host name of the vCenter server. The representative name that identifies the vCenter
<code>fqdn</code>	string	-	FQDN of the vCenter server. Collected only when it is registered

Field	Type	Unit	Description
			in DNS
ip	string	-	Representative IP address of the vCenter server
allIPs	string	-	List of all IP addresses that can reach the vCenter server
connectionAddress	string	-	Address (IP or host name) the agent used to connect to the vCenter API
uuid	string	-	Unique identifier of the vCenter server. Identifies the same instance regardless of a server migration or a name change
overallHealth	string	-	Overall status of the entire vCenter inventory. Displays the highest severity among <code>green</code> , <code>yellow</code> , and <code>red</code>
redAlarmCount	string	count	Number of currently active critical (red) alarms
yellowAlarmCount	string	count	Number of currently active warning (yellow) alarms
alarmCount	string	count	Number of all currently active alarms
hostCount	string	count	Number of ESXi hosts managed by vCenter
vmCount	string	count	Number of virtual machines managed by vCenter
datacenterCount	string	count	Number of registered data centers
clusterCount	string	count	Number of registered clusters
datastoreCount	string	count	Number of registered datastores
networkCount	string	count	Number of registered networks
resourcePoolCount	string	count	Number of registered resource pools
folderCount	string	count	Number of folders in the vCenter inventory

Field	Type	Unit	Description
certificateDaysRemaining	string	count	Days remaining until the vCenter HTTPS certificate expires
certificateExpiration	string	-	Expiration date and time of the vCenter HTTPS certificate
certificateValidFrom	string	-	Date and time from which the vCenter HTTPS certificate is valid
certificateSubject	string	-	Subject (DN) information of the vCenter HTTPS server certificate
certificateIssuer	string	-	Issuer (CA) information that issued the vCenter HTTPS server certificate
certificateThumbprint	string	-	SHA-256 thumbprint of the vCenter HTTPS certificate
licenseExpiration	string	-	Expiration date and time of the vCenter license. Can be empty for a perpetual license
licenseEdition	string	-	Edition of the applied vCenter license
licenseName	string	-	Name of the applied vCenter license
licenseTotal	string	count	Total quantity provided by the license
licenseUsed	string	count	License quantity currently in use
version	string	-	vCenter server version
build	string	-	vCenter server build number
fullName	string	-	Full product name and version string provided by VMware
apiVersion	string	-	vSphere Web Services API version
deploymentType	string	-	vCenter deployment type (for example, Embedded PSC)
timeZone	string	-	Time zone configured on the vCenter appliance

Field	Type	Unit	Description
agentName	string	-	Host name of the server where the agent collecting the VMware information is installed

Some items are provided only by the appliance API of vSphere 7 or later. In environments where that API is not available, the values are empty or displayed as substitute values.

Adjust whether to collect with the following option.

whatap.conf

```
vm.inventory.vcenter.enabled=true # Whether to collect (Default: true)
```

vm_vhost_inventory

Configuration information of ESXi hosts. One host becomes one row, and it is sent only when a value changes. It provides everything from the vCenter, data center, and cluster the host belongs to, through the CPU, memory, GPU, and network configuration, to the operational status.

- Collection target: VMware ESXi

Field	Type	Unit	Description
vCenterServerIp	string	-	IP address of the vCenter server that collected this ESXi host information. In environments that collect from multiple vCenters, it identifies which vCenter the host belongs to
dataCenterName	string	-	Name of the data center that contains the host in the vCenter inventory structure
clusterName	string	-	Name of the cluster the host belongs to. Empty or <code>None</code> for a standalone host or a host that is not in a cluster
folder	string	-	Host folder path where the host is located in the vSphere inventory. Displayed based on the host folder under the data

Field	Type	Unit	Description
			center or a management folder created by the user. Interpret it separately from the actual resource configuration location, such as the cluster, resource pool, or datastore
vHostName	string	-	ESXi host name or connection address identified in vCenter. Usually in the form of a host IP or FQDN
esxiHostNameFqdn	string	-	Fully qualified domain name of the ESXi host. Displayed according to the DNS or host name settings
esxiHostNameShort	string	-	Short host name of the ESXi host. The name with the domain removed from the FQDN
esxiHostNameDomain	string	-	Domain area configured in the ESXi host name. Empty when no domain is configured
esxiManagementIp	string	-	Management IP address of the ESXi host. The representative management IP used for vCenter connections, the management network, and API communication
esxiVmIps	string	-	IP information of the VMkernel interfaces configured on the host. Used for management, vMotion, vSAN, iSCSI/NFS, and other network purposes
manufacturer	string	-	Host hardware manufacturer (for example, HP, Dell, Lenovo)
model	string	-	Host hardware model name. A server model or workstation model is displayed
hostType	string	-	Configuration type of the host. <code>cluster</code> when it belongs to a cluster, <code>standalone</code> when it is a standalone host
serialNumber	string	-	Serial number of the host equipment. Used as identifying information for asset management, maintenance contracts, and vendor support requests

Field	Type	Unit	Description
<code>uuid</code>	<code>string</code>	-	Unique identifier of the host system. Serves as supplementary identification even when names are duplicated or the IP changes
<code>cpuModel</code>	<code>string</code>	-	CPU model name installed on the host (for example, Intel Xeon, Intel Core, AMD EPYC)
<code>cpuSockets</code>	<code>string</code>	-	Number of physical CPU sockets. Used for licensing, capacity planning, and checking the physical CPU configuration
<code>cpuCores</code>	<code>string</code>	-	Total number of physical CPU cores on the host. The basis for allocatable virtual machine resources and CPU capacity planning
<code>cpuThreads</code>	<code>string</code>	-	Number of logical CPU threads including hyper-threading. The logical CPU unit that can actually be scheduled for virtual machines
<code>hyperThreadingEnabled</code>	<code>string</code>	-	Whether hyper-threading is enabled on the host. Used to interpret the difference between the <code>cpuCores</code> and <code>cpuThreads</code> values
<code>cpuTotalMhz</code>	<code>string</code>	<code>MHz</code>	Total processing capacity of all CPUs on the host. Calculated from the number of CPU cores and the clock information
<code>totalMemoryGb</code>	<code>string</code>	<code>GiB</code>	Total memory capacity installed on the host. The basis for the capacity available for virtual machine placement and memory resource planning
<code>gpuCount</code>	<code>string</code>	<code>count</code>	Number of physical GPU devices identified on the host. Calculated from the VGA, 3D Controller, and Display Controller class devices among the PCI device information
<code>gpuDevices</code>	<code>string</code>	-	List of GPU devices identified on the host. Displays the manufacturer, model name, and PCI device identifier (for example, NVIDIA A100, NVIDIA T4, NVIDIA L40)
<code>gpuProfiles</code>	<code>string</code>	-	List of vGPU profile or shared passthrough GPU type candidates available on the host. The value can be empty depending on the

Field	Type	Unit	Description
			NVIDIA vGPU Manager, the ESXi and vCenter versions, the license, and how the API exposes them
vmGuestCount	string	-	Number of virtual machine objects connected to the host. Used to identify the virtual machine placement and density per host
datastoreCount	string	-	Number of datastores accessible from the host. Used to identify the number of storage connections and the configuration scale available to the host
dataStores	string	-	List of datastores accessible from the host. The storage resources where virtual machine disks can be stored
totalCapacityGb	string	GiB	Total capacity of the datastores accessible from the host. Shared datastores can be included, so do not treat it as the capacity of the host alone
networkCount	string	-	Number of networks or port groups connected to the host
networks	string	-	List of networks or port groups connected to the host. Used to check the network connection range available to virtual machines
nicCount	string	-	Number of physical NICs configured on the host. The basis for checking the scale of physical connections for the management, virtual machine, vMotion, and storage networks
nics	string	-	List of physical NICs configured on the host
vSwitchCount	string	-	Number of virtual switches configured on the host. Used to identify whether standard or distributed switches are used and the network configuration scale
vSwitches	string	-	List of vSwitches configured on the host
networkInfo	string	-	vSwitch or network configuration information on the host

Field	Type	Unit	Description
vMotionEnabled	string	-	Whether vMotion is enabled on the host. Used to determine whether virtual machines can be migrated without downtime
haEnabled	string	-	Whether HA is enabled in the cluster the host belongs to. <code>None</code> or <code>false</code> for a standalone host or when cluster HA is disabled
productVersion	string	-	ESXi product version. Used to determine the patch level, feature support, and compatibility
buildNumber	string	-	ESXi build number. The basis for checking the detailed patch version, responding to security vulnerabilities, and reviewing compatibility
license	string	-	Type of VMware license applied to the host. Used to determine the available feature range and operational constraints
bootTime	string	-	Time the host was last booted. Used to check long-term operation and the recovery point after an incident
upTime	string	-	Time the host has been running since the last boot. Used to identify servers that have not been restarted for a long time
timezone	string	-	Time zone information based on the host or vCenter. The reference time for interpreting event times and performance data times
connectionState	string	-	Host connection status as seen by vCenter. Displayed as <code>connected</code> , <code>disconnected</code> , <code>notResponding</code> , and so on, and directly related to whether collection is possible
powerState	string	-	Power state of the host. <code>poweredOn</code> is normally the operational state
overallStatus	string	-	Overall host status determined by vCenter. Displayed as <code>green</code> , <code>yellow</code> , <code>red</code> , <code>gray</code> , and so on, and used to quickly determine whether there is a warning or failure

Field	Type	Unit	Description
<code>maintenanceMode</code>	<code>string</code>	-	Whether the host is in maintenance mode. When <code>true</code> , the host may be excluded from operation for virtual machine migration, patching, or other work

You can join the virtual machine list with the host list in this table using `vHostKey` in `vm_vmguest_inventory`.

vm_vmguest_inventory

Configuration information of virtual machines. One virtual machine becomes one row, and it is sent only when a value changes.

- Collection target: VMware ESXi

Field	Type	Unit	Description	Remarks
<code>vHostKey</code>	-	-	Identification key of the ESXi host on which the virtual machine is running. The same value as the key in <code>vm_vhost_inventory</code>	New in 2.9.23
<code>vmName</code>	<code>string</code>	-	Display name of the virtual machine managed in vCenter or ESXi. The representative name operators use to identify it in the vSphere environment	
<code>guestHostname</code>	<code>string</code>	-	Host name recognized inside the guest operating system. Collected through VMware Tools, so the value can be empty depending on the Tools status or the guest settings	
<code>vCenterServerInfo</code>	<code>string</code>	-	Information about the vCenter server that collected this virtual machine information. In environments that collect from multiple vCenters, it identifies which vCenter the virtual machine belongs to	
<code>datacenter</code>	<code>string</code>	-	Name of the vCenter data center that contains the	

Field	Type	Unit	Description	Remarks
			virtual machine	
folder	string	-	VM folder path where the virtual machine is located in the vSphere inventory. A logical folder location from the VMs and Templates perspective, separate from the execution or resource allocation location such as the host, cluster, or resource pool	
cluster	string	-	Name of the ESXi cluster the virtual machine belongs to. The value can be empty in a standalone host environment or when it is not in a cluster	
vHostInfo	string	-	Information about the ESXi host where the virtual machine is running or registered. The display format can vary between the host name, IP, and FQDN depending on the collection source	
resourcePool	string	-	Name of the resource pool the virtual machine belongs to. Used to identify the resource policy, allocation structure, and service-level groups	
ipAddress	string	-	Representative IP address identified in the guest operating system. Depending on the VMware Tools status, network settings, and multi-NIC configuration, the value can be empty or show only the representative IP	
guestOS	string	-	Guest operating system type configured for the virtual machine. Based on the OS type selected in vSphere, so it can differ from the actual OS	
guestFullName	string	-	Detailed name of the guest operating system. Identified through VMware Tools or the virtual machine settings	
guestId	string	-	Guest operating system identifier used by vSphere	

Field	Type	Unit	Description	Remarks
			(for example, <code>rhel8_64Guest</code> , <code>windows9Server64Guest</code>)	
<code>vcpu</code>	<code>string</code>	<code>count</code>	Number of virtual CPUs allocated to the virtual machine	
<code>sockets</code>	<code>string</code>	<code>count</code>	Number of virtual CPU sockets configured for the virtual machine. Used to check the vCPU configuration and how the guest operating system recognizes the CPUs	
<code>coresPerSocket</code>	<code>string</code>	<code>count</code>	Number of cores per socket configured for the virtual machine. Used to check the CPU topology recognized by the guest operating system and the licensing policy	
<code>totalMemory</code>	<code>string</code>	<code>MiB</code>	Total memory capacity allocated to the virtual machine. The basis for judging the resource allocation scale and whether memory needs to be expanded	
<code>vGpuCount</code>	<code>string</code>	<code>count</code>	Number of vGPUs allocated to the virtual machine. Calculated from the vGPU devices connected as Virtual PCI Passthrough devices. <code>0</code> for a virtual machine with no vGPU allocated	
<code>vGpuProfiles</code>	<code>string</code>	-	vGPU profile information allocated to the virtual machine (for example, <code>grid_a100-1-5c</code> , <code>grid_t4-2q</code>). The value can be empty or displayed differently depending on the vCenter and ESXi versions, the NVIDIA vGPU Manager configuration, and how the API exposes it	
<code>provisionedSpace</code>	<code>string</code>	<code>GiB</code>	Total storage capacity provisioned for the virtual machine. Can be larger than the actual usage in a thin provisioning environment	

Field	Type	Unit	Description	Remarks
<code>usedSpace</code>	<code>string</code>	<code>GiB</code>	Storage capacity actually in use by the virtual machine. Used for datastore usage analysis and capacity management	
<code>datastore</code>	<code>string</code>	-	List of datastores in use by the virtual machine. Used to check the storage where the disk files and configuration files are located	
<code>network</code>	<code>string</code>	-	Network or port group information connected to the virtual machine. Used to check the network connection configuration and to distinguish the service and management networks	
<code>vmPathName</code>	<code>string</code>	-	Datastore path where the virtual machine configuration file is located. Usually in the <code>[datastore] path/vm.vmx</code> form	
<code>vmToolsVersion</code>	<code>string</code>	-	VMware Tools version installed on the virtual machine. Related to whether guest information collection, IP identification, and heartbeat status checks are possible	
<code>template</code>	<code>string</code>	-	Whether the object is a VM template. When <code>true</code> , it is a deployment template rather than a running virtual machine	
<code>firmware</code>	<code>string</code>	-	Firmware type of the virtual machine. Displayed as <code>bios</code> or <code>efi</code> and used to check the boot method, secure boot, and compatibility	
<code>version</code>	<code>string</code>	-	Virtual hardware version of the virtual machine. Used to determine ESXi and vCenter compatibility, supported features, and whether migration is possible	
<code>cpuHotAddEnabled</code>	<code>string</code>	-	Whether adding vCPUs while the virtual machine is	

Field	Type	Unit	Description	Remarks
			running is enabled. When <code>true</code> , CPUs can be added without powering off the virtual machine	
<code>cpuHotRemoveEnabled</code>	<code>string</code>	-	Whether removing vCPUs while the virtual machine is running is enabled. Whether it actually works depends on the supported operating system and the settings	
<code>memoryHotAddEnabled</code>	<code>string</code>	-	Whether adding memory while the virtual machine is running is enabled. When <code>true</code> , memory can be expanded without powering off the virtual machine	
<code>maxMksConnections</code>	<code>string</code>	<code>count</code>	Maximum number of MKS connections that can access the virtual machine console at the same time. Used to check console access limits and operational access control	
<code>creationDate</code>	<code>string</code>	-	Date and time the virtual machine was created. Used for lifecycle management and tracking newly created virtual machines	
<code>uuid</code>	<code>string</code>	-	BIOS UUID of the virtual machine. It can also be referenced inside the guest operating system, and it can be duplicated depending on how the virtual machine is cloned or recreated	
<code>instanceUuid</code>	<code>string</code>	-	UUID that vCenter uses to identify the virtual machine instance. The main basis for identifying a virtual machine within the same vCenter	
<code>annotation</code>	<code>string</code>	-	Description or note entered in the vSphere virtual machine settings	
<code>powerState</code>	<code>string</code>	-	Power state of the virtual machine. Displayed as <code>poweredOn</code> , <code>poweredOff</code> , <code>suspended</code> , and so on	

Field	Type	Unit	Description	Remarks
<code>overallStatus</code>	<code>string</code>	-	Overall status of the virtual machine object. Displayed as <code>green</code> , <code>yellow</code> , <code>red</code> , <code>gray</code> , and so on, and used to quickly check for alarms or configuration problems	
<code>vmGuestState</code>	<code>string</code>	-	State seen from the VMware Tools or guest operating system perspective. Displayed as <code>running</code> , <code>notRunning</code> , <code>shuttingDown</code> , <code>standby</code> , and so on	
<code>guestHeartbeatStatus</code>	<code>string</code>	-	Guest operating system heartbeat status identified through VMware Tools. Displayed as <code>green</code> , <code>yellow</code> , <code>red</code> , <code>gray</code> , and so on	

You can join the virtual machine list and the host list with `vHostKey`. Use it to find all virtual machines placed on a specific host at once.

ⓘ The existing vCenter monitoring (feature project) will end service on December 31, 2026 and will be replaced by the vCenter feature of the server monitoring product. For more information, see the [vCenter monitoring](#) document.

GPU Metrics

WhaTap server monitoring collects and provides GPU performance and status metrics from NVIDIA GPU environments across four perspectives: the GPU server, individual GPUs, GPU-using processes, and GPU fault events. You can check GPU utilization, compute resource activity, memory usage, power, temperature, clock speeds, PCIe/NVLink throughput, ECC errors, performance throttle states, and Xid events — enabling step-by-step analysis from server-wide GPU utilization to per-GPU bottleneck analysis, process-level root cause tracing, and driver/hardware fault event inspection. Some metrics are collected only for Physical GPUs, and in MIG environments the collection targets and interpretation criteria may differ, so it is recommended to check the GPU Type and collection targets together. In addition, Xid events are collected based on the fault/event history that the NVIDIA driver records in the system log. When analyzing the root cause, check the event timestamp, GPU identification, Xid Code, PID, Task Name, and detailed message together.

ⓘ Agent configuration is required to enable GPU monitoring. For more information, see [GPU Feature Settings](#).

GPU performance metrics are provided in 4 categories based on the analysis perspective: server-wide GPU aggregation, per-GPU detail, GPU-using process, and GPU Xid events.

- **infra_gpu**

A category that aggregates the status of all GPUs installed in the server, allowing you to **check the performance and status from a server-wide GPU perspective**. It provides GPU utilization, memory, power, temperature, clock speeds, PCIe/NVLink, ECC, and Violation states based on averages, sums, and maximums, making it easy to quickly assess the overall utilization level and potential bottlenecks of the GPU server.

- **infra_gpu_id**

A **category that provides performance and status at the individual GPU level**. You can check per-GPU utilization, compute resource activity, memory, temperature, power, clock speeds, PCIe/NVLink, ECC, and Violation states to analyze bottlenecks or throttle conditions occurring on a specific GPU. Collectible metrics may differ between Physical GPUs and MIG instances, so GPU Type and collection targets must be checked together.

- **infra_gpu_process**

A **category that provides resource usage information at the per-process level for GPU-using processes**.

Along with PID, process name, user, and GPU UUID, you can check SM utilization, memory controller utilization, encoder/decoder utilization, FB memory usage, and more to trace which processes are consuming GPU resources. Note that some per-process GPU utilization metrics may have different measurement criteria than the overall GPU metrics, so when assessing individual process contributions, interpret them together with the overall GPU metrics.

- **infra_gpu_xid**

A **category that provides GPU Xid event information recorded by the NVIDIA driver**. Based on the `NVRM: Xid` logs recorded in `/var/log/messages` or the system log, it collects the Xid Code, GPU Index, PCI Bus ID, occurrence time, PID, Task Name, and detailed message, allowing you to trace GPU anomaly events such as GPU driver, hardware, PCIe, memory, and application errors. Since not every Xid event immediately indicates a hardware failure, it is recommended to interpret them together with the meaning of the Xid Code, occurrence frequency, concurrent GPU performance metrics, process information, and system logs.

GPU Metrics Detail

The following sections describe the detailed metrics collected per category.

infra_gpu

A **category that collects the aggregated usage status of all GPUs installed in the server**. It provides GPU utilization, memory usage, power, temperature, clock speeds, PCIe/NVLink throughput, ECC, and violation states recalculated as averages, sums, and maximums from a server-wide GPU perspective. Use this to understand the overall utilization level, bottlenecks, and throttle states of the entire GPU server.

- To collect DCP (Datacenter Profiling) metrics, the `nvidia-dcgm` service (or `nv-hostengine` process) must be running.
- For Agent 2.9.16 and later, when `nvidia2.collector=true` is set, metric data is collected automatically via the connectable NVML and DCGM Library. (`nvidia-smi` or `dcgmi` queries are not used.)
- Default collection interval: 30s

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
name_50	GPU Name	string	Represents the model name of the GPU. The display string may vary slightly depending on the driver version or collection path.			2.9.3
model_family	GPU Model Family	string	A normalized value that maps GPU model names — which may appear differently depending on driver version, firmware, or collection path — to representative family names such as A100, H100, or L40S. Used as an identifier for consistent per-model aggregation, filtering, and comparison.			2.9.7
gpu_count	GPU Count	string	The number of Physical GPUs recognized by the server. Used as the baseline when			2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			interpreting average, total, and ratio metrics across all GPUs.			
gpu_utilization	GPU Utilization	%	The average utilization across individual GPUs in the server. In Physical GPU environments, this represents the standard GPU Util value; in MIG environments, the GR Engine Active value per instance can be used as a representative GPU utilization. Shows the overall utilization level of all GPUs in the server.	AVG of installed Physical GPUs	Required to run nvidia-dcgm service or nv-hostengine process in MIG Enabled environments	2.9.6
enc_util_206	Encoder Util	%	The average encoder engine utilization across individual GPUs in the server. Shows how much video encoding work	AVG of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			is being handled by the GPUs. Higher values indicate greater use of NVENC-family engines.			
dec_util_207	Decoder Util	%	The average decoder engine utilization across individual GPUs in the server. Shows how much video decoding work is being handled by the GPUs. Higher values indicate greater use of NVDEC-family engines.	AVG of installed Physical GPUs		2.9.3
gr_engine_active_1001	GR Engine Active Util	%	The average Graphics/Compute Engine active ratio across individual GPUs in the server. In MIG instance environments, this metric can be used as a representative GPU utilization for	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			each instance.			
sm_active_1002	SM Active Util	%	The average SM (Streaming Multiprocessor) active ratio across individual GPUs in the server. Shows the overall utilization level of SM resources used in actual computation.	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
sm_occupancy_1003	SM Occupancy	%	The average SM Occupancy across individual GPUs in the server. Indicates how fully the SMs are loaded with active warps — unlike a simple active ratio, this shows the density of SM resource utilization.	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
tensor_active_1004	Tensor Core Util	%	The average Tensor Core active ratio across individual GPUs in the server. Useful for	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			understanding the workload mix that uses Tensor Cores, such as AI/deep learning computation.		must be running	
fp64_active_1006	FP64 Compute Util	%	The average FP64 compute active ratio across individual GPUs in the server. Indicates the double-precision compute load level.	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
fp32_active_1007	FP32 Compute Util	%	The average FP32 compute active ratio across individual GPUs in the server. Indicates the single-precision compute load level.	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
fp16_active_1008	FP16 Compute Util	%	The average FP16 compute active ratio across individual GPUs in the server. Indicates the half-precision compute load level, which can	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			be important for AI inference/training workloads.			
memory_copy_util_204	Memory Copy Util	%	The average memory copy engine utilization across individual GPUs in the server. Shows how actively data movement paths — such as GPU memory to internal compute areas, Host ↔ Device copies, and Device ↔ Device copies — are being used. Higher values may indicate that data movement load is greater than computation itself, or that memory copy operations are becoming a bottleneck. Note that this value represents the activity level of the memory copy/transfer engine, not memory capacity	AVG of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			utilization.			
dram_active_1005	DRAM Active Util	%	The average DRAM active ratio across individual GPUs in the server. Represents the proportion of time the GPU DRAM memory subsystem was active. Shows how actively the GPU memory bandwidth is being used; tends to be high in memory-intensive or memory-access-heavy workloads. Sustained high values may indicate that GPU memory bandwidth is a bottleneck rather than compute performance. Note that this value is an activity indicator showing how busy the GPU DRAM was — not a memory	AVG of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			capacity metric like FB Memory Used.			
fb_total_250	FB Total Memory	MiB	The sum of FB (Frame Buffer) total memory capacity across individual GPUs in the server. Represents the total GPU memory capacity of the entire server.	SUM of installed Physical GPUs		2.9.3
fb_free_251	FB Free Memory	MiB	The sum of FB free memory capacity across individual GPUs in the server. Represents the total available GPU memory of the entire server.	SUM of installed Physical GPUs		2.9.3
fb_used_252	FB Used Memory	MiB	The sum of FB used memory capacity across individual GPUs in the server. Represents the total GPU memory currently in use across the entire	SUM of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			server.			
fb_resv_253	FB Reserved Memory	MiB	The sum of FB reserved memory capacity across individual GPUs in the server. Represents the total GPU memory reserved by the system, drivers, firmware, etc.	SUM of installed Physical GPUs		2.9.3
fb_USDP_254	FB Memory Usage	%	The overall FB memory utilization for the entire server. Rather than a simple average of individual GPU utilization rates, this is recalculated as FB Used sum / FB Total sum * 100.	AVG of installed Physical GPUs		2.9.3
bar1_total_90	BAR1 Total Memory	MiB	The sum of BAR1 total capacity across individual GPUs in the server. BAR1 is the memory-mapped region used by the host to access GPU	SUM of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			memory; this represents the total BAR1 capacity of the entire server.			
bar1_used_92	BAR1 Used Memory	MiB	The sum of BAR1 usage across individual GPUs in the server. Represents the total BAR1 memory currently in use across the entire server.	SUM of installed Physical GPUs		2.9.3
bar1_free_93	BAR1 Free Memory	MiB	The sum of BAR1 free capacity across individual GPUs in the server. Represents the total available BAR1 capacity of the entire server.	SUM of installed Physical GPUs		2.9.3
bar1_usdp	BAR1 Memory Usage	%	The overall BAR1 utilization for the entire server. Rather than a simple average of individual GPU ratios, this is	AVG of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			recalculated as BAR1 Used sum / BAR1 Total sum * 100.			
ecc_sbe_aggregate	ECC SBE Total	count	The sum of cumulative ECC Single Bit Error counters across individual GPUs in the server. Represents the total accumulated single- bit ECC errors across the entire server.	SUM of installed Physical GPUs		2.9.3
ecc_dbe_aggregate	ECC DBE Total	count	The sum of cumulative ECC Double Bit Error counters across individual GPUs in the server. Represents the total accumulated double- bit ECC errors across the entire server.	SUM of installed Physical GPUs		2.9.3
ecc_sbe_aggregate	ECC SBE Delta	count	The sum of ECC Single Bit Error increments across individual GPUs in	SUM of installed Physical GPUs		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			the server during the collection interval. Represents the number of new single-bit ECC errors that occurred in the recent interval.			
ecc_dbe_aggregate_delta	ECC Double Bit Error Delta	count	The sum of ECC Double Bit Error increments across individual GPUs in the server during the collection interval. Represents the number of new double-bit ECC errors that occurred in the recent interval.	SUM of installed Physical GPUs		2.9.16
gpu_temp_150	GPU Temperature	°C	The maximum temperature among individual GPUs in the server. Represents the highest temperature across all GPUs in the server, used to quickly identify whether a specific	MAX of installed Physical GPUs		2.9.10

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			GPU is overheating.			
memTemp_140	GPU Memory Temperature	°C	The maximum memory temperature among individual GPUs in the server. Typically the HBM or GDDR memory temperature; used to check for memory thermal issues or cooling anomalies. In memory-intensive workloads, this may rise before the GPU core temperature, so it is recommended to monitor both together.	MAX of installed Physical GPUs		2.9.16
power_usage_155	Power Usage	watt	The sum of power consumption across individual GPUs in the server. Represents the total power currently consumed by all GPUs in the server, in watts (W).	SUM of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
energy_cumulative	Energy daily_kwh Cumulative	kWh	The cumulative GPU energy consumption of the entire server since 00:00 of the current day. This is a reference estimate calculated by integrating the powerUsage values collected via NVIDIA drivers/DCGM at each collection interval, and may differ from precise power measurements from a dedicated energy meter. Use for reference to track GPU power usage trends and daily usage patterns.	SUM of installed Physical GPUs		2.9.16
fan_speed_191	Fan Speed	%	The average fan speed across individual GPUs in the server. Indicates the cooling device activity level for GPU models equipped	AVG of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			with fans.			
sm_clock_100	SM Clock	MHz	The average SM clock speed across individual GPUs in the server. Represents the average core clock level at which the compute units are operating.	AVG of installed Physical GPUs		2.9.3
memory_clock_101	Memory Clock	MHz	The average memory clock speed across individual GPUs in the server. Represents the operating speed level of GPU memory.	AVG of installed Physical GPUs		2.9.3
video_clock_102	Video Clock	MHz	The average video engine clock speed across individual GPUs in the server. Represents the operating speed level of video processing engines.	AVG of installed Physical GPUs		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
pcie_tx_bytes_1009	PCIE TX	bytes/s	The sum of PCIe transmit throughput across individual GPUs in the server. Represents the total PCIe data transmitted from the GPUs to the host or external destinations, in bytes/s.	SUM of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
pcie_rx_bytes_1010	PCIE RX	bytes/s	The sum of PCIe receive throughput across individual GPUs in the server. Represents the total PCIe data received by the GPUs from the host or external sources, in bytes/s.	SUM of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
pcie_tx_rx_bytes	PCIE Total	bytes/s	The sum of PCIe transmit and receive throughput across individual GPUs in the server. Represents the total PCIe data movement (PCIE TX + PCIE RX)	SUM of installed Physical GPUs	nvidia-dcgm service or nv-hostengine process must be running	2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			across the entire server.			
nvlink_tx_bytes_101	NVLink TX	bytes/s	The sum of NVLink transmit throughput across individual GPUs in the server. Represents the total data transmitted between GPUs via NVLink, in bytes/s.	SUM of installed Physical GPUs	nvidia-dcgm service or nv- hostengine process must be running	2.9.3
nvlink_rx_bytes_101	NVLink RX	bytes/s	The sum of NVLink receive throughput across individual GPUs in the server. Represents the total data received between GPUs via NVLink, in bytes/s.	SUM of installed Physical GPUs	nvidia-dcgm service or nv- hostengine process must be running	2.9.3
nvlink_tx_rx_bytes	NVLink Total	bytes/s	The sum of NVLink transmit and receive throughput across individual GPUs in the server. Represents the total GPU-to-GPU communication (NVLink TX + NVLink	SUM of installed Physical GPUs	nvidia-dcgm service or nv- hostengine process must be running	2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			RX) across the entire server.			
power_violation_rate_240	Power Violation Rate	%	The representative value showing the highest power cap violation rate among individual GPUs in the server. Used to quickly identify if any GPU is experiencing performance throttling due to power constraints.	MAX of installed Physical GPUs		2.9.16
thermal_violation_rate_241	Thermal Violation Rate	%	The representative value showing the highest thermal limit violation rate among individual GPUs in the server. Used to quickly identify if any GPU is experiencing performance throttling due to overheating.	MAX of installed Physical GPUs		2.9.16
low_util_violation_rate_244	Low Util Violation Rate	%	The representative value showing the highest low-	MAX of installed Physical		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			utilization violation rate among individual GPUs in the server. Used to check if any GPU is in a state that meets the configured low-utilization conditions.	GPUs		
board_limit_violation_rate_243	Board Limit Violation Rate	%	The representative value showing the highest board-level hardware limit violation rate among individual GPUs in the server. Used to identify the GPU most affected by board-level constraints such as power delivery and board design limits.	MAX of installed Physical GPUs		2.9.16
sync_boost_violation_rate_242	Sync Boost Violation Rate	%	The representative value showing the highest sync boost violation rate among individual GPUs in the server. Used to identify the GPU	MAX of installed Physical GPUs		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			most affected by constraints imposed to synchronize the operating speeds of multiple GPUs.			
reliability_violation_rate_245	Reliability Violation Rate	%	The representative value showing the highest reliability protection violation rate among individual GPUs in the server. Used to identify the GPU most affected by constraints for stability and hardware protection.	MAX of installed Physical GPUs		2.9.16
power_violation_246	Power Violation	ns	The sum of time individual GPUs in the server spent in a power cap throttled state during the collection interval. In nanoseconds (ns); shows the total duration of performance throttling due to	SUM of installed Physical GPUs		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			power constraints across the entire server.			
thermal_violation_241	Thermal Violation	ns	The sum of time individual GPUs in the server spent in a thermal limit throttled state during the collection interval. In nanoseconds (ns); shows the total duration of performance throttling due to overheating across the entire server.	SUM of installed Physical GPUs		2.9.16
low_util_violation_244	Low Util Violation	ns	The sum of time individual GPUs in the server spent in a low-utilization throttled state during the collection interval. In nanoseconds (ns); shows the total duration of low- utilization condition	SUM of installed Physical GPUs		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			violations across the entire server.			
board_limit_violation_243	Board Limit Violation	ns	The sum of time individual GPUs in the server spent in a board-level hardware limit throttled state during the collection interval. In nanoseconds (ns); used to identify the GPU most affected by board-level constraints.	SUM of installed Physical GPUs		2.9.16
sync_boost_violation_242	Sync Boost Violation	ns	The sum of time individual GPUs in the server spent in a sync boost throttled state during the collection interval. In nanoseconds (ns); used to identify the GPU most affected by constraints imposed to synchronize the operating speeds of multiple GPUs.	SUM of installed Physical GPUs		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
reliability_violation_245	Reliability Violation	ns	The sum of time individual GPUs in the server spent in a reliability protection throttled state during the collection interval. In nanoseconds (ns); used to identify the GPU most affected by constraints for stability and hardware protection.	SUM of installed Physical GPUs		2.9.16
response_timeout	Response Timeout	string	If no data is collected through NVML for a certain period, this value is recorded as False.			2.9.19

infra_gpu_id

A category that collects per-GPU performance and status information for each individual GPU installed in the server. Provides per-GPU utilization, SM/DRAM/Tensor/FP compute utilization, memory capacity/usage, temperature, power, clock speeds, PCIe/NVLink, ECC, and violation states. Use this to check for overheating, power throttling, memory shortage, or interconnect bottlenecks on a specific GPU.

- Criteria for distinguishing collection targets as Physical/MIG: For metrics that represent the overall hardware status of the board/card, only Physical metrics are collected where Physical-only collection suffices.

- To collect DCP (Datacenter Profiling) metrics, the nvidia-dcgm service (or nv-hostengine process) must be running.
- For Agent 2.9.16 and later, when nvidia2.collector=true is set, metric data is collected automatically via the connectable NVML and DCGM Library. (nvidia-smi or dcgmi queries are not used.)
- Default collection interval: 30s

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
index	GPU Index	string	The index value used to identify the GPU within the server, representing the sequence number assigned by the OS or driver.	Physical, MIG		2.9.3
name_50	GPU Name	string	Represents the model name of this GPU. The display string may vary slightly depending on the driver version or collection path.	Physical, MIG		2.9.3
uuid_54	GPU UUID	string	The unique identifier (UUID) of this GPU. Used to reliably identify the GPU regardless of reboots or device order changes.	Physical, MIG		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
gpuType	GPU Type	string	Displays the type of this GPU as Physical or MIG.	Physical, MIG		2.9.3
modelFamily	GPU Model Family	string	Represents the upper model family for this GPU. Since the same GPU may display slightly different model names depending on driver version or collection path, this is provided as a simplified model name normalized to a common standard such as A100, H100, or L4. Used for consistent classification, aggregation, and comparison of GPUs of the same family.	Physical, MIG		2.9.7
gpuUtilization	GPU Utilization	%	The utilization of this GPU. In Physical GPU environments, this represents the standard GPU Util	Physical, MIG	Required to run nvidia-dcgm service or nv-	2.9.6

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			value; in MIG environments, the GR Engine Active value can be interpreted as a representative GPU utilization.		hostengine process in MIG Enabled environments	
encUtil_206	Encoder Util	%	The encoder engine utilization of this GPU. Higher values indicate greater use of NVENC-family engines.	Physical		2.9.3
decUtil_207	Decoder Util	%	The decoder engine utilization of this GPU. Higher values indicate greater use of NVDEC-family engines.	Physical		2.9.3
grEngineActive_1001	GR Engine Active Util	%	The Graphics/ Compute Engine active ratio of this GPU. In MIG instance environments, this metric can be used as a representative GPU utilization for	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			each instance.			
smActive_1002	SM Active Util	%	The SM (Streaming Multiprocessor) active ratio of this GPU. Shows the overall utilization level of SM resources used in actual computation.	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
smOccupancy_1003	SM Occupancy	%	The SM Occupancy of this GPU. Shows how fully the SMs are loaded with active warps — better represents the density of SM resource utilization than a simple active ratio.	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
tensorActive_1004	Tensor Core Util	%	The Tensor Core active ratio of this GPU. Useful for understanding the proportion of workloads that use Tensor Cores, such as AI/deep learning	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			computation.			
fp64Active_1006	FP64 Compute Util	%	The FP64 compute active ratio of this GPU. Shows the double-precision compute load level.	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
fp32Active_1007	FP32 Compute Util	%	The FP32 compute active ratio of this GPU. Shows the single-precision compute load level.	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
fp16Active_1008	FP16 Compute Util	%	The FP16 compute active ratio of this GPU. Shows the half-precision compute load level, which can be important for AI inference/training workloads.	Physical, MIG	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
memoryCopyUtil_204	Memory Copy Util	%	The memory copy engine utilization of this GPU. Shows how	Physical		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			actively data movement paths — such as GPU memory to internal compute areas, Host ↔ Device copies, and Device ↔ Device copies — are being used. Higher values may indicate that data movement load is greater than computation itself, or that memory copy operations are becoming a bottleneck. Note that this value represents the activity level of the memory copy/ transfer engine, not memory capacity utilization. MIG-enabled Physical GPUs cannot collect data.			
dramActive_1005	DRAM Active Util	%	The proportion of time the DRAM memory subsystem of this GPU was	Physical, MIG	nvidia-dcgm service or nv-hostengine	2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			active. Shows how actively the GPU memory bandwidth is being used; tends to be high in memory-intensive or memory-access-heavy workloads. Sustained high values may indicate that GPU memory bandwidth is a bottleneck rather than compute performance. Note that this value is an activity indicator showing how busy the GPU DRAM was — not a memory capacity metric like FB Memory Used.		process must be running	
fbTotal_250	FB Total Memory	MiB	The total FB (Frame Buffer) memory capacity of this GPU.	Physical, MIG		2.9.3
fbFree_251	FB Free Memory	MiB	The available FB memory capacity of this GPU. Represents	Physical, MIG		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			the amount of GPU memory currently available.			
fbUsed_252	FB Used Memory	MiB	The used FB memory capacity of this GPU. Represents the amount of GPU memory currently occupied.	Physical, MIG		2.9.3
fbResv_253	FB Reserved Memory	MiB	The reserved FB memory capacity of this GPU. Represents the GPU memory area reserved by the system, drivers, firmware, etc.	Physical, MIG		2.9.3
fbUsdp_254	FB Memory Usage	%	The FB memory utilization of this GPU. Calculated as $\text{FB Used} / \text{FB Total} * 100$.	Physical, MIG		2.9.3
bar1Total_90	BAR1 Total Memory	MiB	The total BAR1 capacity of this GPU. BAR1 is the memory-mapped region used by the host to access	Physical, MIG		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			GPU memory.			
bar1Used_92	BAR1 Used Memory	MiB	The BAR1 memory currently in use on this GPU.	Physical, MIG		2.9.3
bar1Free_93	BAR1 Free Memory	MiB	The available BAR1 capacity on this GPU.	Physical, MIG		2.9.3
bar1Usdp	BAR1 Memory Usage	%	The BAR1 utilization of this GPU. Calculated as BAR1 Used / BAR1 Total * 100.	Physical, MIG		2.9.3
eccSbeAggregateTotal_94	ECC SBE Total	count	The total number of accumulated ECC Single Bit Errors on this GPU. As a cumulative counter, continuous trend observation is important.	Physical		2.9.3
eccDbeAggregateTotal_95	ECC DBE Total	count	The total number of accumulated ECC Double Bit Errors on this GPU. Double-bit errors indicate a more severe	Physical		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			condition than single-bit errors.			
eccSbeAggregateDelta_SBE	ECC_SBE Delta	count	The number of newly incremented ECC Single Bit Errors during the most recent collection interval on this GPU. Used to check new single-bit error occurrences in the interval.	Physical		2.9.16
eccDbeAggregateDelta_DBE	ECC_DBE Delta	count	The number of newly incremented ECC Double Bit Errors during the most recent collection interval on this GPU. Used to check new double-bit error occurrences in the interval.	Physical		2.9.16
gpuTemp_150	GPU Temperature	°C	The current temperature of this GPU. The primary indicator for assessing	Physical		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			overheating at the individual GPU level.			
memTemp_140	GPU Memory Temperature	°C	The current memory temperature of this GPU. Typically the HBM or GDDR memory temperature; used to check for memory thermal issues or cooling anomalies. In memory-intensive workloads, this may rise before the GPU core temperature, so it is recommended to monitor both together.	Physical		2.9.16
powerUsage_155	Power Usage	watt	The current power consumption of this GPU, in watts (W).	Physical		2.9.3
powerRate	Power Rate	%	The power utilization rate of this GPU. Interpreted as Power Usage / Power Limit * 100; shows how close the current	Physical		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			power usage is to the power limit.			
energyCumulativeDaily	Energy Cumulative	kWh	The daily energy usage of this GPU, accumulated since 00:00 of the current day. This is a reference estimate calculated by integrating the powerUsage values collected via NVIDIA drivers/DCGM at each collection interval, and may differ from precise power measurements from a dedicated energy meter. Use for reference to track GPU power usage trends and daily usage patterns.	Physical		2.9.16
pstate_190	Performance State	P	The current performance state (P-state) of this GPU. Lower numbers	Physical		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			indicate higher performance states, with P0 being the highest. The range is typically P0 to P15; when stored, only the numeric portion without the P prefix is used (e.g., P0 is stored as 0, P2 as 2).			
fanSpeed_191	Fan Speed	%	The current fan speed of this GPU as a percentage. Shows the cooling device activity level on models equipped with fans.	Physical		2.9.3
smClock_100	SM Clock	MHz	The SM clock speed of this GPU in MHz. Shows the core clock level at which the compute units are operating.	Physical		2.9.3
memoryClock_101	Memory Clock	MHz	The memory clock speed of this GPU in MHz. Shows the operating speed of	Physical		2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			the GPU memory.			
videoClock_102	Video Clock	MHz	The video engine clock speed of this GPU in MHz. Shows the operating speed level of video processing engines.	Physical		2.9.3
pcieTxBytes_1009	PCIE TX	bytes/s	The data throughput transmitted from this GPU to external destinations via PCIe, in bytes/s.	Physical	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
pcieRxBytes_1010	PCIE RX	bytes/s	The data throughput received by this GPU from external sources via PCIe, in bytes/s.	Physical	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
pcieTxRxBytes	PCIE Total	bytes/s	The total PCIe transmit and receive throughput of this GPU in bytes/s. Equals PCIE TX +	Physical	nvidia-dcgm service or nv-hostengine process	2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			PCIE RX.		must be running	
nvlinkTxBytes_1011	NVLink TX	bytes/s	The data throughput transmitted from this GPU to other GPUs via NVLink, in bytes/s.	Physical	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
nvlinkRxBytes_1012	NVLink RX	bytes/s	The data throughput received by this GPU from other GPUs via NVLink, in bytes/s.	Physical	nvidia-dcgm service or nv-hostengine process must be running	2.9.3
nvlinkTxRxBytes	NVLink Total	bytes/s	The total NVLink transmit and receive throughput of this GPU in bytes/s. Equals NVLink TX + NVLink RX.	Physical	nvidia-dcgm service or nv-hostengine process must be running	2.9.16
powerViolationRate_240	Power Violation Rate	%	The proportion of time this GPU was in a power cap	Physical		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			throttled state. Shows how much performance throttling occurred due to the power limit.			
thermalViolationRate_241	Thermal Violation Rate	%	The proportion of time this GPU was in a thermal limit throttled state. Shows how much performance throttling occurred due to overheating.	Physical		2.9.16
lowUtilViolationRate_244	Low Util Violation Rate	%	The proportion of time this GPU was in a low-utilization throttled state. Shows how much of the configured low- utilization conditions were triggered.	Physical		2.9.16
boardLimitViolationRate_243	Board Limit Violation Rate	%	The proportion of time this GPU was in a board-level hardware limit throttled state.	Physical		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			Shows how much performance throttling occurred due to board-level constraints such as power delivery and board design limits.			
syncBoostViolation	Sync Boost Rate_242 Violation Rate	%	The proportion of time this GPU was in a sync boost throttled state. Shows how much performance throttling occurred due to constraints imposed to synchronize the operating speeds of multiple GPUs.	Physical		2.9.16
reliabilityViolation	Reliability Rate_245 Violation Rate	%	The proportion of time this GPU was in a reliability protection throttled state. Shows how much performance throttling occurred due to constraints for stability and	Physical		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			hardware protection.			
powerViolation_240	Power Violation	ns	The time this GPU spent in a power cap throttled state during the collection interval. In nanoseconds (ns); shows the duration of performance throttling due to the power limit.	Physical		2.9.16
thermalViolation_241	Thermal Violation	ns	The time this GPU spent in a thermal limit throttled state during the collection interval. In nanoseconds (ns); shows the duration of performance throttling due to overheating.	Physical		2.9.16
lowUtilViolation_244	Low Util Violation	ns	The time this GPU spent in a low-utilization throttled state during the collection interval. In nanoseconds (ns);	Physical		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			shows how long the low-utilization condition violation persisted.			
boardLimitViolation ₂₄₃	Board Limit Violation	ns	The time this GPU spent in a board-level hardware limit throttled state during the collection interval. In nanoseconds (ns); shows the duration of performance throttling due to board-level constraints.	Physical		2.9.16
syncBoostViolation ₂₄₂	Sync Boost Violation	ns	The time this GPU spent in a sync boost throttled state during the collection interval. In nanoseconds (ns); shows the duration of performance throttling due to constraints imposed to synchronize the operating speeds of	Physical		2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Target	Collection Condition	Min. Agent Version
			multiple GPUs.			
reliabilityViolation ₂₄₅	Reliability Violation	ns	The time this GPU spent in a reliability protection throttled state during the collection interval. In nanoseconds (ns); shows the duration of performance throttling due to constraints for stability and hardware protection.	Physical		2.9.16

infra_gpu_process

A category that collects GPU and system resource usage information for each individual process using the GPU. Provides GPU index/UUID, PID, process name/type, SM utilization, memory controller utilization, encoder/decoder utilization, FB memory usage, and protected memory (CCPM). It also provides CPU utilization, RSS memory, and process I/O information, enabling process-level analysis of not only GPU resource usage causes but also potential CPU, memory, and I/O bottlenecks.

- From Agent 2.9.16 onward, when `nvidia2.collector=true` is set, the mapping between processes and MIG instances can be confirmed in MIG environments.
- Default collection interval: 30s

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
index	GPU Index	string	The index of the GPU on which this process is running. Used as a sequence number to identify GPUs within the server.	Agent 2.9.16+: Can identify processes assigned to MIG		2.9.3
uuid	GPU UUID	string	The UUID of the GPU on which this process is running. As the GPU's unique identifier, it reliably identifies the GPU regardless of device order changes, unlike the index.			2.9.3
pid	Process ID	string	The PID (Process ID) of the process using the GPU.			2.9.3
process_name	Process Command	string	The name of the process using the GPU. For example, python, cuda_program, Xorg, etc. may be displayed.			2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
process_type	Process Type	string	The type of the running GPU process. C indicates a Compute process (CUDA, AI/compute tasks), and G indicates a Graphics process (GUI, rendering, display acceleration). Both types may coexist depending on the environment.			2.9.3
owner	Owner	string	The user account that launched this process.			2.9.16
args	Arguments	string	The command-line arguments of this process.			2.9.16
stime	StartTime	string	The start time of this process in epoch time.			2.9.16
gpu_utilization	GPU Utilization	%	Displays the overall GPU Utilization of the GPU device to which this process is			2.9.16

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			bound at the time of collection. This is not a directly measured per-process GPU utilization; rather, it shows the overall utilization of the connected GPU alongside the process information. As a result, multiple processes using the same GPU may show identical or similar values, and this value alone should not be used to determine how much each individual process contributed to GPU utilization.			
sm_util	SM Util	%	The proportion of the GPU's SM (Streaming Multiprocessor) compute resources being used by this process. Shows the utilization level of			2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			GPU core resources used in actual computation; higher values indicate the process is making heavy use of GPU compute resources. Note that this is a process-perspective SM utilization and may use different measurement criteria than the device-wide SM Active Util (%). Therefore, even if a specific process shows high SM Util, the GPU's SM Active Util may appear lower or different depending on idle time in the overall observation window, inter-kernel gaps, scheduling with other processes, MIG/Time-slicing effects, etc. (For processes using MIG			

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			Instances, per-process SM Util values may not be collected or displayed due to limitations in the driver and monitoring interface.)			
mem_controller_util	Memory Controller Util	%	The proportion of the GPU memory controller being used by this process. This is not how much VRAM capacity is occupied, but rather how actively read/write and data movement requests to GPU memory were made. Higher values may indicate the process is making heavy use of memory bandwidth or memory access paths. Note that this is a process-perspective memory			2.9.15

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			controller utilization and differs in meaning from the device-wide Memory Copy Util (%) or DRAM Active Util (%). Memory Copy Util is closer to the overall GPU memory copy/transfer engine utilization, while DRAM Active Util is closer to the proportion of time GPU DRAM was active with read/write requests during the observation window. (For processes using MIG Instances, per-process SM Util values may not be collected or displayed due to limitations in the driver and monitoring interface.)			

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
enc_util	Encoder Util	%	The proportion of the GPU encoder engine being used by this process. Primarily increases during video encoding, streaming, and video save/compression tasks.			2.9.3
dec_util	Decoder Util	%	The proportion of the GPU decoder engine being used by this process. Primarily increases during decoding of compressed video streams, playback, and video analysis preprocessing.			2.9.3
fb_used	FB Used	MiB	The Frame Buffer (GPU memory) size being used by this process, in MiB. Can be interpreted as the VRAM usage occupied by this process.			2.9.3

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
ccpm	CCPM	MiB	The Confidential Computing Protected Memory usage of this process, in MiB. Available on GPUs that support Confidential Computing, such as the NVIDIA H100. May display as 0 on GPUs that do not support this feature.			2.9.3
cpu	CPU	%	The total CPU utilization used by this process.			2.9.15
rss	RSS	Byte	The RSS (Resident Set Size) of this process; the amount of memory currently loaded in system memory.			2.9.15
total_char_io	Process IO Total	Bytes/s	The total character-based I/O (char I/O) throughput generated per second by this			2.9.15

Field Category field name	Title Metric name in Performance Summary / Performance Detail / Metrics Explorer	Unit	Description	Notes	Collection Condition	Min. Agent Version
			process. May include I/O to pipes, terminals, buffers, and cache layers in addition to disk, so it can be larger than actual physical disk I/O.			

infra_gpu_xid

A category that collects GPU fault/event history by parsing the `NVRM: Xid` logs that the NVIDIA driver records in `/var/log/messages`. It provides the Xid Code, GPU Index, PCI Bus ID, related PID/process name, and raw message, and is used to trace when a GPU error occurred, the affected GPU, the related process, and the error type. Since Xid can involve a wide range of causes — from simple application errors to driver, hardware, and PCIe issues — it is important to analyze the code meaning, occurrence frequency, GPU status metrics, and system logs together.

- Supported from Agent 2.9.19 or later.
- Default collection interval: 60s

Field Category field name	Title Performance summary/detail/ Metrics Explorer name	Unit	Description	Note	Collection condition	Minimum supported Agent Version
index	GPU Index	string	The index of the GPU			2.9.19

Field Category field name	Title Performance summary/detail/ Metrics Explorer name	Unit	Description	Note	Collection condition	Minimum supported Agent Version
			where the Xid event occurred. A number used to identify the GPU device within the server; matched with the GPU Index shown in nvidia-smi to identify the GPU where the fault occurred.			
pcieBusId	PCI Bus ID	string	The PCI Bus ID of the GPU where the Xid event occurred. Can be used as a secondary identifier for the physical GPU in environments where the GPU Index may change, and is referenced when checking the GPU device location within the same server.			2.9.19
agentTimestamp	Agent Timestamp	string	The actual event occurrence time recorded in /var/log/ messages.			2.9.19
xidCode	XID Code	string	The Xid error code			2.9.19


Field Category field name	Title Performance summary/detail/ Metrics Explorer name	Unit	Description	Note	Collection condition	Minimum supported Agent Version
			reported by the NVIDIA driver. A code that distinguishes various GPU anomaly states such as GPU driver, hardware, PCIe, memory, and application errors; used together with the meaning of each code to analyze the fault cause and determine whether an RMA (Return Merchandise Authorization) is needed.			
pid	PID	string	The process ID included in the Xid log. Used to trace the responsible process when a GPU error occurs while a specific process is running. The value may be empty if the log has no PID information or the event occurred at the driver/hardware			2.9.19

Field Category field name	Title Performance summary/detail/ Metrics Explorer name	Unit	Description	Note	Collection condition	Minimum supported Agent Version
			level.			
taskName	Task Name	string	The Task Name recorded by the NVIDIA driver when the Xid occurred. A short name corresponding to the Linux /proc/{pid}/comm or per-thread /proc/{pid}/task/{tid}/comm value, it does not mean the full command line (cmdline). In multi-threaded applications, even for the same PID, different names such as python3 or pt_nccl_watchdg may appear depending on the task where the error was recorded. Therefore, when identifying the actual application, it is recommended to interpret it together with the PID, GPU PCI address, GPU			2.9.19

Field Category field name	Title Performance summary/detail/ Metrics Explorer name	Unit	Description	Note	Collection condition	Minimum supported Agent Version
			UUID, and cmdline information.			
message	Message	string	The detailed message of the NVRM: Xid log parsed from /var/ log/messages. In addition to the Xid code, it may include the GPU UUID, PCI information, channel information, process information, and detailed error text; used for checking the raw text during fault cause analysis.			2.9.19

Metrics Chart

In **Metrics Chart**, you can view the metrics data collected from the monitoring targets in a chart as follows. Specifying the time and metrics is required.



Options at the top

In the 1 area, you can specify the time range of the chart and the agent to be monitored through the options at the top of **Metrics Chart**.

- **Time:** You can specify the start and end of the X-axis as the total range in time.
- **Interval:** You can specify the X-axis data interval as the time interval.
- **Time merge:** As one of the data merging methods, you can merge data within the time specified by intervals.

For example, **Avg** indicates the average value of data within 1 hour.

-  **Agent:** You can specify the agents to query. If not specified, all items are queried.

 For more information on how to use the time selector.

Metrics list

The  area is a list of metrics to query for options. First select a **Category**. After querying the metrics under the selected **Category**, select a desired metric. If you select the **Category** and metric, Based on the data within the time range specified in the top menu of the area , you can see the data via the chart widgets in the  area.

Data Merge

Data merge provides the following methods: **Object Merge** and **Time merge**.

- **Time merge** merges data with the same field values at regular intervals in the original data.
- **Object Merge** merges data with matching tags among data with different field values.

Chart widget

On the top left of , you can see the metric name. On the top right of the chart widget, you can see the following options:

CPUUtilization.Average

◀ ▶ Default ⌵ ⚙ 📷 CSV  Detail ✕

- Time movement: ◀ ▶ Using the left arrow button or right arrow button, it can be moved by -1 or +1 as much as the selected time range.

For example, if the time range is between 00:00 on February 13 and 00:00 on February 14, when you select the ◀ left arrow button, you can see the data within the time range between 00:00 on February 12 and 00:00 on February 13.

- Interval/Time Merge: In the menu at the top of , you can modify the specified interval and time merge.
- By selecting the monitoring target: , you can specify the monitoring targets. If not selected, it searches for all.

- Time comparison: If you select , you can compare the current trend of the same metric with that of the previous time zone.
- By selecting the snapshot: , you can snapshot charts except for widget options.
- By selecting the CSV: , you can download the plot data on the chart as a CSV file.
- **Detail:** By selecting the , the detailed view can be seen. When there are many monitoring targets, you can see the metrics trends for each monitoring target individually.

ⓘ If you do not see any options at the top of the metrics chart widget, select the  icon.

Deleting a widget

To delete all widgets placed on the screen, select  **Delete All Widgets** on the screen.

Metrics Search

Home > Select Project > **Analysis** > **Metrics Search**

You can query specific metrics based on tags in **Metrics Search**.

1 Metrics Search

Time selector: < 2023/08/17 07:57 ~ 2023/08/17 08:57 60Min > Max count: 100 Category: app_counter

2 Tags: All container containerKey host_ip okindName oname onodeName pid type 8/8 Filter: Filter by tag value

Field: All _clock_delta active_tx_0 active_tx_3 active_tx_8 active_tx_count apdex_satisfied apdex_tolerated apdex_total arrival_rate httpc_count httpc_error httpc_time meterlog resp90 resp95 resp_time sql_count sql_error sql_fetch_count sql_fetch_time sql_time tps tx_count tx_doc_time tx_error tx_http_time tx_sql_time tx_time tx_time_sql_sum 5/29

3

	Time	Oid	Tags								Fields				
			container	containerKey	host_ip	okindName	oname	onodeName	pid	type	active_tx_3	arrival_rate	sql_time	tps	tx_error
1	2023-08-17 07:57:40	633280970	agent.service	1395943798	10.21.1.26	demo-okind-1	demo-8103	node-1	3285013	java	10	20.57942	7.822286	20.94554	0
2	2023-08-17 07:57:40	-857948929	agent.service	1395943798	10.21.1.26	demo-okind-1	demo-8105	node-1	3284721	java	10	17.974834	10.626543	20.609276	0

1 Time and category selection

In the **1** area, you can specify the **Time selector**, **Max count**, and **Category** for the collected metrics. **Time selector** and **Category** must have been specified.

- **Time selector:** You can query by specifying the time when metrics were collected. The default value is **1 hour**. In addition to the lookup time provided as the default option, you can specify the date and time after selecting the **Time selector** tab.

! For more information on how to use the **Time selector** option.

- **Max count:** You can specify the maximum number of metrics to search in the **3** list. You can set the count up to **10**, **50**, **100**, **200**, **300**, **1000**, **2000**, and **3000**.
- **Category:** It is a classification unit for the related metrics. You can set the desired **Category** after selecting the **Category** tab.
- **Refresh categories, tags, and field options:** **Category**, **Tags**, and **Field** options can be fetched.

2 Tag and field selection

In the 2 area, select **Tags** and **Field**. The default setting is "Select All," unless specified separately by user.

- **Tags:** Unique data that can identify the collected targets.
- **Field:** Metrics collected from the monitored targets.
- **Filter:** Select  **Filter by tag value** and then set the tag value for filtering.

For example, you can search filtered data by setting the value of `oname` to `demo-8101` .

-  **Search:** If you set the conditions and then select **Search**, you can search the source data of the metric in the 3 area.
-  : You can download the source data of the metric as a CSV file.

3 Metrics table

Because the metrics to be collected cannot be specified in advance, it is important to check the source data for all metrics being collected. After setting the desired conditions in the condition area above, you can search the original data of the metric in a table format in the 3 area. Columns in the table are changed when you specify the conditions for **Tags** and **Field**.

- To search metrics, **Time selector** and **Category** must have been specified.
- To search metrics, **Tags** and **Field** are optional.

Log Monitoring

Logs are the files that have recorded the events and messages that occurred in applications and systems. It is important to check the logs to find anomalies for preventing system deterioration or detecting the failure causes.

Modern IT service building environments are changing to MSA or Kubernetes. As a result, the targets to be managed are increasing. In typical log monitorings, the `tail` command or editor is used for checking after connecting the servers. However, in this case, it is difficult to use simple methods such as accessing individual servers one by one. How can you view the logs from hundreds or thousands of servers in severe cases?

Through the WhaTap log monitoring service, you can easily manage numerous logs.

! The Linux `tail` command is mainly used to check the logs whose content is added over time.

Key features

- **Central integrated management**

WhaTap can manage a large number of logs from the center. It is very convenient because you can check the log details without accessing individual servers.

- **Collection of all logs**

While collecting the logs selectively, important data may be missing. In general, WhaTap collects all logs. The collected logs can be checked in real time through the [Live Tail](#) menu.

- **Visibility**

You can check the logs through various charts provided by WhaTap. Through this visibility, you can increase access to errors and issues and detect failures as soon as possible.

- **Flexible capacity management**

While loading logs on individual servers, the logs may cause an excessive file system capacity being occupied. However, if the logs are collected through WhaTap, there is no need to maintain log files that occur on individual servers, and only the log maintenance period can be managed centrally.

- **Various analysis perspectives**

To find and predict failures, you have to analyze logs from various perspectives. WhaTap can check the log count trends with specific tags or the logs filtered by specific tags. For frequently used patterns, you can save them as charts so that you can search them at any time.

- **Pattern notification**

Once you have found a pattern of failures, you can set this as an alert to recognize or prevent problems as soon as possible. WhaTap provides [Real-time Log Alert](#) to receive notifications when specific keywords are included on an individual log basis and [Composite Log Alert](#) to receive notifications based on the count trend of logs containing specific tags.

Analyzing logs

Live Tail

Through [Live Tail](#), the logs collected in real time can be viewed as they flow through the screen, just like using the `tail` command. For more information about [Live Tail](#), see [the following](#).

Log Trend

Through [Log Trend](#), you can see all the collected logs or the log count trend with specific tags. If the log counts are closely related to the times of failure occurrence and resolution, the failure cause analysis and response can be accelerated through the log occurrence trend. For more information about the [log trend](#), see [the following](#).

Log Search


Through [Log Search](#), you can view all the collected logs or the logs with specific tags. The logs generated in a specific time period or specific server can be viewed based on tags. The Adjacent Log function that checks the logs that occurred before and after a selected log, can be used to see the situation in which a specific error or exception occurred. For more information about [Log Search](#), see [the following](#).

Applying The Log Monitoring


The following guides you to the basic application method for using the WhaTap log monitoring service. WhaTap log monitoring does not require configuring additional agents or applying a parser according to the logs. A simple configuration gets you started quickly.

Operating principle of log monitoring

Common log integration services



WhaTap log monitoring service



A typical log integration service consists of the collector, processor, repository, and UI module. The process of building each module is cumbersome and additional costs are incurred because it requires step-by-step configuration and structuring.

WhaTap log monitoring is easy to apply. Because the previous monitoring agent acts as a collector, the log monitoring can be started simply by turning on the agent option.

- [Java](#) [PHP](#) [Python](#) [Go](#) [Server](#) [Kubernetes](#)

- It directly collects the logs that are forwarded to the log library of your Java application rather than reading the output file.
- Because the logs are collected directly, it does not cause file I/O so the performance impact on the system is very low.
- You can check logs from the transaction traces by obtaining the connection traceability for transaction traces and logs.

! Java log library

Representative Java log libraries are Apache Log4j and Logback.

! They are available in Java Agent 2.1.1 or later.

- The log collection function has been added to the existing agent. The monitoring agent collects by using the output logs in addition to log files.

! They are available in PHP Agent 2.3.2 or later.

- The log collection function has been added to the existing agent. The monitoring agent collects by using the output logs in addition to log files.
- If the transaction ID is output to the log, you can check logs from the transaction traces by obtaining the connection traceability for transaction traces and logs.

! They are available in Python Agent 1.2.2 or later.

- The log collection function has been added to the existing agent. The monitoring agent collects by using the output logs in addition to log files.
- The log collection function has been added to the existing agent. The monitoring agent collects by using the output logs in addition to log files.

! They are available in Server Agent 2.1.2 or later.

- Logs can be collected in the Kubernetes container.
- You can collect logs from the applications in Kubernetes containers.

! They are available in Kubernetes Agent 1.1.35 or later.

Applying the log monitoring

The application method according to the used application is as follows. Before application of Log Monitoring, check the [supported version](#).

1. Check the version of the supported agent and perform [update](#).
2. Set the [options](#) for Log Monitoring.
3. **Enable** the log monitoring.

Java

Learn how to collect logs from your Java application.

PHP

Learn how to collect logs from your PHP application.

 Node.js

Learn how to collect logs from your Node.js application.

 Python

Learn how to collect logs from your Python application.

 Go

Learn how to collect logs from your Go application.

 Server

Learn how to collect logs from your server application.

 Kubernetes

The following guides you through collecting logs from Kubernetes containers and the applications in Kubernetes containers.

Java

To collect logs from your Java application, complete the following 3 steps:

Agent Update

It is available from Java agent 2.1.1. For more information about update, see [the following](#).

Checking the agent settings

Home > Select Project > **Log** > **Log Setting**

Select a project in the initial screen of the WhaTap monitoring service, and then select **Log** > **Log Setting** under Project Menu. Proceed with the process according to the instructions in the **Check agent settings** tab of the **Getting started with log monitoring** section.

1. Under Project Menu, select **Management** > **Agent CONFIG..**
2. In **Write options**, select **Direct Input**. Add the `logsink_enabled=true` option to the input field as follows:
3. Restart the application to apply the log monitoring.

Key options

- **hooklog_enabled** Boolean

Default `false`

It enables the log monitoring by hooking the log library.

 Before running the application, this option must have been enabled in `whatap.conf` to dynamically turn on or off the monitoring through the `logsink_enabled` setting. If the `logsink_enabled` option is set to `true` before running the application, log monitoring is possible without a separate setting.

ⓘ If there is a possibility to use the log monitoring, it is recommended to set this option in advance.

- **hooklog_custom_methods**

Register a custom log. An arbitrary log framework content is delivered. Use this to trace logs of the log module created on the website.

Java

```
package io.home.test;

public class MyLog {
    public void customLog(String log) { ... }
}
```

whatap.conf

```
hooklog_custom_methods=io.home.test.MyLog.customLog
```

- **logsink_enabled** Boolean

Default false

It turns on or off Log Monitoring.

ⓘ Before running the application, if in [whatap.conf](#), the `hooklog_enabled` option is set, this option can dynamically turn on or off the log monitoring.

- **logsink_trace_enabled** Boolean

Default false

Set whether or not to expose the log tab of the transaction trace by inserting the transaction ID into the logs.

Enabling the log monitoring

Home > Select Project > [Log](#) > [Log Setting](#)

You can enable or disable the WhaTap log monitoring with the toggle button in the [Activate log monitoring](#) tab of the [Starting the log monitoring](#) section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

Role

After agent installation, you can enable the log monitoring settings only when you have the [edit permission](#) on the project. For more information about the roles, see [the following](#).

PHP

To collect logs by your PHP application, check the following:

Agent Update

It is available from PHP agent 2.3.2. For more information about update, see [the following](#).

Checking the agent settings

Home > Select Project > [Log](#) > [Log Setting](#)

Select a project in the initial screen of the WhaTap monitoring service, and then select [Log](#) > [Log Setting](#) under Project Menu. Proceed with the process according to the instructions in the [Check Agent Settings](#) tab of the [Start Log Monitoring](#) section.

Checking the configuration file path

Check the **configuration file path** of the file for log monitoring. You can check it in the CLI environment or on the web.

Checking in the CLI environment

```
$ php -i | grep 'Scan'
Scan this dir for additional .ini files => /etc/php8.0.d
```

ⓘ If the path is "(None)," the configuration file path is [/usr/whatap/php](#).

Checking on the web

If the PHP settings of the CLI and Apache environments are different, check the results of the `phpinfo()`

function on the web. .

Entering the command

Log monitoring starts when you enter the command including the configuration file path.

```
export LOGFILES=/some/path/file1,/some/other/file2
echo "whatap.logsink_enabled=true" | sudo tee -a [configuration file path]/whatap.ini
echo "whatap.logsink.files=$LOGFILES" | sudo tee -a [configuration file path]/whatap.ini
```

Enabling the log monitoring

Home > Select Project > [Log](#) > [Log Setting](#)

You can enable or disable the WhaTap log monitoring with the toggle button in the [Activate log monitoring](#) tab of the [Starting the log monitoring](#) section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

Role

After agent installation, you can enable the log monitoring settings only when you have the [edit permission](#) on the project. For more information about the roles, see [the following](#).

Node.js

To collect logs by your Node.js application, check the following.

Agent Update

It is available from Node.js agent 0.5.1. For more information about the update, see [the following](#).

Checking the agent settings

Home > Select Project > Log > Log Setting

Select a project in the initial screen of the WhaTap monitoring service, and then select **Log > Log Setting** under Project Menu. Proceed with the process according to the instructions in the **Check agent settings** tab of the **Getting started with log monitoring** section.

1. Under Project Menu, select **Management > Agent CONFIG..**
2. In **Write options**, select **Direct Input**. Add the `logsink_enabled=true` option to the input field as follows:
3. Restart the application to apply the log monitoring.

Key options

- **logsink_enabled** Boolean

Default `false`

Set whether or not to enable the log monitoring.

- **logsink_trace_txid_enabled** Boolean

Default `true`

Set whether or not to expose the log tab of the transaction trace by inserting the transaction ID into the logs.

- **logsink_limit_content_enabled** Boolean

Default `true`

Set whether or not to limit the maximum size for log messages.

- `logsink_limit_content_length` **Number**

Default `10,000`

Set the maximum length of the log message. This value is applied only when the `logsink_limit_content_enabled` option is set to `true`.

Enabling the log monitoring

Home > Select Project > [Log](#) > [Log Setting](#)

You can enable or disable the WhaTap log monitoring with the toggle button in the [Activate log monitoring](#) tab of the [Starting the log monitoring](#) section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

Role

After agent installation, you can enable the log monitoring settings only when you have the [edit permission](#) on the project. For more information about the roles, see [the following](#).

Python

To collect logs by your Python application, check the following.

Agent Update

It is available from Python agent 1.2.2. For more information about update, see [the following](#).

Checking the agent settings

Home > Select Project > [Log](#) > [Log Setting](#)

Select a project in the initial screen of the WhaTap monitoring service, and then select [Log](#) > [Log Setting](#) under Project Menu. Proceed with the process according to the instructions in the [Check Agent Settings](#) tab of the [Start Log Monitoring](#) section.

Enabling the collection of logs

If you enter the command including the configuration file path ([WHATAP_HOME](#)), log collection starts.

```
export LOGFILES={Log file's full path},...  
echo "logsink_enabled=true" | sudo tee -a {Configuration file path}/whatap.conf  
echo "logsink.files=$LOGFILES" | sudo tee -a {Configuration file path}/whatap.conf
```

Linking the log and web transaction

Version 1.3.6 or later

Starting from the version 1.3.6, transactions and logs can be integrated using the following methods. Currently, WhaTap supports the Python's logging and loguru libraries. Configure [whatap.conf](#) according to your Python log libraries.

- logging module

logging

```
echo "trace_logging_enabled=true" | sudo tee -a {configuration file path}/whatap.conf
```

- loguru module

loguru

```
echo "trace_loguru_enabled=true" | sudo tee -a {configuration file path}/whatap.conf
```

Version 1.3.6 or earlier

The transaction ID `{txid}` is output to the log so that the logs generated by each transaction can be viewed separately. WhaTap Monitoring automatically injects `{txid}` into Python LogRecord so that `{txid}` can be output to the log file when setting the formatter.

```
settings.py
...
LOGGING = {
...
    'formatters': {
        ...
    },
    'handlers': {
        ...
    },
    'loggers': {
        ...
        '{logger_name}': {
            'handlers': [...],
            ...
        },
    },
}
```



```

try:
    import whatap.trace.mod.logging as whatap_logging
    if whatap_logging.logging_injection_processed:
        LOGGING['formatters']['whatap.formatter']={
            '()': 'django.utils.log.ServerFormatter',
            'format': '[{server_time}] -- {{ "@txid" : "{txid}" }} -- {message}',
            'style': '{',
        }
        LOGGING['handlers']['whatap']={
            'level': 'DEBUG',
            'class': 'logging.handlers.RotatingFileHandler',
            'filename': os.path.join(BASE_DIR, 'logs','whatap_log.log'),
            'formatter': 'whatap.formatter',
        }
        LOGGING['loggers'][Logger name]['handlers'].append('whatap')
except:
    pass
...

```

Enabling the log monitoring

Home > Select Project > [Log](#) > [Log Setting](#)

You can enable or disable the WhaTap log monitoring with the toggle button in the [Activate log monitoring](#) tab of the [Starting the log monitoring](#) section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

! Role

After agent installation, you can enable the log monitoring settings only when you have the [edit permission](#) on the project. For more information about the roles, see [the following](#).

Go

To collect logs by your Go application, check the following.

Checking the agent settings

Home > Select Project > [Log](#) > [Log Setting](#)

Select a project in the initial screen of the WhaTap monitoring service, and then select [Log](#) > [Log Setting](#) under Project Menu. Proceed with the process according to the instructions in the [Check Agent Settings](#) tab of the [Start Log Monitoring](#) section.

Enabling the collection of logs

If you enter the following command, log collection is started immediately.

```
export LOGFILES=/some/path/file1,/some/other/file2
echo "logsink_enabled=true" | sudo tee -a /usr/whatap/agent/whatap.conf
echo "logsink.files=$LOGFILES" | sudo tee -a /usr/whatap/agent/whatap.conf
```

Enabling the log monitoring

Home > Select Project > [Log](#) > [Log Setting](#)

You can enable or disable the WhaTap log monitoring with the toggle button in the [Activate log monitoring](#) tab of the [Starting the log monitoring](#) section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

Role

After agent installation, you can enable the log monitoring settings only when you have the [edit](#)

! permission on the project. For more information about the roles, see [the following](#).

Server

To collect logs by your server application, check the following:

Agent Update

It is available from Server agent 2.1.2. For more information about update, see [the following](#).

Checking the agent settings

Home > Select Project > Log > Log Settin_

Select a project in the initial screen of the WhaTap monitoring service, and then select **Log > Log Setting** under Project Menu. In the OS selection tab at the top of **Check Agent CONF.**, select a server OS (**Linux (shell)**, **Windows (Powershell)**). See **Basic Installation** or **Install with Category** for the process.

! The following configuration requires restart of the agent.

Linux Shell

- [Basic installation](#)

Linux Shell

```
export LOGFILES=/some/path/file1,/some/other/file2
echo "logsink.files=$LOGFILES" | sudo tee -a /usr/whatap/infra/conf/whatap.conf
```

- [Install with Category](#)

Linux Shell

```
1 cd /usr/whatap/infra
```

```

2  sudo mkdir extension
3
4  cat <<EOL | sudo tee extension/logsink.conf > /dev/null
5  [[inputs.logsink]]
6    category = "serverlog"
7    ## Whether to transfer the log generation statistics as separate data
8    stats_enabled = true
9    ## Log generation statistics category
10   stats_category = "logsink_stats"
11   ## Set the log file name to be excluded if the log file path contains asterisk (*).
12   excludeNames = [ ".gz", ".zip" ]
13   [[inputs.logsink.file]]
14     ## When setting a log file, you can set the date pattern (strftime.org).
15     path = "/some/path/%Y-%m-%d/.log"
16     disabled = false
17     encoding = "utf-8"
18
19   [[inputs.logsink.file]]
20     path = "/some/other/log"
21     disabled = false
22     encoding = "utf-8"
23
24   [[inputs.logsink.file]]
25     ## When searching for the keyword in the line-by-line log, merge it with the previous log.
26     nowrap_keywords = ["Caused by:", "Test"]
27
28   [[inputs.logsink.file]]
29     ## When entering the project code and project license, the logs can be transferred to a desiged project in the
30     same region.
31     # pcode = Project code
32     # license = "Project license"
33
34   [[inputs.logsink.file]]
35     # whatap_host_ip = "Collection server IP1/Collection server IP2"
36     # whatap_host_port = 6600
37
38   [[inputs.logsink.file]]
39     path = "/tmp/mixed.log"
40     disabled = false

```

```

40 encoding = "utf-8"
41 ## The [{ "A":1,"B":2},{ "C":3,"D":4},...] format can be divided into individual logs.{"A":1,"B":2} {"C":3,"D":4}
42 multiline_pattern = "simple_json_array"
43 EOL
44
45 sudo service whatap-infra restart

```

Windows Powershell

- [Basic installation](#)

Windows Powershell

```

$LOGFILES="c:\\whatap\\logs\\%Y-%m-%d\\*.log,c:\\whatap\\logs\\*.log"
Add-Content "c:\\Program Files\\WhatapInfra\\whatap.conf" -Value "logsink.files=$LOGFILES"

```

- [Install with Category](#)

Windows Powershell

```

1  # The administrator role required
2  New-Item -type "Directory" -Path "C:\\Program Files\\WhatapInfra\\extension"
3
4  $contentToAdd = @"
5  [[inputs.logsink]]
6  category = "serverlog"
7  ## Whether to transfer the log generation statistics as separate data
8  stats_enabled = true
9  ## Log generation statistics category
10 stats_category = "logsink_stats"
11 ## Set the log file name to be excluded when the log file path contains asterisk (*).
12 excludeNames = [ ".gz",".zip" ]
13 [[inputs.logsink.file]]
14 ## The date pattern (strftime.org) can be entered when setting the log file.
15 path = "c:\\whatap\\logs%Y-%m-%d_.log"
16 disabled = false
17 encoding = "euc-kr"
18

```

```

19 [[inputs.logsink.file]]
20 ## When searching for the keyword in the line-by-line logs, it merges with the previous log.
21 nowrap_keywords = ["Caused by:", "Test"]
22
23 [[inputs.logsink.file]]
24 ## When entering the project code and project license, logs can be transferred to a desired project in the same
   region.
25 # pcode = Project code
26 # license = "Project license"
27
28 [[inputs.logsink.file]]
29 # whatap_host_ip = "Collection server IP1/Collection server IP2"
30 # whatap_host_port = 6600
31
32 [[inputs.logsink.file]]
33 path = "/tmp/mixed.log"
34 disabled = false
35 encoding = "utf-8"
36 ## The [{"A":1,"B":2},{ "C":3,"D":4},...] format can be divided into individual logs.{"A":1,"B":2} {"C":3,"D":4}
37 multiline_pattern = "simple_json_array"
38
39 "@
40
41 New-Item -path "C:\Program Files\WhatapInfra\extension" -name "logsink.conf" -type "file" -value $contentToAdd
   -Force
42
43 Restart-Service "Whatap Infra"

```

Configuring options

- `stats_enabled` : Set whether or not to collect the collection status data. The default value is `false` . The value must be set to `true` . If set to `true` , statistical data is generated in the category (`logsink_stats`) set in the following `stats_category` .
- `stats_category` : Set the metrics category to store the collection status data in. The value must be set to `logsink_stats` . Statistical data fields are as follows:
 - `file`

- `checkInterval`
- `encoding`
- `filepos`
- `checkedLocalTime`
- `lastupdatedLocalTime`
- `fileSize`
- `error`
- `firstCheck`
- `transferBytes`
- `excludeNames` : By including asterisks (*) in the log file path, you can set the file name to exclude non log files. You can set multiple values using comma (,) as a separator.

Example

```
excludeNames = [ ".gz", ".zip" ]
```

- `nowrap_keywords` : While searching logs by line, if any keyword is detected by a specified option value, it is merged with the previous log.
- `pcode` : Enter the project code of the same region to which you want to send the logs.
- `license` : Enter the license code of the project in the same region to which you want to send the logs.

ⓘ In case of the `pcode` and `license` options, it is supported from Server Agent 2.7.4 or later.

Configuring the log options for Windows events

When collecting Windows event logs, you can set the options as follows:

```
1 # The administrator role required
2 New-Item -type "Directory" -Path "C:\Program Files\WhatapInfra\extension"
3
4 $contentToAdd = @"
5 [[inputs.win_eventlog]]
```



```

6  category = "win_event_log"
7  stats_category = "win_event_log_stats"
8  stats_enabled = true
9  enabled = true
10 [[Inputs.win_eventlog.file]]
11     #true | false
12     enabled = true
13     # Application, Security, Setup, System, ForwardedEvents
14     file = "Application"
15     #1: Information, 2: Warning 3: Error 4: Audit Success 5 Audit Fail
16     #event_type =
17     #event_id =
18     #event_id =
19     #event source name
20     #source_name = ""
21     #When entering the project code and project license, logs can be transferred to the desired project in the same region.
22     #pcode = Project code
23     #license = "Project license"
24     #whatap_host_ip = "Collection server IP1/Collection server IP2"
25     #whatap_host_port = 6600
26
27 "@
28
29 New-Item -path "C:\Program Files\WhatapInfra\extension" -name "win_eventlog.conf" -type "file" -value $contentToAdd
30 -Force
31 Restart-Service "Whatap Infra"

```

- It is required to specify the `category` .
e.g. `win_event_log`
- It is required to specify the statistics category (`stats_category`).
e.g. `win_event_log_stats`
- The statistics category On/Off (`stats_enabled`) is required.
e.g. `true` or `false`
- The collection On/Off (`enabled`) is required.

e.g. `true` or `false`

- The file-based collection On/Off (`enabled`) is required.

e.g. `true` or `false`

- The file (`file`) is required.

e.g. `Application` , `Security` , `Setup` , `System` , `Forwarded`

- The event type (`event_type`) is optional.

e.g. `1` , `2` , `3` , `4` , `5`

❗ Event type

1. Information
2. Warning
3. Error
4. Audit Success
5. Audit Fail

- The event ID (`event_id`) is optional.
- The event source name (`source_name`) is optional.
- Project code (`pcode`) optional
- Project license (`license`) optional

- ❗ • For Windows event log collection, the Server Agent 2.5.2 or later is supported.
- Among the Windows event log options, `pcode` and `license` are supported in the Server Agent 2.7.4 or later.

Specifying and separating multiple files

Specifying multiple files

To categorize multiple log files under the same category, add each log file path into the `[[inputs.logsink.file]]` entry and specify the file path by using the `path` setting. Additionally, you can configure additional settings, such as log generation statistics and options such as `category` , `stats_enabled` , and `stats_category` . The example is where the `/root/test1` and `/root/test2` files are assigned the same category (`serverlog`).

```
[[inputs.logsink]]
  category = "serverlog"
  stats_enabled = true
  stats_category = "logsink_stats"
  excludeNames = [ ".gz", ".zip" ]
```

```
[[inputs.logsink.file]]
  path = "/root/test1"
  disabled = false
  encoding = "utf-8"
```

```
[[inputs.logsink.file]]
  path = "/root/test2"
  disabled = false
  encoding = "utf-8"
```

❗ For information about options such as `category` , `stats_enabled` , `stats_category` , and `nowrap_keywords` , see the guidance [above](#).

In this example, you can see the following in **Live Tail**.

serverlog			
▼	oname		
▼	APPSVR	2024-10-11 18:14:30.325	file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709 kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh
▼	APPSVR	2024-10-11 18:14:30.325	file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709 kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh
▼	APPSVR	2024-10-11 18:14:30.325	file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709 kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh
▼	APPSVR	2024-10-11 18:14:30.325	file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709 kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh
▼	APPSVR	2024-10-11 18:14:30.325	file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709 kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh
▼	APPSVR	2024-10-11 18:14:30.325	file /root/test1 pcode 34 oname APPSVR agenttime 1728638069515 oid 1988910709 kdjfhakdhgldkjggha;dfghad;fkjggha;dfkjggha;dfgjha;fkghj;fgjh
▼	APPSVR	2024-10-11 18:14:40.326	file /root/test2 pcode 34 oname APPSVR agenttime 1728638079521 oid 1988910709 sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbvcmbn.bcmn

File separation

To specify different categories for log files, create separate `logsink{number}.conf` files and then set different `category` values in each file. The following example shows how to set different log categories through `logsink.conf` and `logsink1.conf`.

logsink.conf

```
# cat logsink.conf
[[inputs.logsink]]
  category = "serverlog1"
  stats_enabled = true
  stats_category = "logsink_stats"
  excludeNames = [ ".gz", ".zip" ]

[[inputs.logsink.file]]
  path = "/root/test1"
  disabled = false
  encoding = "utf-8"
```

logsink1.conf

```
#cat logsink1.conf
[[inputs.logsink]]
  category = "serverlog2"
  stats_enabled = true
  stats_category = "logsink_stats"
  excludeNames = [ ".gz", ".zip" ]

[[inputs.logsink.file]]
  path = "/root/test2"
  disabled = false
  encoding = "utf-8"
```

ⓘ For information about options such as `category` , `stats_enabled` , `stats_category` , and `nowrap_keywords` , see the guidance [above](#).

In this example, you can see the following in **Live Tail**.

FILTERS		VIEW	
All logs		SHOW ALL LOGS	
▼	oname	TIME	LOG
▼	APPSVR	2024-10-11 18:20:55.084	file /root/test1 pcode 34 oname APPSVR agenttime 1728638454313 oid 1988910709 category serverlog1 kdjfhakdhgldkjg ha,dfghad;fkjggha;dfkjggha;dfgjha;fkgjha;fgjh
▼	APPSVR	2024-10-11 18:20:55.087	file /root/test1 pcode 34 oname APPSVR agenttime 1728638454313 oid 1988910709 category serverlog1 kdjfhakdhgldkjg ha,dfghad;fkjggha;dfkjggha;dfgjha;fkgjha;fgjh
▼	APPSVR	2024-10-11 18:21:05.103	file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2 sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn
▼	APPSVR	2024-10-11 18:21:05.104	file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2 sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn
▼	APPSVR	2024-10-11 18:21:05.104	file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2 sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn
▼	APPSVR	2024-10-11 18:21:05.104	file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2 sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn
▼	APPSVR	2024-10-11 18:21:05.104	file /root/test2 pcode 34 oname APPSVR agenttime 1728638464315 oid 1988910709 category serverlog2 sadlkfj;kgj'fdnc/mbnv.c,mbn.cvmbn.vc,mnbcvmbn.bcmn

Enabling the log monitoring

Home > Select Project > Log > Log Setting

You can enable or disable the WhaTap log monitoring with the toggle button in the **Activate log monitoring** tab of the **Starting the log monitoring** section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

Role

After agent installation, you can enable the log monitoring settings only when you have the [edit permission](#) on the project. For more information about the roles, see [the following](#).

Kubernetes

To collect logs from Kubernetes containers and the applications in Kubernetes containers, check the following.

Agent Update

It is available from Kubernetes agent 1.1.35. For more information about update, see [the following](#).

Checking the agent settings

Home > Select Project > [Log](#) > [Log Setting](#)

Select a project in the initial screen of the WhaTap monitoring service, and then select [Log](#) > [Log Setting](#) under Project Menu. Proceed with the process according to the instructions in the [Check Agent CONF.](#) tab of the [Set up the agent and enable log monitoring](#) section.

Enabling the container log collection

To collect logs in Kubernetes containers, select [Apply log settings](#) in the [Check agent configuration](#) tab.

Enabling application log collection in the container

It is available from Java 2.1.1 and Python 1.2.2. To collect logs of the applications running on Kubernetes containers, see the following.

1. On the lower part of [Check agent settings](#), select **1**  [Agent CONFIG..](#)
2. Go to [Application Agent CONFIG.](#), and then select [Direct Input](#) in the **2** [Write Options](#) window.
3. Enter the [key](#), `logsink_enabled` and the [value](#), `true` of the following agent configuration command in the **3** field.

```
logsink_enabled=true
```

4. Select  **Apply**. You can collect application logs in the Kubernetes containers.

- ! • For more information about the collection of Java application logs, see [the following](#).
- For more information about the collection of Python application logs, see [the following](#).
- For more information about the collection of Go application logs, see [the following](#).

Enabling the log monitoring

Home > Select Project > [Log](#) > [Log Setting](#)

You can enable or disable the WhaTap log monitoring with the toggle button in the [Activate log monitoring](#) tab of the [Starting the log monitoring](#) section.

- If you turn on , the log monitoring is enabled. You can try for free for 15 days from the activation date.
- If you turn off , the log monitoring is disabled. Logs are no longer saved.

! Role

After agent installation, you can enable the log monitoring settings only when you have the [edit permission](#) on the project. For more information about the roles, see [the following](#).

Kubernetes log monitoring category guide

You can check various logs related to Kubernetes. The monitoring categories provided by WhaTap Kubernetes are as follows. Depending on the settings, duplicate log contents may be stored. Be sure to check for duplicates.


Category	Description
#K8sEvent	• Logs that store the events that occur in Kubernetes • Generation by default regardless of user settings

Category	Description
#WhatapEvent	• Logs that store the events that occur by WhaTap event settings • Generation by default regardless of user settings
containerStdout	• Container Standard Out logs • Generation by user settings • When <code>logsink_enabled=true</code> setting is added to the node agent
AppLog	• Application logs in the container • Generation by user settings • When <code>logsink_enabled=true</code> setting is added to the application agent

Setting the containerStdOut log filter

Home > Select Project > **Management** > **Agent CONF.**

The following guides you to filtering options to select information needed in the Kubernetes logs. You can set which logs to allow or exclude through filtering. In **Agent CONF.**, select **Direct Input** in the **Write options** manner.



Log naming conventions

The Kubernetes containerStdOut logs are generated under the `/var/log/containers/` path on the node. They keep the following naming conventions.

```
<podName>_<namespace>_<containerName>-<containerId>.log
```

Log filtering options

- **log_filter_enabled** `bool`

Default `false`

It sets whether to use log filters.

- **log_filter_allow_list** `list`

It sets the list of logs to allow. It indicates the items to be included for filtering.

- **log_filter_deny_list** `list`

It sets the list of logs to exclude. It indicates the items to be excluded for filtering.

Log filtering operation

Log filters work only when the `log_filter_enabled` option is active. You can selectively collect logs via `log_filter_allow_list` and `log_filter_deny_list` when the value of the option is set to `true`.

Log filters work on a blacklist basis. For example, if you enable `log_filter_enabled` and do not set the items to be added to the `log_filter_allow_list` option, the log collection for all containers is blocked by default.

Additionally, in the log filters, deny rules take precedence over allow rules. If the same items are set in both allow and deny rules, the deny rule takes precedence and logs are not collected.

Log filter example

Collecting logs for a single namespace

- Allow `log_filter_allow_list` `*_infra_*`

All containers that exist in the `infra` namespace are targeted for collection.

Collecting multiple namespace logs

- Allow `log_filter_allow_list` `*_infra_*, *_monitoring_*`

All containers that exist in the `infra` and `monitoring` namespaces are specified as collection targets.

Pods containing specific terms are excluded.

- Allow `log_filter_allow_list` `*_*_*`
- Deny `log_filter_deny_list` `*prod*_**`

All containers are targeted for collection, but the logs that contain `prod` in `podName` are excluded.

Pods containing specific terms are collected.

- Allow `log_filter_allow_list` `*prod*_**`

Only the logs that contain `prod` in `podName` are collected.

Multi-namespace logs are collected and Pods containing specific terms are excluded.

- Allow `log_filter_allow_list` `*_infra_*, *_monitoring_*`
- Deny `log_filter_deny_list` `*prod*_**`

The logs for the `infra` and `monitoring` namespaces are collected, and the logs that contain `prod` in `podName` are excluded.

The container logs that start with a specific term in the container name are excluded.

- Allow `log_filter_allow_list` `*_*_db-*`

The container logs whose `containerName` starts with `db` in all namespaces are excluded.

Live Tail

! Without the **Log Read** role, you cannot enter the menu.

Home > Select Project > **Log** > **Live Tail**

In the **Live Tail** menu, you can directly check the log data stream on the monitoring screen without connecting to the server console. Even with large volumes of complex logs, you can filter and quickly identify them using the **Filter** and **Highlight** features. The log data retrieval cycle is 2 seconds.

The major terms are as follows:

- **Category:** Unit for log collection and inquiry
- **Content:** Log messages
- **Search Key:** It is generated through the log parser settings.
- **Tag:** Key to search the collected logs

! You can change the column width by dragging an edge of the log table column.

Agent option

If the agent option is set, the log levels are collected and the colors for each log level appear as follows:

2023-12-18	14:31:02.563	[level]	INFO	[pcode]	2277	[agenttime]	1702877462042	[oid]	778873916	[category]	AppLog	[loggerName]	io.home.test.logback02starter.base.web.LogbackControllerGreeting
													INFO log in our greeting method.
2023-12-18	14:31:02.563	[level]	WARN	[pcode]	2277	[agenttime]	1702877462043	[oid]	778873916	[category]	AppLog	[loggerName]	io.home.test.logback02starter.base.web.LogbackControllerError
												[threadName]	http-nio-19090-exec-2
													io.home.test.logback02starter.base.errors.exception.ApiException: [2204] Process failure. Please try again.
2023-12-18	14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462043	[oid]	778873916	[category]	AppServer	[_event_status_]	error
												[_event_status_literal_name_]	level
													Servlet.service() for servlet [dispatcherServlet] in context with path [] threw exception [Request processing failed: io.home.test.logback02starter.base.error]
2023-12-18	14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462057	[oid]	778873916	[category]	AppServer	[_event_status_]	error
												[_event_status_literal_name_]	level
													Servlet.service() for servlet [dispatcherServlet] in context with path [] threw exception [Request processing failed: io.home.test.logback02starter.base.error]
2023-12-18	14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462081	[oid]	778873916	[category]	AppServer	[_event_status_]	error
												[_event_status_literal_name_]	level
													Servlet.service() for servlet [dispatcherServlet] in context with path [] threw exception [Request processing failed: io.home.test.logback02starter.base.error]
2023-12-18	14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462087	[oid]	778873916	[category]	AppServer	[_event_status_]	error
												[_event_status_literal_name_]	level
													Servlet.service() for servlet [dispatcherServlet] in context with path [] threw exception [Request processing failed: io.home.test.logback02starter.base.error]

! Setting the agent options



- The agent options are as follows:

```
# whatap.conf
weaving=log4j-2.17
weaving=logback-1.2.8
```

- In Java agent 2.2.22 or later, it can be used when setting log4j-2.17 or logback-1.2.8 in the weaving setting. It is required to restart the agent.
- The log level is determined based on the `level` and `type` among the parsed keywords. When there are keys parsed as `level` and `type`, and the parsed values have **FATAL**, **CRITICAL**, **ERROR**, **WARN**, **WARNING**, and **INFO**, the log level color appears.

1 Filter area

Apply Filter


When a filter is applied, the logs that match the specified conditions are filtered. You can enter multiple filters. OR (`||`) if the filters have the same tags, otherwise AND (`&&`).

You can specify a filter by directly entering a value in the input field or clicking the **Filter** field. Filter tags are entered in the following order: `Search key`, `Operator`, and `Search value`. If you select  **Search**, the data with the filter applied can be viewed log list area.



Guide UI

The **Guide UI** is provided under the input field as follows: You can close the **Guide UI** via the ESC key when your mouse cursor is over the input field or filter tag.



Entry of the search key, operator, and search value

- Upon entry of the **search key**, recommended values are provided by distinguishing for general index, reserved word index, and numeric index.
- When entering an **operator**, for a general index search key, **==** and **!=** options are provided at the bottom. In case of the numeric index, the **>**, **<**, **<=**, **>=**, **==**, and **!=** options are provided.
- Upon entry of the **search value** for matched search (**>**, **<**, **<=**, **>=**, **==**), it is highlighted in **blue**. In case of excluded search (**!=**), it is highlighted in **red**. Upon entry of the **search value**, you can search using the case-sensitive option.

- If the filter tags exceed 2 lines, you can fold it by selecting the **^ Fold** icon.
- After entering a filter tag, you can use the Shift and Tab keys simultaneously in the input field to go to the previous filter tag.



- After entering a filter tag, you can use the Tab key in the input field to go to the next filter tag.

Adding filter tags

To add a filter tag, type the desired text in the input field and press **Enter** or **Tab**.


You can also select a recommended value from the **Guide UI** below the input field by clicking it or navigating with the arrow keys.

Deleting filter tags

To remove a single tag, press **Backspace** or click the **X** icon on the tag. To clear all tags at once, click the **X** icon located on the right side of the input field.

Filter favorites

The filter **favorites** feature is provided. You can select the ☆ **Favorites** icon on the right of the input window to **add**, **delete**, and select desired filter search conditions in the favorites. **Favorites** can store up to 50 cases.



Exceptional cases for application of filters

- In the tag with a numeric index (`search key` that ends with `.n`), only numbers can be input for the `search value` .
- Duplicate 'search key' and 'search value' cannot be entered.
- Searching is not possible if there is a tag whose `search key` or `search value` is missing. Invalid tags appear in gray.

ⓘ You cannot enter a `category` with the **Live Tail** `search key` .
The expression under the entered filter value is a preview of the filter expression for querying the log data.

Applying non-parsing keyword filters

You can search the logs containing keywords that have not been parsed, or for which no index has been created.

In this case, all logs in the specified range are fully scanned. Therefore, the search speed may be slower compared to the indexed key. For standardized log data, it is recommended to search using the index key through **Log parser setting**.

Filter ⓘ

🔍
oid != -1009464250 ✕
okind == -398596773 ✕
content == *select* ✕
Enter Filter and press enter to add.
✕

oid != -1009464250 && okind == -398596773 && content == *select*

1. Select **Category**.
2. In the **Filter** field, enter a desired keyword for a search after leaving spaces on the **content** basis.
e.g. `content *select*`
3. Click 🔍 **Search** to view the logs.

- ⓘ • In case of **Live Tail**, all logs can be viewed. Thus, the category is not required to be specified.
- For more information about the parser configuration, see [the following](#).

Modifying the filter

Enter a value in the filter and then click the specified value to modify it.



- You can modify the filter by entering text again in the input field.
- You can modify the recommended values by selecting them in the **Guide UI** below the input field.

Search Key

In the following image, the blue boxed part is the search key for parsing. In **Log Configuration**, you can register and set the parsing logic in **Log parser setting** tab with the search key.

Search key

Filter input syntax

A tag consists of a search key and a search value. In the following example, the search key is `exception` and the search value is `UnknownHostException`. The example searches for the log data that includes an exception (`UnknownHostException`) that is generated when the server cannot connect the host because the IP address and domain address do not match among the collected log data.

exception == UnknownHostException

Search key types

Search key types	Search key format	Meaning	Example for Search key and search value	Search example
String keyword	keyword	File Name	- Key: fileName - Value: /data/whatap/logs/yard.log	fileName:/data/whatap/logs/yard.log
Numeric keyword	keyword.n	Response Time	- Key: response_time.n - Value: 2945	response_time.n>=2945
Reserved keyword (predefined keyword)	@keyword	Transaction ID	- Key: @txid - Value: 85459614215434144	-
Log body's keyword	content	Log body	- Key: content - Value: user input	content: *ERROR*

ⓘ Content search key

- The content search key searches the bodies of unindexed logs. For example, if you enter `content *ERROR*`, the logs containing `ERROR` are searched among the log bodies.
- If you do not know which keywords need to be used for indexing, use the content search keys to identify the logs containing problematic keywords. Afterwards, you can improve the search performance by setting a parser with the keyword through the log parser settings in **Log Configuration** and creating indexes.

Common syntaxes

Syntax type	Description	Example
<code>==searchValue</code>	Searches for the logs that match the search value.	<code>exception==RuntimeExceptionexception</code>

Syntax type	Description	Example
<code>!=searchValue</code>	Searches for the logs that exclude the search value.	<code>exception!=RuntimeException</code>
<code>*searchValue</code>	Searches for the logs ending with the search value.	<code>word==*hello</code>
<code>searchValue*</code>	Searches for the logs starting with the search value.	<code>word==hello*</code>
<code>*searchValue*</code>	Searches for the logs that contain the search value in the middle.	<code>word==*hello*</code>
<code>*search*Value*</code>	Searches for the logs that contain the search value.	<code>word==*he*llo*</code>
<code>re:{regex}</code>	Searches for the logs that match a regular expression.	<code>caller==re:^(i\.w\.a\.w\.s\.v\.r\.</code>
<code>**</code>	Searches all logs that match the search key.	

Syntax in case the search key has a numeric value (keyword.n)

The following syntax is supported only when the search key is in `keyword.n` format.

- Search values must only be numbers.
- No prefix is attached to the value of the `.n` keyword. All keywords other than `.n` must be attached with a prefix.

For example, `+>searchValue` is invalid.

Syntax type	Description	Example
<code>>searchValue</code>	Searches for the logs containing values greater than the search value.	<code>response_time.n>3000</code>

Syntax type	Description	Example
<code>>=searchValue</code>	Searches for the logs containing values greater than or equal to the search value.	<code>response_time.n>=3000</code>
<code>==searchValue</code>	Searches for the logs containing the same search value.	<code>response_time.n==3000</code>
<code>!=searchValue</code>	Searches for the logs that contain values different from the search value.	<code>response_time.n!=3000</code>
<code><searchValue</code>	Searches for the logs containing values less than the search value.	<code>response_time.n<3000</code>
<code><=searchValue</code>	Searches for the logs containing values less than or equal to the search value.	<code>response_time.n<=3000</code>

Viewing personal information

If you have the **log personal information query** role, you can display de-identified personal information by enabling the **Display personal information**  toggle button.

- For more information about the **log personal information query** role, see [the following](#).
- For more information about **De-identification of personal information**, see [the following](#).

2 Content highlight area

The highlight function allows you to easily identify desired keywords in log content. You can filter with single or multiple keywords.

Live Tail

Category

ApplLog

Filter

content

+*select*

×

Enter Filter and press enter to add.

Q

content == +*select*

a

select

Q

⏏

↺

↻

⌵

⚙

	oname	Timestamp	Log
b	demo-8105	2023-11-24 14:08:35.620	@txid 5935183317842444630 pcode 5490 oname demo-8105 onodeName node-1 oid -857948929 okindName demo-okind-1 okind -628198688 select
	demo-8105	2023-11-24 14:08:35.624	@txid 4922255873619405363 pcode 5490 oname demo-8105 onodeName node-1 oid -857948929 okindName demo-okind-1 okind -628198688 select productmodelid, avg(listprice) as [average list price]
	demo-8105	2023-11-24 14:08:35.630	@txid 4922255873619405363 pcode 5490 oname demo-8105 onodeName node-1 oid -857948929 okindName demo-okind-1 okind -628198688 select productid, avg(unitprice) as [average price]
	demo-8105	2023-11-24 14:08:35.635	@txid 4922255873619405363 pcode 5490 oname demo-8105 onodeName node-1 oid -857948929 okindName demo-okind-1 okind -628198688 select

- Enter the keyword you want to highlight in the keyword input field.
 - Example: `select`
- The searched keyword will be highlighted in the Content of the log list.
 - If you click the `⌵` **Log full screen** icon, you can view **Log** and **Timestamp** in fullscreen.
 - You can navigate between highlighted keywords using shortcuts.
 - `Enter` : Move to the next row
 - `Shift+Enter` : Move to the previous row
 - `^` / `↓` icons: Move between rows
 - `Cmd + F` (Mac) / `Ctrl + F` (Windows): Search keywords

Multi-keyword condition

When highlighting with multiple keywords, write as follows:

String to input	Description	Result
a b c	Separate each keyword with a space.	a, b, c
"Whatap is good."	To include spaces in keywords, wrap them in <code>'</code> or <code>"</code> .	Whatap is good.
"Whatap\\ is good."	If <code>\</code> is included in the keyword surrounded by <code>"</code> , enter it as <code>\\</code> .	Whatap\ is good.

Setting the highlight color

By selecting  **Highlight**, you can set the keyword and its color for highlighting.

- In addition, enter the keyword to set the color for.
- If you click **Color** on the left of the input field, available colors appear.
- Basically, different highlighting by the level (**WARN**, **ERROR**, **FATAL**) is applied.
- The settings are saved in **Project Rules**.

Table settings

Column settings

Right of the  area, select  **Column settings**. You can add columns or set order.

- **Adding columns**

You can add columns to the table by selecting tags. If you unselect the **log** column, you cannot see **Log view detail setting**. Be sure to select at least one column.

- **Setting the column order**

When you add a column, the column is added to **Display columns**. Drag a desired column to change the order of the columns.

Log view detail setting

On the right of , select  **Log view detail setting**. **content** and **Tag** are all checked and two items are also displayed. Be sure to select any of **content** and **Tag**.

Unchecked items are not displayed in the table. If you deselect the **Tag** as follows, the log's **Tag** does not appear in the table.

Log table tag invisible

When you add tags in the **Tag management** list, the tags in the log are listed in the order you added. You can change the order of the tags by dragging them. If you disable an added tag, the disabled tag does not appear.

- **Column settings** and **Log view detail setting** can be used in **Live Tail**, **Log Search**, and **Log Trend**.
- In the same project, the **Live Tail**, **Log Search**, and **Log Trend** menus share **Column settings** and **Log view detail setting**.

Log Explorer


ⓘ Without the **Log Read** role, you cannot enter the menu.

Home screen > Select Project > **Log** > **Log Explorer** *New*

The **Log Explorer** feature allows you to quickly search, filter, and analyze log data. You can understand the field structure, perform query-based filtering to efficiently search for logs, and visualize the results. This feature is based on the Lucene index and offers the following key functions:

1. **Log trend analysis and search:** Enables both log trend visualization and search, helping you understand log patterns over time.
2. **Fast search and filtering:** Search for problematic logs quickly using simple keywords. Supports both WhaTap log search query and Lucene query syntax for defining various search conditions.
3. **Autocomplete functionality:** Helps users write and learn queries easily through auto-suggestions.
4. **Data visualization:** Uses stacked graphs to visualize the volume of logs over time. Drill-down capability allows deep analysis of specific logs.
5. **Favorites and sharing:** Save frequently used filters as favorites and share them easily via URL.

Basic Screen Overview



Time Range Selection

Set the time range for the log search using the **Time Selector** in the top option bar.

Category

Select the type of log you want to explore in **Select Category** on the top option bar (e.g., AppLog).

Filter

In the **Filter** on the top option bar, select the query language and enter the search conditions to apply the filter. The search query syntax supports both WhaTap Log Search Query and Lucene Query.

-  : WhaTap Log Search Query
-  : Lucene Query

Search Query Syntax

WhaTap log search query

WhaTap log search queries support autocomplete and allow quick log searches with simple keyword input. You can define various search conditions using operators like `and` , `or` , etc.

- Example: Entering `error` in the search filter will search within the log content.
 1. In the top option bar, click the filter icon and select a field from the query search box (fields are shown based on selected time and category).
 2. Choose an operator:
 - a. Default: `:` , `.*` , `and` , `or`
 - b. If the field is of type `Number` : Default + `<=` , `>=` , `<` , `>`
 - c. If the field is of type `String` : Default only
 3. Select a value (values shown based on selected time, category, and field).

Search Index Field Types

Only three field types are supported: `Number` , `String` , and `Boolean` .

- `Number` : Numeric values like long, double, float
- `String` : Full text such as email body
- `Boolean` : A value of either True or False

 For more details about WhaTap log search query syntax, refer to the [documentation](#).

Lucene

Use the standard syntax of Apache Lucene version 8.11.1. Autocomplete is not supported.

 For more information on the Lucene QueryParser Classic package, see the [official documentation](#).

Field List

Can be used for drill-down analysis.

Apply Field Info to Filter

Check field names in the **Field List** and click the `=` or `!=` buttons to apply them as filter conditions.

- Hovering over a field shows the `=` and `!=` buttons, which can be clicked to apply filters.
- Use the field search function to quickly find the desired field.

Trend Chart

Use **Trend Chart** to view logs in graph form with **interval** and **breakdown** applied.

- **Graph Interaction:** Zoom in using mouse drag or the  button, and click bars to view logs in that area
- **Interval:** Choose from Auto, 1s, 5s, 30s, 1m, 5m, 10m, 30m, 1h
- **Breakdown:** Visualize distribution of Top 3 values + Others for a specific field, and apply drill-down filters

Breakdown

Breakdown is a visualization feature that analyzes log count trends by a specific field.

- Displays a stacked graph separating the Top 3 values and Others for a given field
- Useful for drill-down analysis. You can click parts of the legend or stacked graph to highlight specific values or add filter conditions.

View Logs

You can view logs that match your search conditions. In the **Field List**, field data and log tags are displayed in matching highlight colors.

Content Highlight

The highlight function allows you to easily identify desired keywords in log content. You can filter with single or multiple keywords.

1. Enter the keyword you want to highlight in the keyword input field.
 - Example: `select`
2. The searched keyword will be highlighted in the Content of the log list.
 - If you click the  **Log full screen** icon, you can view **Log** and **Timestamp** in fullscreen.
 - You can navigate between highlighted keywords using shortcuts.
 - `Enter` : Move to the next row
 - `Shift+Enter` : Move to the previous row
 - `^` / `↓` icons: Move between rows
 - `Cmd + F` (Mac) / `Ctrl + F` (Windows): Search keywords

Multi-keyword condition

When highlighting with multiple keywords, write as follows:

String to input	Description	Result
<code>a b c</code>	Separate each keyword with a space.	a, b, c
<code>"Whatap is good."</code>	To include spaces in keywords, wrap them in <code>'</code> or <code>""</code> .	Whatap is good.
<code>"Whatap\\ is good."</code>	If <code>\</code> is included in the keyword surrounded by <code>""</code> , enter it as <code>\\</code> .	Whatap\ is good.

Setting the highlight color

By selecting  **Highlight**, you can set the keyword and its color for highlighting.

- In addition, enter the keyword to set the color for.
- If you click **Color** on the left of the input field, available colors appear.
- Basically, different highlighting by the level (**WARN**, **ERROR**, **FATAL**) is applied.
- The settings are saved in **Project Rules**.

Column settings

Click the  button to add columns or change their order.

- **Add Columns**

In **Column Settings**, select tags to add them as columns in the table.

 If no log columns are selected, **Log view detail setting** will not be visible. At least one column must be selected.

- **Reorder Columns**

Added columns will appear in **Display columns**. Drag them to change the display order.

Log view detail setting

Click the  button to configure log display options. By default, both **content** and **Tag** are selected and shown.

- Either **content** or **Tag** must be selected.
- Tags added in **Tag management** appear in order in the log display. You can reorder tags by dragging. Disabled tags are not shown.
- If an item is unchecked, it is not shown in the table. For example, if **Tag** is unchecked, log tags are not displayed.

log table contents remove tag sc

! Scope of Column and Log Display Settings

- **Column settings** and **Log view detail setting** can be used in **Live Tail**, **Log Search**, and **Log Trend**.
- In the same project, **Live Tail**, **Log Search**, and **Log Trend** share **Column settings** and **Log view detail setting**.

Download Log File

You can download the collected logs that match your search results in [CSV](#) or [TXT](#) file formats.

1. Click the  **Download** button at the right center of the screen.
2. Click **Download CSV** or **Download Log** to download the file in your preferred format.

WhaTap Log Search Query Syntax

✓ This document describes the WhaTap log search query syntax used in the Log Explorer.

1. Equal

- `field : value`

2. Not

- `not field : value`

3. Wildcard (Only applicable to `String`-type fields; do not use `"` around the value)

- `field : val*`

4. Exist

- `field : *`

5. Range (Only applicable to `Number`-type fields)

- Operators supported: `<`, `>`, `<=`, `>=`
- Example: `field < 30 and field >= 10`

6. Combining Multiple Queries

- AND example: `field < 30 and field >= 10`
- OR example: `field < 30 or field >= 10`

7. Grouping

- Example: `(field < 30 or field >= 10) and other : 70`

8. Escaping Special Characters

- Special characters supported by Lucene escape: `&` `&` `|` `|` `!` `(` `)` `{` `}` `^` `"` `~` `*` `?` `:` `\` `/`
- Use `\` before the character to escape
- Example: To search `(1+1):2`, use `\\(1\\+1\\)\\:2`

Search Index Field Types

Supported field types: `Number` and `String`

- `Number` : Numeric values such as long, double, float
 - Use syntax like `field : "value"`
- `String` : Full text such as email body
 - Use syntax like `field : value`
- `Boolean` : True or False
 - Use as `field : true` or `field : false`

Log Trend


ⓘ Without the **Log Read** role, you cannot enter the menu.

Home > Select Project > **Log** > **Log Trend**

In **Log Trend**, you can check the occurrence pattern of a specific error type through the occurrence count trend of the logs classified by type and see the log data details by time. You can quickly see the desired log through the highlighting feature. You can also check the trends of the collected logs by category.

The major terms are as follows:

- **Category:** Unit for log collection and inquiry
- **Content:** Log messages
- **Search Key:** It is generated through the log parser settings.
- **Tag:** Key to search the collected logs



Request Data

- If the scroll reaches the bottom, the next data is retrieved.
- In 1, specify the **time** and collection unit, **category**, of the log data to be searched.
- If the **category** is changed, the logs that correspond to the selected category are searched. You can see the logs in 2 bar chart and 3 log table.
- If you click a bar in 2 bar chart, the logs are viewed that correspond to the time range of the bar.
- 3 You can see the total number of logs viewed from the upper left of the log table.
- On the upper right of the 3 log table, select the  **full screen** icon, you can see the **log** and **collection time** in the full screen.

- **4** In Side menu, you can see the logs filtered by tags. Up to two search keys can be selected, and multiple search values can be selected.
- If the agent option is set, the log levels are collected and the colors for each log level appear as follows:

2023-12-18 14:31:02.563	level	INFO	pcode	2277	agenttime	1702877462042	oid	778873916	category	AppLog	loggerName	io.home.test.logback02starter.base.web.LogbackControllerGreeting
INFO log in our greeting method.												
2023-12-18 14:31:02.563	level	WARN	pcode	2277	agenttime	1702877462043	oid	778873916	category	AppLog	loggerName	io.home.test.logback02starter.base.web.LogbackControllerError
io.home.test.logback02starter.base.errors.exception.ApiException: [2204] Process failure. Please try again.												
2023-12-18 14:31:02.586	level	error	pcode	2277	agenttime	1702877462043	oid	778873916	category	AppServer	event_status	error
Servlet.service() for servlet [DispatcherServlet] in context with path [/] threw exception [Request processing failed: io.home.test.logback02starter.base.error]												
2023-12-18 14:31:02.586	level	error	pcode	2277	agenttime	1702877462057	oid	778873916	category	AppServer	event_status	error
Servlet.service() for servlet [DispatcherServlet] in context with path [/] threw exception [Request processing failed: io.home.test.logback02starter.base.error]												
2023-12-18 14:31:02.586	level	error	pcode	2277	agenttime	1702877462081	oid	778873916	category	AppServer	event_status	error
Servlet.service() for servlet [DispatcherServlet] in context with path [/] threw exception [Request processing failed: io.home.test.logback02starter.base.error]												
2023-12-18 14:31:02.586	level	error	pcode	2277	agenttime	1702877462087	oid	778873916	category	AppServer	event_status	error
Servlet.service() for servlet [DispatcherServlet] in context with path [/] threw exception [Request processing failed: io.home.test.logback02starter.base.error]												

① Setting the agent options

- The agent options are as follows:

```
# whatap.conf
weaving=log4j-2.17
weaving=logback-1.2.8
```

- In Java agent 2.2.22 or later, it can be used when setting log4j-2.17 or logback-1.2.8 in the weaving setting. It is required to restart the agent.
- The log level is determined based on the `level` and `type` among the parsed keywords. When there are keys parsed as `level` and `type`, and the parsed values have **FATAL**, **CRITICAL**, **ERROR**, **WARN**, **WARNING**, and **INFO**, the log level color appears.

- ① You can change the column width by dragging an edge of the log table column.

Viewing personal information

If you have the **log personal information query** role, you can display de-identified personal information by enabling the **Display personal information**  toggle button.

- For more information about the **log personal information query** role, see [the following](#).
- For more information about **De-identification of personal information**, see [the following](#).

Checking the log content

! What is content?

Content means the log messages.

- The first line of the log column is the log's parsed key and value, and the second line is the log's content.
- 3 For each row (log) of the log table, the ▶ **More** button exists. If you select the ▶ **More** button, you can check the entire content of the log like 5.

Configure Highlight

- The highlight function is provided to easily identify the desired keywords.
- 3 In the [input field](#) on the right, enter the keyword to highlight. You can check the highlighted keywords like this.

Content Highlight

The highlight function allows you to easily identify desired keywords in log content. You can filter with single or multiple keywords.

1. Enter the keyword you want to highlight in the keyword input field.
 - Example: `select`
2. The searched keyword will be highlighted in the Content of the log list.
 - If you click the [] **Log full screen** icon, you can view **Log** and **Timestamp** in fullscreen.
 - You can navigate between highlighted keywords using shortcuts.
 - `Enter` : Move to the next row
 - `Shift+Enter` : Move to the previous row
 - `^ / v` icons: Move between rows
 - `Cmd + F` (Mac) / `Ctrl + F` (Windows): Search keywords

Multi-keyword condition

When highlighting with multiple keywords, write as follows:


String to input	Description	Result
a b c	Separate each keyword with a space.	a, b, c
"Whatap is good."	To include spaces in keywords, wrap them in "'" or "\"".	Whatap is good.
"Whatap\\ is good."	If \ is included in the keyword surrounded by "", enter it as \\.	Whatap\ is good.

Setting the highlight color

By selecting  **Highlight**, you can set the keyword and its color for highlighting.

- In addition, enter the keyword to set the color for.
- If you click **Color** on the left of the input field, available colors appear.
- Basically, different highlighting by the level (**WARN**, **ERROR**, **FATAL**) is applied.
- The settings are saved in **Project Rules**.

Viewing the logs with a chart



- You can click or drag the desired time range bar on the bar chart (a) to view the logs for the time range.
- You can use the (b) time selection option on the upper left of the log table under the bar chart to further refine your log search in the selected time zone:
 - 1min: interval (spacing between bars on the chart)
 - Select a time option: Time zone whose time range is divided into 6 segments.

Setting the table

Column settings

Right of the (3) area, select **Column settings**. You can add columns or set order.

- **Adding columns**

You can add columns to the table by selecting tags. If you unselect the **log** column, you cannot see **Log view detail setting**. Be sure to select at least one column.

- **Setting the column order**

When you add a column, the column is added to **Display columns**. Drag a desired column to change the order of the columns.

Log view detail setting

On the right of 3, select  **Log view detail setting**. **content** and **Tag** are all checked and two items are also displayed. Be sure to select any of **content** and **Tag**.

Unchecked items are not displayed in the table. If you deselect the **Tag** as follows, the log's **Tag** does not appear in the table.

Log table tag invisible

When you add tags in the **Tag management** list, the tags in the log are listed in the order you added. You can change the order of the tags by dragging them. If you disable an added tag, the disabled tag does not appear.

- **Column settings** and **Log view detail setting** can be used in **Live Tail**, **Log Search**, and **Log Trend**.
- In the same project, the **Live Tail**, **Log Search**, and **Log Trend** menus share **Column settings** and **Log view detail setting**.

Log Search

ⓘ Without the **Log Read** role, you cannot enter the menu.

Home > Select Project > **Log** > **Log Search**

In **Log Search**, you can search a large number of integrated logs under various conditions and specify the desired logs. Multiple search conditions can be specified with parsed keys and values, so only the log data that matches the desired conditions are extracted.

The searched log data is imported to the dynamic page line by line. If the scroll reaches the bottom, the next data is automatically retrieved and displayed.

The major terms are as follows:

- **Category:** Unit for log collection and inquiry
- **Content:** Log messages
- **Search Key:** It is generated through the log parser settings.
- **Tag:** Key to search the collected logs

1 Log Search

Time: < 2023/08/30 11:22 ~ 2023/08/30 12:22 60Min > Filter: category AppLog Enter Filter and press enter to add. category == AppLog

2 AppLog (4,517,902)

AppLog @txid oid okind okindName oname onode onodeName

AppStdErr (461)

AppStdErr @txid oid okind okindName oname onode onodeName

AppStdOut (19,839)

AppStdOut oid okind okindName oname onode onodeName

3 Of the total 4,538,718 results, 9999+ match your search. Oldest Newest Display tag Please input the Keyword.

▶	oname	Timestamp	Log
▶	demo-8101	2023-08-30 11:22:00.019	@txid -2940368816900096605 pcode 5490 oname demo-8101 onodeName node-1 http://127.0.0.1:8101/remote/account/pickup/dept/daegu status=200
▶	demo-8101	2023-08-30 11:22:00.025	@txid -6539276394531757588 pcode 5490 oname demo-8101 onodeName node-1 select ename, job, sal + 100 from emp
▼	demo-8101	2023-08-30 11:22:00.020	@txid -877561626 pcode 5490 oname demo-8101 onodeName node-1 oid -877561626 category AppLog okindName demo-okind-1 okind -628198688 onode 336434079 select distinct pp.lastname, pp.firstname from person.person pp join humanresources.employee e on e
▶	demo-8101	2023-08-30 11:22:00.032	@txid -6539276394531757588 pcode 5490 oname demo-8101 onodeName node-1 select productid, avg(orderqty) as averagequantity, sum(linetotal) as total
▶	demo-8101	2023-08-30 11:22:00.038	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select distinct ename, deptno, sal, job from emp
▶	demo-8101	2023-08-30 11:22:00.032	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select e.ename, d.dname
▶	demo-8101	2023-08-30 11:22:00.043	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select productmodelid, name
▶	demo-8101	2023-08-30 11:22:00.049	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select ename, sal from emp
▶	demo-8101	2023-08-30 11:22:00.054	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select ename, sal+1000 from emp
▶	demo-8101	2023-08-30 11:22:00.062	@txid -129020984714615108 pcode 5490 oname demo-8101 onodeName node-1 select employee_id, dept, name, age, salary
▶	demo-8101	2023-08-30 11:22:00.060	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select ename, job, sal + 100 from emp
▶	demo-8101	2023-08-30 11:22:00.067	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select p.name as productname,
▶	demo-8101	2023-08-30 11:22:00.067	@txid -129020984714615108 pcode 5490 oname demo-8101 onodeName node-1 select e.ename ' departments's name is:' d.dname
▶	demo-8101	2023-08-30 11:22:00.072	@txid -3664464458700254001 pcode 5490 oname demo-8101 onodeName node-1 select salesorderid, sum(linetotal) as subtotal

4

Request Data

- If the scroll reaches the bottom, the next data is retrieved. You can view 10,000 logs at a time.
- **3** You can see the total number of logs viewed from the upper left of the log table.
- Log data can be queried in chronological order. On the upper right of the **3** log table, select a desired search method among **Oldest** and **Newest**.
- After specifying the time range, select **Apply** to finish the search time setting, and then select **Search** to search for data.
- On the upper right of the **3** log table, select the **Log full screen** icon, you can see the **Log** and

Timestamp in the full screen.

- If the agent option is set, the log levels are collected and the colors for each log level appear as follows:

2023-12-18 14:31:02.563	[level]	INFO	[pcode]	2277	[agenttime]	1702877462042	[oid]	778873916	[category]	AppLog	[loggerName]	io.home.test.logback02starter.base.web.LogbackControllerGreeting
2023-12-18 14:31:02.563	[level]	WARN	[pcode]	2277	[agenttime]	1702877462043	[oid]	778873916	[category]	AppLog	[loggerName]	io.home.test.logback02starter.base.web.LogbackControllerError
2023-12-18 14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462043	[oid]	778873916	[category]	AppServer	[event_status]	error
2023-12-18 14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462057	[oid]	778873916	[category]	AppServer	[event_status]	error
2023-12-18 14:31:02.586	[level]	error	[pcode]	2277	[agenttime]	1702877462087	[oid]	778873916	[category]	AppServer	[event_status]	error

! Setting the agent options

- The agent options are as follows:

```
# whatap.conf
weaving=log4j-2.17
weaving=logback-1.2.8
```

- In Java agent 2.2.22 or later, it can be used when setting log4j-2.17 or logback-1.2.8 in the weaving setting. It is required to restart the agent.
- The log level is determined based on the `level` and `type` among the parsed keywords. When there are keys parsed as `level` and `type`, and the parsed values have **FATAL**, **CRITICAL**, **ERROR**, **WARN**, **WARNING**, and **INFO**, the log level color appears.

- ! You can change the column width by dragging an edge of the log table column.

Checking the log content

! What is content?

Content means the log messages.

- The first line of the log column is the log's parsed key and value, and the second line is the log's content.

- **3** For each row (log) of the log table, the **More** button exists. If you select the **More** button, you can check the entire content of the log like **4**.
- If you select a tag in the log, a drop-down menu appears that allows you to perform **Copy**, **Search**, **Exclusion search**, and **Adjacent Log**.

Filter


Apply Filter

In **Time selector** on the left of **1**, you can specify the time range. On the right, When a filter is applied, the logs that match the specified conditions are filtered. You can enter multiple filters. OR (||) if the filters have the same tags, otherwise AND (&&).

You can specify a filter by directly entering a value in the input field or clicking the **Filter** field. Filter tags are entered in the following order: **Search key** , **Operator** , and **Search value** . If you select **Search**, the data with the filter applied can be viewed log list area.

Guide UI

The **Guide UI** is provided under the input field as follows: You can close the **Guide UI** via the ESC key when your mouse cursor is over the input field or filter tag.



Entry of the search key, operator, and search value

- Upon entry of the **search key**, recommended values are provided by distinguishing for general index, reserved word index, and numeric index.
- When entering an **operator**, for a general index search key, **==** and **!=** options are provided at the bottom. In case of the numeric index, the **>**, **<**, **<=**, **>=**, **==**, and **!=** options are provided.
- Upon entry of the **search value** for matched search (**>**, **<**, **<=**, **>=**, **==**), it is highlighted in **blue**. In case of excluded search (**!=**), it is highlighted in **red**. Upon entry of the **search value**, you can search using the case-sensitive option.

- If the filter tags exceed 2 lines, you can fold it by selecting the **^ Fold** icon.
- After entering a filter tag, you can use the Shift and Tab keys simultaneously in the input field to go to the previous filter tag.



- After entering a filter tag, you can use the Tab key in the input field to go to the next filter tag.

Adding filter tags

To add a filter tag, type the desired text in the input field and press **Enter** or **Tab**.


You can also select a recommended value from the **Guide UI** below the input field by clicking it or navigating with the arrow keys.

Deleting filter tags

To remove a single tag, press **Backspace** or click the **X** icon on the tag. To clear all tags at once, click the **X** icon located on the right side of the input field.

Filter favorites

The filter **favorites** feature is provided. You can select the ☆ **Favorites** icon on the right of the input window to **add**, **delete**, and select desired filter search conditions in the favorites. **Favorites** can store up to 50 cases.



Exceptional cases for application of filters

- In the tag with a numeric index (`search key` that ends with `.n`), only numbers can be input for the `search value` .
- Duplicate 'search key' and 'search value' cannot be entered.
- Searching is not possible if there is a tag whose `search key` or `search value` is missing. Invalid tags appear in gray.

Applying non-parsing keyword filters

You can search the logs containing keywords that have not been parsed, or for which no index has been created. In this case, all logs in the specified range are fully scanned. Therefore, the search speed may be slower compared to the indexed key. For standardized log data, it is recommended to search using the index key through **Log parser setting**.

Filter ⓘ

🔍 oid != -1009464250 ✕ okind == -398596773 ✕ content == *select* ✕ Enter Filter and press enter to add.

oid != -1009464250 && okind == -398596773 && content == *select*


1. Select **Category**. It is required to set the category.
2. In the **Filter** field, enter a desired keyword for a search after leaving spaces on the **content** basis.
e.g. `content *select*`
3. Click **Search** to view the logs. Several logs are viewed first. You can view a maximum of 10,000 logs per each search.
4. You can make additional inquiries by scrolling down and selecting **View more** at the bottom.

The screenshot shows the Log Search interface. At the top, there's a filter bar with the same filters as the previous image. Below it, a message states: "You have viewed 10,500 of the total 36,841 results, and the results retrieved are 9999+." The main area displays a list of logs with columns for Time, Category, and Content. A modal dialog is open, asking "Want to view more logs?" and providing instructions: "Adding a log parser enables search without additional lookups." It includes buttons for "Setting up a log parser" and "View more".

- Among all logs, 10,000 logs are viewed for a server search range. The default range for a server search is 200,000, but the ratio may differ depending on the total log volume.
- For more information about the parser configuration, see [the following](#).

Modifying the filter

Enter a value in the filter and then click the specified value to modify it.



- You can modify the filter by entering text again in the input field.
- You can modify the recommended values by selecting them in the **Guide UI** below the input field.

Search Key

The search key Indicates an identifier for accessing a desired value within the log data. The actual data for the search key is the search value. The tags on the left ② area are the search keys parsed for each category. You can enter filters by selecting tags. The orange tag is the category, and the blue tag is the search key.

For example, AppLog and AppStdOut in ② are categories. The tags below such as oid are search keys parsed. In **Log > Log Configuration**, you can register and set the parsing logic in **Log parser setting** tab with the search key.

Filter input syntax

A tag consists of a search key and a search value. In the following example, the search key is exception and the search value is UnknownHostException . The example searches for the log data that includes an exception (UnknownHostException) that is generated when the server cannot connect the host because the IP address and domain address do not match among the collected log data.



exception

==

UnknownHostException

×

Search key types

Search key types	Search key format	Meaning	Example for Search key and search value	Search example
String keyword	keyword	File Name	- Key: fileName - Value: /data/whatap/logs/yard.log	fileName:/data/whatap/logs/yard.log
Numeric keyword	keyword.n	Response Time	- Key: response_time.n - Value: 2945	response_time.n>=2945
Reserved keyword (predefined keyword)	@keyword	Transaction ID	- Key: @txid - Value: 85459614215434144	-

Common syntaxes

Syntax type	Description	Example
<code>==searchValue</code>	Searches for the logs that match the search value.	<code>exception==RuntimeExceptionexception</code>
<code>!=searchValue</code>	Searches for the logs that exclude the search value.	<code>exception!=RuntimeException</code>
<code>*searchValue</code>	Searches for the logs ending with the search value.	<code>word==*hello</code>
<code>searchValue*</code>	Searches for the logs starting with the search value.	<code>word==hello*</code>
<code>*searchValue*</code>	Searches for the logs that contain the search value	<code>word==*hello*</code>

Syntax type	Description	Example
	in the middle.	
*search*Value*	Searches for the logs that contain the search value.	word==*he*llo*
re:{regex}	Searches for the logs that match a regular expression.	caller==re:^(i\.w\.a\.w\.s\.v\.r\.
**	Searches all logs that match the search key.	

Syntax in case the search key has a numeric value (keyword.n)

The following syntax is supported only when the search key is in `keyword.n` format.

- Search values must only be numbers.
- No prefix is attached to the value of the `.n` keyword. All keywords other than `.n` must be attached with a prefix.

For example, `+>searchValue` is invalid.

Syntax type	Description	Example
>searchValue	Searches for the logs containing values greater than the search value.	response_time.n>3000
>=searchValue	Searches for the logs containing values greater than or equal to the search value.	response_time.n>=3000
==searchValue	Searches for the logs containing the same search value.	response_time.n==3000
!=searchValue	Searches for the logs that contain values different from the search value.	response_time.n!=3000
<searchValue	Searches for the logs containing values less than the search value.	response_time.n<3000

Syntax type	Description	Example
<code><=searchValue</code>	Searches for the logs containing values less than or equal to the search value.	<code>response_time.n<=3000</code>


Viewing personal information

If you have the **log personal information query** role, you can display de-identified personal information by enabling the **Display personal information** ☒ toggle button.

- For more information about the **log personal information query** role, see [the following](#).
- For more information about **De-identification of personal information**, see [the following](#).

Log tagging options

When you select a log tag, the following drop-down menu appears. You can find the **Search**, **Exclusion search**, **Adjacent Log** and **Transaction information** options.



- Copy
- **Search**

If you select the **Search** option, the tags that correspond to **Filter** are entered with the inclusion (`==`) condition.

- **Exclusion search**

If you select the **Exclusion search** option, the tags that correspond to **Filter** are entered with the exclusion (`!=`) condition.

- **Adjacent Log**

If you select the **Adjacent Log** option, the **Adjacent Log** detail window appears. Search for the logs of the target server and the logs in time zones adjacent to the selected log. You can click **Time selector** to see the logs in adjacent time zones. Reference logs appear with a blue background.

- **Transaction**

If the log contains the `@txid` and `@mtid` tags, you can see the information in the transaction Information details window by clicking the transaction information option.

Content highlighting

The highlight function allows you to easily identify desired keywords in log content. You can filter with single or multiple keywords.

1. Enter the keyword you want to highlight in the keyword input field.
 - Example: `select`
2. The searched keyword will be highlighted in the Content of the log list.
 - If you click the  **Log full screen** icon, you can view **Log** and **Timestamp** in fullscreen.
 - You can navigate between highlighted keywords using shortcuts.
 - `Enter` : Move to the next row
 - `Shift+Enter` : Move to the previous row
 - `^ / v` icons: Move between rows
 - `Cmd + F` (Mac) / `Ctrl + F` (Windows): Search keywords

Multi-keyword condition

When highlighting with multiple keywords, write as follows:

String to input	Description	Result
<code>a b c</code>	Separate each keyword with a space.	a, b, c

String to input	Description	Result
"Whatap is good."	To include spaces in keywords, wrap them in ' ' or " ".	Whatap is good.
"Whatap\\ is good."	If \ is included in the keyword surrounded by "", enter it as \\.	Whatap\ is good.

Setting the highlight color

By selecting  **Highlight**, you can set the keyword and its color for highlighting.

- In addition, enter the keyword to set the color for.
- If you click **Color** on the left of the input field, available colors appear.
- Basically, different highlighting by the level (**WARN**, **ERROR**, **FATAL**) is applied.
- The settings are saved in **Project Rules**.

Setting the table

Column settings

Right of the  area, select  **Column settings**. You can add columns or set order.

- **Adding columns**

You can add columns to the table by selecting tags. If you unselect the **log** column, you cannot see **Log view detail setting**. Be sure to select at least one column.

- **Setting the column order**

When you add a column, the column is added to **Display columns**. Drag a desired column to change the order of the columns.

Log view detail setting

On the right of , select  **Log view detail setting**. **content** and **Tag** are all checked and two items are also displayed. Be sure to select any of **content** and **Tag**.

Unchecked items are not displayed in the table. If you deselect the **Tag** as follows, the log's **Tag** does not appear in the table.

Log table tag invisible

When you add tags in the **Tag management** list, the tags in the log are listed in the order you added. You can change the order of the tags by dragging them. If you disable an added tag, the disabled tag does not appear.

- **Column settings** and **Log view detail setting** can be used in **Live Tail**, **Log Search**, and **Log Trend**.
- In the same project, the **Live Tail**, **Log Search**, and **Log Trend** menus share **Column settings** and **Log view detail setting**.

Log Setting

Home > Select Project > Log > Log Configuration

You can configure the log monitoring related settings in **Log Configuration**. Using the tab at the top, you can use menus for checking the agent settings, determining whether to enable log monitoring, setting the log data retention period and lookup password, registering the log parser, and setting the quick indexes.



- To use the **Activate log monitoring** function, the **Edit project** role is required.
- The **Log edit** role allows you to modify the **Log Configuration** menu other than **Activate log monitoring**.

Starting the log monitoring

Getting started with log monitoringLog filter settingsLog primary parser settingLog secondary parser settingFast index settingLog long-term archive statisticsDe-identification of personal information

Set up the agent and enable log monitoring ⓘ

View plans

You can easily check the log of the application with the Live Tail and Log Trend.

▼ 1. Check agent settings

To collect logs from Java applications, please check the following:

1. Checking and Upgrading Agent Versions (available from v2.1.1 or later)

2. Agent Settings > Direct Input > Enter the agent setting command below.

3. Restart

Agent settings

logsink_enabled=trueCopy

▶ 2. Activate log monitoring

Log monitoring data setting ⓘ

InitializedSave

Log lookup password

Turn on to use the log lookup password.

Log retention period (initial value) ⓘ

15days

You can enter a value between 1 days and 40 days.

Log usage

Category	Data retention	Cumulative number of logs (based on the retention period) ⓘ	Average daily log count ⓘ	Yesterday's log count ⓘ	Expected log count per week ⓘ	Expected log count per month ⓘ
Total		3,523,236,480	234,882,433	124,084,216	1,644,177,010	7,046,472,900
AppLog	15days	3,494,072,628	232,938,175	123,070,829	1,630,567,225	6,988,145,250
AppStdErr	15days	338,233	22,548	12,222	157,836	676,440
AppStdOut	15days	13,911,268	927,417	483,168	6,491,919	27,822,510
VirtualLog	15days	14,914,351	994,290	517,997	6,960,030	29,828,700

At the top, select the **Getting started with log monitoring** tab. If you select the ⓘ **View guide** icon and the **View plans** button, the corresponding guide screen appears.

Set up the agent and enable log monitoring

In the **1** area, check the agent settings, and then click the **Activate log monitoring** toggle button to set whether or not to enable the log monitoring.

Check agent configuration

Check the agent version and settings to start the log monitoring. Select **Check Agent Configuration** to follow the onscreen instructions.

- **Server Monitoring**

Under **Apply**, check the application guides for each application. See the following guides: [Java](#), [PHP](#), [Node.js](#), [Python](#), [Go](#).

- **DB Monitoring**

Under **Apply**, check each individual application guide. See [the following](#). According to the instructions, add the settings for the log monitoring target file in [whatap.conf](#).

- **Kubernetes Monitoring**

Under **Apply**, check each individual application guide. See [the following](#).

Activate log monitoring

Select **Activate log monitoring** to set whether or not to enable the log monitoring.

- If you turn on the  toggle button, the log monitoring is enabled. You can try it for free for 15 days from the activation date.
- If you turn off the  toggle button, the log monitoring is disabled. Logs are no longer saved.

Log monitoring data setting

In the **2** area, you can see the **Log usage**. You can also change the **Log retention period** and **Log lookup password** settings.

Log retention period

This is the default data retention period to be applied commonly. If not specified, the default value is 1 day. In addition to the option to select a log retention period, you can enter a desired period of time. If not set separately in the **Log usage** list, this log data retention period is applied by default. If you set the data retention periods for each category in the **Log usage** list and select **Initialized**, the default data retention period is reset.

Log usage amount

You can specify the log data retention periods for each category in the **Log usage** list. **Log Count** means the number of logs stacked for the period. For example, **Daily log count** is the number of logs accumulated during the day, and **Expected log count** is the number of logs of the today's log count multiplied by the days of data retention.

You can specify the log data retention period as follows. You can free the space by deleting old data according to the specified period.

- **Trial Project**

You can select 1, 2, or 3 days for the data retention period.

- **Paid Project**

For the data retention period, you can select one of 1 day, 2 days, 3 days, 4 days, 5 days, 6 days, 7 days, 10 days, 30 days, and 40 days.

- **Data size-based billing**

The price differs depending on the data retention period.

For example, if an average of 2 million logs per day accumulates and the data retention period is set to 3 days, an average of 6 million logs is kept on the collection server and subject to billing.

Log lookup password

To enhance security, set the **Log lookup password**. The log lookup password is optional. If you are using a log lookup password, enter the password to go to the log screen.

ⓘ In case you forgot the password



If you have the **Edit Log** role, you can modify it with a new password in **Log Configuration**.

Log Filter Settings

Log Filter Settings allow you to store only the logs that match the filter conditions specified by the user. By storing only the necessary logs, you can reduce costs.

! Permission

Only users with log edit permissions can configure these settings.

Add a Log Filter Setting

This section explains how to add a new filter setting. When configuring log filter settings, **Filter Name**, **Show in Live Tail**, **Category**, and **Filter Condition** are required fields.

Even logs that are not stored by the filter can still be viewed in **Live Tail** if the **Show in Live Tail** toggle is enabled.

! Only one filter can be set per category.

1. Go to **Logs > Log Settings > Log Filter Settings** tab.
2. Click the + **Add** button to create a new filter setting.
3. In **Add Log Filter Setting**, enter **Filter Name**, **Show in Live Tail**, **Category**, and **Filter Condition**.
 - **Filter Name:** A name to identify the filter.
 - **Show in Live Tail:** Select whether to display in Live Tail regardless of storage.
 - **Category:** The log category to which the filter is applied (**multiple selection allowed**).
 - **Filter Condition:** Conditions for logs to be stored by the filter.

! For more details on filter conditions, see [WhaTap log search query syntax](#).

- **Example**

- **Include:** `oid : 123`
 - Collect only logs whose oid is 123.
 - **Exclude:** `not oid : 123`
 - Collect only logs whose oid is **not** 123.
 - **Include all:** leave the filter condition empty
 - **Exclude all:** `not content : *`
4. Click **Add** to review and edit the filter settings in the **Preview Results** screen.
 5. Once confirmed, click **Apply** to save the log filter settings.

Simulation

You can preview the log results based on the entered category and filter conditions before applying the log filter settings.

Click the **Simulation** button to display the **Log Search** screen with the applied filter settings.

Edit a Log Filter Setting

Registered log filters can be modified by clicking the edit icon. The **Edit Log Filter Setting** screen shows the existing values, which you can change and apply as needed.

 When editing log filter settings, only one category can be selected.

Delete a Log Filter Setting

Registered filters can be deleted by clicking the delete icon. After deletion, make sure to click the **Save** button.

 If changes are not saved, the filter will not be deleted.

Import/Export Log Filter Settings as JSON

Log filter settings can be imported or exported in JSON format.

This is useful if you want to apply the same log filter settings to another project.

Export JSON

1. Click the  button at the top right of the screen.
2. In the **Export JSON** window, click the **Export** button.
3. The log filter settings will be downloaded as a JSON file.

Import JSON

1. Click the  button at the top right of the screen.
2. Select the JSON file to import.
3. In the **Import JSON** window, click either **Add to List** or **Overwrite** to import the log filter settings.
4. Click the **Save** button to apply the changes.

Validation Rules

The following validation rules are applied when saving settings:

- **type:** Only `GENERAL` is supported.
- **category:** Required, supports `string` type.
- **condition**
 - **language:** Only `jsonDsl` and `lucene` are supported.
 - **description:** Invalid filter conditions cannot be saved.
- **showLiveTail:** Only `boolean` type is supported.
- **name:** Required, supports `string` type.

Parsing The Logs

The log parser helps you transform the irregular logs into a structured form. WhaTap Log Monitoring provides two types of parsers.

- **GROK parser:** The logs collected in arbitrary forms are parsed using regular expressions and GROK syntax.
- **JSON parser:** The logs collected in the JSON form are parsed.

! Common precautions

- If multiple parsers are registered in the same category, only the first matching parser is applied.
- WhaTap has the role to disable the parsers that may affect the stability of the WhaTap service.

GROK parser

If the logs are collected in an irregular form, you can use the GROK parser to parse them. GROK syntax provides named regular expressions, allowing you to use regular expressions more easily.

Starting GROK

GROK provides two types of syntaxes.

1. `%{SYNTAX:SEMANTIC}` : Syntax provided in the GROK library. Using the **named regular expressions**, you can extract tags. For usage examples, see [the following](#).
 - **SYNTAX:** Specify the named regular expressions provided by GROK.
 - **SEMANTIC:** Specify the name for the matched value.

! named regular expressions

The syntax is provided by GROK. This function provided by GROK allows you to assign names to complex regular expressions.



name	regular expression
WORD	<code>\b\w+\b</code>
SPACE	<code>\s*</code>
NOTSPACE	<code>\S+</code>
UUID	<code>[A-Fa-f0-9]{8}-(?:[A-Fa-f0-9]{4}-){3}[A-Fa-f0-9]{12}</code>

To see all named regular expressions provided by WhaTap, see the following [link](#).

2. `(?<SEMANTIC>REGX)` : The **named capturing group** syntax for regular expressions. You can use regular expressions to extract tags according to the user's intent. For usage examples, see [the following](#).

- **SEMANTIC**: Specify the name for matched values.
- **REGX**: Enter regular expressions for matching.



named capturing group

It is the syntax provided by regular expressions.

- capturing group: This indicates the function to bundle multiple tokens into one as a single matching unit.
- named capturing group: Name assigned to the capturing group.
- Let's look at the string matching example. [dev@whatap.io](#)
 - Example 1 `(\w+)@(\w+\.\w+)`
 - Example 2: When the entire email is matched and the username and domain are matched additionally, `(?<username>\w+)@(?<domain>\w+\.\w+)`

%{SYNTAX:SEMANTIC} Usage example

The following is an example that uses the `%{SYNTAX:SEMANTIC}` syntax.

Sample log

```
[2023-08-08 02:02:30,101 GMT][INFO ][i.w.y.l.c.LogSinkDexScheduleThread.realProcess(159)] 8 VirtualLog 20230808
02:01:00.000 {area=4, city=5} 56ms
```

You can look at the sample log to infer what each word means. When replacing each part with semantic words, it can be expressed as follows:

semantic replace

```
[date][logLevel][caller] projectCode logCategory dexBuildStartTime {area=areaEnum, city=cityEnum} dexBuildElapsed
```

Any semantic word can be replaced with a regular expression. The GROK parser allows you to use the predefined named regular expressions. `TIMESTAMP_ISO8601`, `LOGLEVEL`, and `DATA` used here are named regular expressions provided by GROK. These values are matched after replacing each of them with the following regular expression.

- name: `TIMESTAMP_ISO8601`
 - regular expression: `%{YEAR}-%{MONTHNUM}-%{MONTHDAY}[T]
%{HOUR}:%{MINUTE}(?:%{SECOND})?%{ISO8601_TIMEZONE}?`
- name: `LOGLEVEL`
 - regular expression: `LOGLEVEL
([Aa]lert | ALERT | [Tt]race | TRACE | [Dd]ebug | DEBUG | [Nn]otice | NOTICE | [Ii]nfo | INFO | [Ww]arn(?:ing)? | WARN(?:ING)?`
- name: `DATA`
 - regular expression: `.*`

GROK parsing pattern

```
\[%{TIMESTAMP_ISO8601:date}\sGMT\]\[%{LOGLEVEL:level}\s\]\[%{DATA:caller}\]
```

If you parse with the above syntax, you can extract tags as follows. In this way, GROK's `%{SYNTAX:SEMANTIC}` syntax helps you apply complex and long regular expressions easily and concisely.

Tag extraction

```
- date : 2023-08-08 02:02:30,101
- caller : i.w.y.l.c.LogSinkDexScheduleThread.realProcess(159)
- level : LEVEL
```

(?<SEMANTIC>REGX) Usage example

The parts that do not match with the named regular expressions can be parsed using the `(?<SEMANTIC>REGX)` pattern. In the [sample log](#) above, the parts that cannot be parsed with the `%{SYNTAX:SEMANTIC}` syntax are as follows.

Unparsed area

```
8 VirtualLog 20230808 02:01:00.000 {area=4, city=5} 56ms
```

When replacing each part of the log with semantic words, it can be expressed as follows.

semantic replace

```
projectCode logCategory dexBuildStartTime {area=areaEnum, city=cityEnum} dexBuildElapsed
```

These irregularly strings can be parsed using the `(?<SEMANTIC>REGX)` syntax.

Regular expressions that match for each sample log parsing keyword

Parsing keyword	(?<SEMANTIC>REGX)
8	(?<projectCode>\d)
VirtualLog	(?<logCategory>\w*)
20230808 02:01:00.000	(?<dexBuildStartTime>\d{8}\s\d{2}:\d{2}:\d{2}\.\d{3})
area=4	area=(?<areaEnum>\d)
city=5	city=(?<cityEnum>\d)

Parsing keyword	(?<SEMANTIC>REGX)
56ms	(?<dexBuildElapsed>\d{2})ms

Basic regular expression syntax

Syntax rule	Meaning	Alias
<code>?</code>	0 or 1	-
<code>+</code>	1 or more	-
<code>*</code>	0 or more	-
<code>a{5}</code>	exactly 5	-
<code>\w</code>	word character	<code>[a-zA-Z_0-9]</code>
<code>\s</code>	white space	-
<code>.</code>	any character except newline	
<code>[abc]</code>	any of	-
<code>[^abc]</code>	not a,b, or c	-
<code>[a-z]</code>	character between a and z	-
<code>[1-3[7-9]]</code>	union (combining two or more character classes)	-
<code>[1-6&&[3-9]]</code>	intersection	-
<code>[0-9&&[^2468]]</code>	subtraction	-
<code>a{2,}</code>	2 or more	-
<code>a{1,3}</code>	between 1 and 3	-



Syntax rule	Meaning	Alias
<code>a+?</code>	match as few as possible	-
<code>{2,3}?</code>	match as few as possible	-
<code>(abc)</code>	capturing group (processes multiple strings as a single unit)	-
<code>\d</code>	digit	<code>[0-9]</code>
<code>\D</code>	non-digit	<code>[^0-9]</code>
<code>\W</code>	non-word character	-
<code>\S</code>	non-white space	-

By connecting the parsed keywords with space (`\s`) and special character escapes (`\{` , `\,` , `\}`), you can apply the pattern as follows.

GROK parsing pattern

```
(?<projectCode>\d)\s(?<logCategory>\w*)\s(?<dexBuildStartTime>\d{8})\s\d{2}:\d{2}:\d{2}\.\d{3})\s\{area=(?<areaEnum>\d),\scity=(?<cityEnum>\d)\}\s(?<dexBuildElapsed>\d{2})ms
```

If you parse with the above syntax, you can extract tags as follows.

Tag extraction

```
- projectCode : 8
- logCategory : VirtualLog
- dexBuildStartTime : 20230808 02:01:00.000
- areaEnum : 4
- cityEnum : 5
- dexBuildElapsed : 56
```

Applying GROK

Log Configuration > Log primary parser setting

1. To apply the GROK pattern parser, go to the **Log primary parser setting** tab in **Log Configuration**.
2. Select **+ Add** and then select the **GROK** parser in the **Parser** field.
3. When you select **Register pattern**, the pattern registration and simulation window appears on the right.
4. Enter **Pattern** and **Log**, and then click **Simulation** to check whether parsing is successful with the pattern to apply.

Pattern e.g.

```
[%{TIMESTAMP_ISO8601:date}\sGMT][%{LOGLEVEL:level}\s][%{DATA:caller}]\s(?:<projectCode>\d)\s(?:<logCategory>\d{2}:\d{2}:\d{2}\.\d{3})\s\{area=(?:<areaEnum>\d),\scity=(?:<cityEnum>\d)\}\s(?:<dexBuildElapsed>\d{2})ms
```

Log example:

```
[2023-08-08 02:02:30,101 GMT][INFO ][i.w.y.l.c.LogSinkDexScheduleThread.realProcess(159)] 8  
VirtualLog 20230808 02:01:00.000 {area=4, city=5} 56ms
```

5. If the simulation is successful, you can view **Simulation result** and **Performance measurement results**.
6. When you click **Apply pattern** after simulation, the pattern that has been entered for the selected parser is applied.
7. After applying the pattern, enter values for **Category**, **Log detection condition**, and **Pattern**.

Parser Add

Parser*

WhaTap reserves the right to disable parsers that may affect the stability of the WhaTap service.

GROK

parser pattern*

The pattern is registered after the simulation is performed to ensure that the log is successfully parsed with the entered pattern.

Success

```

\[%{TIMESTAMP_ISO8601:date}\sGMT\}\[%{LOGLEVEL:level}\s\}\[%{DATA:caller}\}\s(?:<projectCode>\d)\s(?:
<logCategory>\w*)\s(?:<dexBuildStartTime>\d{8}\s\d{2}:\d{2}:\d{2}\.\d{3})\s\{area=(?:<areaEnum>\d),\scity=(?:
<cityEnum>\d)\}\s(?:<dexBuildElapsed>\d{2})ms

```

Modify parser pattern

Category*

If multiple parsers are registered in the same category, only the first matching parser is applied.

AppLog

Log detection condition

If you don't enter any filter, or if you enter only one of the search keys/values, the pattern is applied to all logs.

oname

demo-8100

Add

- **Category**

Select a log category. **Category** is required.

- **Log detection condition**

- Only the logs that meet the conditions are applied to the parser.
- Select or enter values for **Search key** and **Search value**.
- **Log detection condition** are applied before all parsers are run. Accordingly, you cannot use the **Tag** that have been added for the parser.

8. Select **Add** to register a parser.

- In the log parser list, you can change the **Priority** or **Enable**, **Edit**, or **Delete** the parser.
- After **parser simulation**, you can register a pattern.

! GROK parser precautions

- GROK parser supports two patterns: `%{SYNTAX:SEMANTIC}` and `%{SYNTAX:SEMANTIC}` .
- When using the `%{SYNTAX:SEMANTIC}` pattern, `SEMANTIC` must be entered.
- When using the `%{SYNTAX:SEMANTIC}` pattern, `SEMANTIC` must be unique in a parser.
- When using the `(?<SEMANTIC>REGX)` pattern, `SEMANTIC` can only contain characters (a-z, A-Z), numbers (0-9), and specified special characters (. , _ , -).
- `SEMANTIC` must start with a character (a-z, A-Z).
- `SEMANTIC` must end with a character (a-z, A-Z) or number (0-9).

Selecting additional GROK patterns

During GROK parser simulation, in addition to the basic patterns (`WORD` , `NUMBER` , `IP` , etc.), you can select **additional patterns** that match a specific log format. 32 pattern types, such as Java, Nginx, and MySQL, are provided grouped into 11 categories.

Additional pattern category dropdown

1. In the parser simulation window, click the dropdown in the **Additional patterns (optional)** area.
2. From the list of pattern types grouped by category, select the pattern type you need. You can also search by entering a pattern type name or description in the search box.
3. The selected pattern type is loaded together with the basic patterns during simulation.

! The basic patterns (Core) are always included automatically. Selecting additional patterns is optional.

Pattern type list by category

Category	Pattern type
Web Servers	haproxy, httpd, iis, nginx, squid
Application Servers & Frameworks	java, log4j, rails, ruby, springboot

Category	Pattern type
Databases	mongodb, mysql, oracle, postgresql, redis
Cloud	aws
System & Infrastructure	docker, kubernetes, linux-syslog
Messaging & Search	elasticsearch, kafka
Network & Security	bind, bro, firewalls, junos, zeek
Mail	exim, postfix
Monitoring & Backup	bacula, mcollective, nagios
Build Tools	maven

Usage example

To parse a Tomcat Catalina 8 log, select the **java** pattern type additionally and then use the `%{CATALINA8_LOG}` pattern.

Sample log

```
31-Jul-2020 16:40:38 INFO [main] org.apache.catalina.startup.VersionLoggerListener.log Server version name: Apache Tomcat/8.5.57
```

GROK parsing pattern

```
%{CATALINA8_LOG}
```

Tag extraction

```
- log_timestamp : 31-Jul-2020 16:40:38
- log.level : INFO
- java.log.origin.thread.name : main
- java.log.origin.class.name : org.apache.catalina.startup.VersionLoggerListener
```

```
- log.origin.function : log
- log_message : Server version name: Apache Tomcat/8.5.57
```

Viewing pattern details

You can check which patterns are included in the selected additional pattern type in the details window. You can look up the pattern definitions (regular expressions) and sample logs and apply them directly to the simulation.

Viewing pattern details

After selecting an additional pattern type, click the selected pattern type tag to open the details window on the right. The details window is displayed side by side without hiding the simulation form, so you can proceed with the simulation while checking the pattern definitions.

You can use the following functions in each pattern item.

Button	Function	Description
Sample	Enter a sample log	Automatically enters an example log that matches the pattern into the Simulation parsing log input box.
Select	Insert a pattern	Inserts it into the Pattern input box in the <code>%{pattern_name}</code> format. The corresponding pattern type is also selected automatically.
Copy	Copy to clipboard	Copies the <code>%{pattern_name}</code> format to the clipboard.

- When you click a pattern item, its pattern definition (regular expression) expands. You can expand multiple patterns at the same time to compare their definitions.
- In the `%{PATTERN:field}` references included in a pattern definition, the pattern name, field name, and literal characters are displayed distinguished by different colors.
- You can run the **simulation** directly with the log entered by the **Sample** button to check the parsing result.

Saving simulation sample logs

The log used in the simulation is saved automatically together with the parser settings. When you edit a parser, the sample log used previously is restored automatically to the simulation window, so you do not need to enter the log again each time.

JSON parser

If logs are collected in JSON format, they can be easily parsed using the JSON parser.

Applying the JSON

Log Configuration > Log primary parser setting

1. To apply the JSON pattern parser, go to the **Log primary parser setting** tab in **Log Configuration**.
2. Select **+** **Add** and then select the **JSON** parser in the **Parser** field.
3. When you select **Register pattern**, the pattern registration and simulation window appears on the right.
4. Enter **Pattern** and **Log**, and then click **Simulation** to check whether parsing is successful with the pattern to apply.

Example

```
2023-08-08 02:43:28,615 -- {"host":"10.21.3.24","method":"POST","status":"200","url":"http://dev.whatap.io/yard/
api/flush"} --
```

In the example log, if **Prefix** and **Postfix** are specified as `--` and `url` is specified in **Ignore**, only three **Tag** (`host` , `method` , `status`) are created.

- **Prefix** Specify the position where the JSON format begins in the log. If the entire log is in JSON format, set an empty value.
- **Postfix** Specify the position where the JSON format ends in the log. If the entire log is in JSON format, set an empty value.
- **Ignore** Specify the key for which tags are not created during JSON formatting.

5. If the simulation is successful, you can view **Simulation result** and **Performance measurement results**.

Parser Simulation

View Guide

Pattern

agenttime, category, content, logContent, message, pcode, time, timestamp is a reserved word.

Reserved words do not work even if registered as log detection conditions or patterns.

Prefix

Postfix

Ignore

Log *

Log simulation is essential and additional pattern registration may be required depending on simulation results.

```
2023-08-08 02:43:28.615 --
{"host":"10.21.3.24","method":"POST","status":"200","url":"http://devote.whatap.io/yard/api/flush"} -

```

[Simulation](#)

Simulation result Success

Key	Value	Result
method	POST	Ok
host	10.21.3.24	Ok
status	200	Ok

Performance measurement results Success

Simulation count	Time min(ns)	Time max(ns)	Time a
1	7,047	7,047	7,047
10	2,663	4,255	2,872
100	2,655	4,574	2,818
1,000	2,651	94,191	2,961
10,000	2,651	146,472	2,987
100,000	2,648	194,424	2,917
1,000,000	2,647	80,465,427	3,136

[Apply pattern](#)

6. When you click **Apply pattern** after simulation, the pattern that has been entered for the selected parser is applied.
7. After applying the pattern, enter values for **Category**, **Log detection condition**, and **Pattern**.

Parser Add

Parser*

WhaTap reserves the right to disable parsers that may affect the stability of the WhaTap service.

JSON

parser pattern*

Simulate that the log is parsed successfully. Simulation must be performed, and additional pattern registration may be required depending on the simulation results.

✓ Success

Prefix --

Postfix --

Ignore url

Simulation verification and pattern modification

Category*

If multiple parsers are registered in the same category, only the first matching parser is applied.

AppLog

Log detection condition

If you don't enter any filter, or if you enter only one of the search keys/values, the pattern is applied to all logs.

oname

demo-8100

Add

- **Category**

Select a log category. **Category** is required.

- **Log detection condition**

- Only the logs that meet the conditions are applied to the parser.
- Select or enter values for **Search key** and **Search value**.
- **Log detection condition** are applied before all parsers are run. Accordingly, you cannot use the **Tag** that have been added for the parser.

8. Select **Add** to register a parser.

- In the log parser list, you can change the **Priority** or **Enable**, **Edit**, or **Delete** the parser.
- Before registering a parser, you can check whether the pattern to register is normal through



Simulation. It is the same as the GROK parser registration simulation process. See [the following](#).

Example of using JSON format

Sample log

```
{"host":"10.21.3.24","method":"POST","status":"200","url":"http://dev.whatap.io/yard/api/flush"}
```

If the above sample log has been collected, select the **JSON** parser in the **Parser Add** window. Without having to write a complex parsing logic, you can extract the **Tag** for log analysis as follows:

Tag extraction

```
- host : 10.21.3.24
- method : POST
- status : 200
- url : http://dev.whatap.io/yard/api/flush
```

Usage example when configuring part of the JSON format

Some JSON format sample log

```
2023-08-08 02:43:28,615 -- {"host":"10.21.3.24","method":"POST","status":"200","url":"http://dev.whatap.io/yard/api/flush"} --
```

If part of the log is in JSON format as in the example, specify **Prefix** and **Postfix**. WhaTap log monitoring recognizes and parses the area between **Prefix** and **Postfix** in JSON format.

Tag extraction

```
- host : 10.21.3.24
- method : POST
- status : 200
- url : http://dev.whatap.io/yard/api/flush
```

Log Fast Index Setting

Fast Index Setting

Select the **Fast index setting** tab at the top of **Log Configuration**. Collecting a large number of logs can significantly decrease the log search performance. The frequently used search conditions are created as **index**, you can improve the log search performance for quick search. The setting items are as follows:

Set value	Required	Description
Category	Mandatory	Category to be set as fast index
Search Key	Mandatory	Search key for fast index setting
Case insensitive	Option	Whether to be case sensitive
Rule	Mandatory	* must be included at least one.
Enabled	Mandatory	Active or inactive (default value is <code>true</code>)

Log Long-term Archive Statistics Setting

Log long-term archive statistics

Select the **Log long-term archive statistics** tab at the top of **Log Configuration**. Log data is so large and difficult to retain for a long time. Using the Set log statistics data function, you can **save information on how many logs that meet specific conditions are collected every 5 minutes**. Even if actual log data has been deleted for a long time, you can check the trend of how many logs that meet the conditions are collected.

Adding the log long-term archive statistics

× Log long-term archive statistics Add

Statistic key*

Please input the Statistic key.

Data retention period(disk usage)

15 Days (About 3 MB) ▾

Category*

Please select the Category.

Log detection condition*

Please select a category first. ▾

Please select a category first.

Exclude ☐

Case Sensitive ☒

+ Add

Add

If you select + **Add** under the **Log long-term archive statistics** tab, the **Log long-term archive statistics Add** window appears. You can add rules by using + **Add** or delete the created rules by using the - icon.

Setting item

Field	Description
Category	Category to apply the rule to.
Statistic Key	The same key cannot be set twice to save when a log that meets the rule is generated.
Log detection condition	Condition for generating log statistical data. Statistical data is generated based on how many logs that meet this condition are collected.
Exclude	If checked, statistical data is generated with values that do not correspond to the entered conditions.
Case Sensitive	Specify case sensitivity for the values of the entered log detection conditions.
Enabled	Active or inactive (default value is <code>true</code>)


Example

If a setting is added as follows, statistical data is generated with a key value of **TotalCount** for the logs whose status is `200` or `300`.


Log long-term archive statistics You can generate statistical data about how many logs have been collected that meet certain conditions.					+ Add	Save
Category	Statistic key	Log detection condition	Data retention period(disk usage)	Enable		
AppLog	TotalCount	status 200,300	15 Days (About 3 MB)	☐		
AppLog	TotalCount	status 200	15 Days (About 3 MB)	☐		
AppLog	TotalCount	status 300	15 Days (About 3 MB)	☒		
AppLog	TotalCount	status 200,300	15 Days (About 3 MB)	☒		

Data Search


1. Create a widget by searching **Log long-term archive statistics** in **Widget templates** of **Integrated Flex Board**.



2. Enter the **Category** and **Statistic key** to view and then select **Apply**.



3. With the added settings, you can check the **Log long-term archive statistics** data as follows.



Log Personal Information De-identification

Personal information de-identification

In **De-identification of personal information**, you can mask personal information in log data or replace it with safe values.

- Encryption: Specified sensitive data is encrypted and stored in the database.
- Masking: When displayed on the log screen, sensitive data is masked (*******) by default.
- Substitution: Sensitive data is replaced with a user-specified value that is safe.

Personal information classification

The items classified as personal information are as follows: If there are any items that need to be treated as sensitive data in addition to personal information, mask or replace the desired items in the log data by using the **De-identification of personal information** feature.

- National ID number
- Passport number
- Driver's license number
- Foreigner registration number
- Credit card number
- Bank account number
- Biometric information

Required roles

De-identified personal information is displayed as masked. Only users with the **log personal information query** role can view original data through the **Display personal information** toggle. Adding or modifying de-identification target search keys requires the **log editing** role.

! For more information about the **log personal information query** and **log editing** roles, see [the following](#).

Setting the personal information de-identification

You can protect sensitive information by applying masking to new log data.

1. In the **De-identification of personal information** tab, select **Add**.
2. In the **De-identification of personal information Add** pane on the right, select **Category** and enter the **Search key** to mask.
3. Click **Add** to add the setting.
4. After checking the de-identified items in the **De-identification of personal information** list, click **Save** on the upper right to save the current settings.

De-identification of personal information If you add a search key containing personal information, the data is encrypted, stored, and masked.				+ Add	Save
Enable	Category	Search key		Edit	
☐	VirtualLog	age.n		✎	✕
☒	VirtualLog	status		✎	✕
☒	ApplLog	agenttime		✎	✕

- The de-identified search keys are output in the format of `search key.pii` in the log content.
- You can perform ☒ **Enable** for de-identified items in the **Enable** column of the **De-identification of personal information** list.
- You can [✎ modify](#) and [✕ remove](#) in the **edit** column of the **De-identification of personal information** list.

✓ If the search key name is specified as `search-key.p` in the **Log primary parser setting** tab, it is automatically registered as a de-identification target. You can check or disable auto-registered items in **Specified search key** after selecting the **Category** in the **De-identification of personal information** tab.

We recommend configuring de-identification directly in the **De-identification of personal information** tab rather than using the `.p` suffix in parser settings.

Personal information display

De-identified data is displayed as masked in the **Live Tail**, **Log Trend**, and **Log Search** menus. Users with the **log personal information query** role can click the **Display personal information** ☒ button in the log list to view the original data.

▼	java-demo2	2024-12-16 16:06:49.988	area	river	!privacy	true	!status.pii	***	pcode	8	city	seoul	onodeName	node-0	myname.p	***
			oid	-1737750367	okind	1152715164	price.n	12000	oname	java-demo2	age.n	33				
			agenttime	1734332809555	category	VirtualLog	okindName	dev-okind-0	onode	1693789385						
			You need not love me, if you choose not to.													
▶	java-demo1	2024-12-16 16:06:49.993	area	mountain	!privacy	true	!status.pii	***	pcode	8	city	junju	onodeName	node-1	myname.p	****
			Waiting does not have to aim for a meeting.													
▼	java-demo2	2024-12-16 16:06:49.988	area	river	!privacy	true	!status.pii	500	pcode	8	city	seoul	onodeName	node-0	myname.p	lee
			oid	-1737750367	okind	1152715164	price.n	12000	oname	java-demo2	age.n	33				
			agenttime	1734332809555	category	VirtualLog	okindName	dev-okind-0	onode	1693789385						
			You need not love me, if you choose not to.													
▶	java-demo1	2024-12-16 16:06:49.993	area	mountain	!privacy	true	!status.pii	200	pcode	8	city	junju	onodeName	node-1	myname.p	hong
			Waiting does not have to aim for a meeting.													
▶	java-demo6	2024-12-16 16:06:50.064	area	sea	!privacy	true	!status.pii	200	pcode	8	city	daegu	onodeName	node-1	myname.p	song
			Waiting does not have to aim for a meeting.													

- ⓘ The search key specified in the **De-identification of personal information** tab appears in the log content of the Log menu in the format of `Search key.pii`.

Example of status.pii ***

Checking the log's personal information de-identification

- ⓘ If you replace the existing data using the **Log's personal information de-identification checking** feature, the data cannot be recovered.

Log's personal information de-identification checking identifies personal information in existing log data and replaces it with the **Value to substitute** specified by the user.

Replacement by target key

If you specify a target key, only the value stored in that key is replaced, and the rest of the log content is kept as is. The target key applies to both the `field` value and the `tag` value of a log. If you do not specify a target key, personal information is identified and replaced across the entire log content, the same as before.

- ⓘ Replacement applies only to the stored log data. The original value may remain in the search index even after replacement.

You specify the target key in the **Target key** input field of the log personal information de-identification check form. On the screen, `field` and `tag` are not distinguished and are entered as a single target key.

Log's personal information de-identification checking It de-identifies the accumulated logs. ⓘ
 ✓ Check inspection history

Time < 2025/03/21 12:26 ~ 2025/03/21 13:26 1 hr >	Category VirtualLog	Category VirtualLog
Personal information entity Resident registration number	Search start time 2025/03/21 12:26	Search end time 2025/03/21 13:26
Personal information pattern (?<!\w)\d{2}(0[1-9] 1[0-2])([0-2][0-9] 3[0-1])[1-4]\d{6}(?!w)	Personal information entity Resident registration number	Personal information pattern (?<!\w)\d{2}(0[1-9] 1[0-2])([0-2][0-9] 3[0-1])[1-4]\d{6}(?!w)
Whether to apply substitution ☐	Value to substitute *****	Value to substitute *****
Value to substitute ⓘ *****	Pattern matching result 0 ✓ Check results	Inspection start time 2025/03/21 13:28
▶ Start inspection ⌛ Stop inspection	Inspection end time 2025/03/21 13:28	Whether to apply substitution No
	Inspection result SUCCESS	

- Specify the start time and end time.
- Set values for **Category**.
- Select **Personal information entity** from the examples or **Direct Input**.

ⓘ When selecting **Direct Input**, directly enter a regular expression pattern to identify personal information.

- Select whether to replace personal information, if it is identified via the **Whether to apply substitution** toggle.
- Enter **Value to substitute** to change the string identified by the regular expression pattern in the existing log.


ⓘ

- Value to substitute** must not be longer than the original value.
- When enabling the **Whether to apply substitution** toggle, be sure to enter **Value to substitute**.

- Enter the conditions and then select **Start inspection**.

ⓘ You can stop this step at any time by clicking **Stop inspection** during the checkup.

7. When applying substitution, enter **Replace** in the **Start of inspection for personal information de-identification** confirmation window, and then click the **OK** button.



Start of inspection for personal information de-identification ✕

The value that matches the selected personal information pattern is replaced.

ⓘ Data cannot be restored after replacement.

If replacement is required, see below. "Replace" Please enter directly.

Replace

Cancel OK

8. After the inspection is finished, you can see **Inspection result** on the right as in the example image.
- **Category:** It displays the categories checked.
 - **Search start time** and **Search end time:** It displays the time range for the checkup.
 - **Personal information entity:** It displays the personal items that have been checked.
 - **Personal information pattern:** It displays the regular expression pattern used.
 - **Value to substitute:** It displays the applied replacement value.
 - **Pattern matching result:** The number of personal data identified is displayed. You can check the sample logs by clicking **Check results**.
 - **Inspection result:** It displays the result where personal information has been identified.

SUCCESS (check success), **FAILED** (check failure), **STOP** (check stop), **TIMED_OUT** (check ended due to timeout)

Check inspection history

To view the results of past checks after performing de-identification check for the log's personal information, click the **Check inspection history** button at the upper right. You can trace the substitution history, inspection time, and applied personal information items. You can also download the list information in **CSV** file format by clicking the **↓ CSV** button at the upper right of the **Check inspection history** window.

Check inspection history

< 2025/03/25 08:17 ~ 2025/03/25 09:17 1 hr

>

Personal infor...

>

Q

↓ CSV

stitute	Pattern matching result	Inspection start time	Inspection end time	Whether to apply substi...	Inspection result
***	0	2025/03/25 09:16	2025/03/25 09:16	No	SUCCESS
***	0	2025/03/25 09:16	2025/03/25 09:16	No	SUCCESS
***	0	2025/03/25 09:16	2025/03/25 09:16	No	SUCCESS

Personal information pattern

Examples of regular expression patterns used to identify personal information in the existing log are as follows:

Personal information pattern	Personal information example	Regular expression pattern	Replacement value
National ID number	123456-1234567	(?<!w)d{2}(0[1-9] 1[0-2])([0-2][0-9] 3[0-1])-[1-4]d{6}(?!w)	*****_ *****
Passport number	M1234567	([MSRODTC][0-9]{7}\b)	*****
Driver's license number	01-23-456789-00	(?<!w)d{2}-d{2}-d{6}-d{2}(?!w)	**_**_ *****_**
Foreigner registration number	987654-1234567	(?<!w)d{6}-?[5-8]d{6}(?!w)	*****_ *****

Personal information pattern	Personal information example	Regular expression pattern	Replacement value
Credit card number	1234-5678-9012-3456	(?<!w)d{4}-d{4}-d{4}-d{4}(?!w)	*****_***** *****_*****
Mobile phone number	010-1234-5678	(?<!w)(+82-?1[016789]-?d{3,4}-?d{4} 01[016789]-?d{3,4}-?d{4})(?!w)	****_*****_ *****
Email address	example@domain.com	([a-zA-Z0-9][a-zA-Z0-9._%+-]*@[a-zA-Z0-9.-]+\.[a-zA-Z]{2,})	***@***.***
Bank account number	123456-01-654321	(?<!w)d{6}-d{2}-d{6}(?!w)	*****_*_*_ *****
Business registration number	111-11-01111	(?<!w)d{3}-d{2}-d{5}(?!w)	***_*_*_*****
Corporate registration number	111111-1111111	(?<!w)d{6}-?d{7}(?!w)	*****_ *****

Importing And Exporting Log Settings

Importing/exporting log settings

You can save parser settings, Fast index setting, and filter settings in JSON file format, and apply them by importing a JSON file from other projects. You can reduce the hassle of creating the settings repeatedly for each project.

1. Add parser settings, Fast index setting, and filter settings to one project.
2. Select **JSON**  on the upper right of each configuration tab.
3. On the upper right of the **Export JSON** window, select **Export**.
4. The JSON configuration file is saved on your PC.
5. Move to another project and then go to the **Log > Log setting** menu.
6. Select the configuration tab where you have exported the JSON configuration file before, and then select .
7. If the file selection window appears, select the JSON configuration file saved on your PC.
8. If the **Import JSON** window appears, check the configuration file and then select **Add to list** or **Overwrite**.
9. Select **Save** on the upper right of the screen.

 After importing the JSON configuration file, you cannot save the imported settings unless you select **Save**.

Alert Notifications

This feature delivers notifications to users based on predefined conditions when issues occur in monitored systems or applications. Notifications can be sent through various channels such as email, SMS, messengers, and mobile push notifications.

For example, you can configure alert thresholds for key performance indicators such as CPU usage, memory usage, disk I/O, network traffic, transaction count, error occurrences, and response time latency. When these thresholds are exceeded or abnormal behavior is detected, notifications are automatically sent.

With WhaTap's AI-based alerting, you do not need to define complex conditions manually. The AI learns normal application behavior patterns and automatically detects and notifies anomalies. Additionally, by using tag-count alerts, you can define conditions on values aggregated by specific tags, enabling precise detection even in complex scenarios.

Key Features

- **Event Configuration:** Configure threshold-based alert conditions and notification channels
- **Event Notification Settings:** Configure notification channels and time schedules per user
- **Event History:** View the history of alert occurrences

Permissions

To add or configure events, the **Alert Configuration** permission is required.

- **View:** To view metrics event configurations, the **View & Analyze** permission is required.
- **Add/Edit/Delete:** To modify configurations, the **Alert Configuration** permission is required.
- **Receive Notifications:** To receive notifications, the **Alert Receive** permission is required.

Reference

For more details, see [Member Permissions](#).

Event Configuration

Home screen > Select project  **Alert** > **Event Configuration**

In **Event setting**, you can configure and manage events based on anomaly conditions that may occur in monitored systems or applications. When the configured conditions are met, notifications can be delivered through your preferred channels such as email, SMS, messengers, or app push notifications. You can tailor event configurations according to your monitoring objectives.

Metrics

Metrics events allow you to monitor system and application performance indicators based on real-time data.

Real-time logs

Sends alert notifications when event conditions configured in real-time collected logs are detected.

Composite log

Configure composite log events that trigger alert notifications based on log content.

Home Screen > Select Project >  **Alert** > **Event Configuration**

Event setting provides features to configure and manage events based on anomaly signals that may occur in monitored systems or applications. When configured conditions are met, you can receive notifications through your preferred channels such as email, SMS, messengers, or app push notifications. You can design event configurations according to your monitoring objectives.

Metrics

Metrics events allow you to monitor system and application performance indicators based on real-time data.

Real-time logs

Sends alert notifications when event conditions configured in real-time collected logs are detected.

Composite log

Configure composite log events that trigger alert notifications based on log content.

Reference

- When configuring events, use appropriate threshold values to prevent unnecessary event generation.
- When setting up notifications, consider enabling recovery notifications (alerts sent when issues are resolved) to support effective operational response.

Common Features

All event configurations can be downloaded in **JSON or Excel format**. By sharing the same type of product configuration, you can reduce repetitive event setup work.


- **Batch JSON Edit:** Edit all event conditions in bulk using JSON format.
- **Batch JSON Download:** Save all currently registered event conditions as a JSON file.
- **Batch Excel Download:** Download the event list as an Excel file for external reporting or backup purposes.

Batch JSON Edit

1. Click the [Batch JSON Edit] button at the top of the screen.
2. Import a JSON file downloaded via the **Batch JSON Download** feature, or edit the content directly.
3. Click the [Overwrite] button to save.
4. When a confirmation message appears, click the [Overwrite] button.

⚠ After overwriting, you cannot restore the previous event configuration.

- Arbitrarily modifying the values of the `productType` or `platform` properties may cause events to malfunction in other projects.
- The JSON editor includes built-in validation. If the format or values are invalid, an error message appears below the editor and the [Overwrite] button is disabled.



JSON Data Structure



```
{
  "productType": "{productType}",
  "platform": "{platform}",
  "events": {
    "HITMAP": {
      ...
    },
    "LOG_REALTIME": [
      {...}
    ],
  },
}
```




```
"METRICS": [
  {...}
],
"COMPOSITE_LOG": [
  {...}
],
"BASIC": [
  {...}
],
"COMPOSITE_METRICS": [
  {...}
],
"ANOMALY": [
  {...}
]
}
}
```

JSON Field	Description
productType	Product type (e.g. APM, Infra)
platform	Platform information (e.g. java, node)
events	Event data
⬠ HITMAP	Hitmap pattern
⬠ LOG_REALTIME	Real-time log events
⬠ METRICS	Metrics
⬠ COMPOSITE_LOG	Composite log events
⬠ BASIC	Basic events
⬠ COMPOSITE_METRICS	Composite metrics

	JSON Field	Description
	◇ ANOMALY	Anomaly detection

Batch JSON Download

Click the [Batch JSON Download] button at the top of the screen.

- The JSON file name follows the format `integrated-event-YYYYMMDD.json`.
- The structure of the JSON data and the sub-properties under the `events` attribute may vary depending on the product (`platform`) or product type (`productType`).

 If you arbitrarily modify the values of the `productType` or `platform` properties in the JSON data, events may not function correctly in other projects.

Batch Excel Download

Click the [Batch Excel Download] button at the top of the screen.

Metrics

Home > Select Project >  **Alert** > **Event Configuration** > **Metrics** tab

You can view the complete list of all metrics-based events configured for the project. You can review, edit, or delete existing metrics events, or add new events as needed.

Default Options

- **Event Level:** Indicates the severity of an event and is classified into three levels: Normal (**Info**), Warning (**Warning**), and Critical (**Critical**).
- **Duplicate Agent Alert Suppression:** If two or more agents with the same name run simultaneously, events are not triggered for the selected duration (30 seconds, 1 minute, 3 minutes, 5 minutes, 10 minutes).
- **Search:** Search events by event name.
- **JSON Download/Upload:** Edit event rules directly in the JSON editor, or download them as a file, modify them, and upload them again.
- **+ Add Event:** Add a new metrics event.
- **Resize Columns:** Adjust column widths by dragging column borders.

Table | Metrics Event List Structure

Item	Description
No.	Event sequence number
	Event activation status – Enabled: An event is triggered when conditions are met – Disabled: The event does not run and conditions are not evaluated
	Edit or delete event settings
Event Name	User-defined event name – Up to 255 characters

Item	Description
Rule	Event trigger conditions
Target	Monitoring targets to which the event applies
Trigger Count	The time window and number of occurrences required to trigger an event – An event is triggered when conditions are met the configured number of times within the selected duration
Pause	Temporarily suppresses the same event after an alert is triggered – Available durations: Disabled, 5 min, 10 min, 15 min, 20 min, 30 min, 1 hr, 2 hrs, 3 hrs, 6 hrs, 12 hrs, 1 day – If recovery notifications are enabled, alerts are not triggered for the selected duration after a RECOVERED notification is sent
Recovery Notification	Whether to receive a RECOVERED (normal) notification when Critical or Warning events are resolved – Enabled: Sends a RECOVERED notification when the status changes from in progress to normal and allows viewing active and resolved states in event history – Disabled: Triggers a one-time event each time the threshold is exceeded
Event Recipients	Users or groups that receive notifications for the event (recipient tags)

Event Configuration

Metrics events allow you to define events based on conditions that occur in monitored targets and configure notifications so you can respond efficiently.

Add Event

To add a metrics event, go to **Event Configuration > Metrics** tab, then click the [+ Add Event] button in the upper-right corner.

How to Add a Metrics Event

Step 1. Select Template: Choose the basic structure of the event.

Step 2. Define Event Conditions: Configure the conditions that determine when the event is triggered.

Step 3. Select Event Target: Select the monitoring targets where the event applies.

Step 4. Basic Information and Notification Settings: Configure the event name, message, and recipients.

Select Template

- Depending on how you want to configure the event, select one of the three templates on the left side of the screen.
 - Create New:** Freely configure categories and metrics from scratch. This option is suitable for advanced users who need complex conditions or fine-grained monitoring rules.
 - Quick Setup:** Provides predefined categories, metrics, and default thresholds. You can apply it quickly by adjusting only the threshold values, which makes it suitable for beginners.
 - Advanced Setup:** Add conditions or freely modify settings based on predefined categories and metrics.

ⓘ Metrics events support multiple configuration approaches. The available templates may vary depending on the monitoring platform.

Define Event Conditions

In the **Define Event Conditions** step, configure the metrics that determine when the event is triggered, along with the trigger count, pause settings, recovery notifications, and event active time.

Metric Settings


- Enter a category in the search field, or select a category (metric) from the category list. You must select a category to proceed with metric configuration.

- The category list shows the category name, data collection interval, and key. It displays metrics data collected in the project within the last 3 hours.
2. If the metric you need is not listed, select **Direct Input** and enter the category manually.

❗ **If you are not sure about categories or metrics, use the Metrics Finder.**

When configuring metrics for the first time, the Metrics Finder helps you locate metrics more quickly.

- Search for a metric in the Metrics Finder to view related categories and detailed information.
- Click the [Select] button for a metric in the search results to automatically apply the selected value to the event configuration screen.



The screenshot shows the 'Metrics finder' window with a search bar containing 'cpu'. Below the search bar, it indicates '13 cases/ Total 150 cases'. A table lists various metrics with columns for Category, Name, Unit, Type, Description, Merge, and a Select button. The metrics listed are related to CPU usage and load.

Category	Name	Unit	Type	Description	Merge	
Application Process Counter	CpuTime	ms	field	CPU usage time	OM: avg TM: avg	Select
JavaOS	ProcessCpuTime	ms	field	Process CPU Time	OM: avg TM: avg	Select
JavaOS	SystemCpuLoad	pct	field	Number of Avg System CpuLoad Percent	OM: avg TM: avg	Select
JavaOS	ProcessCpuLoad	pct	field	Number of Avg Process CpuLoad Percent	OM: avg TM: avg	Select
JavaOS	CpuLoad	pct	field	Number of Avg CpuLoad Percent	OM: avg TM: avg	Select
Application Host Resource	CpuProc	pct	field	CPU utilization in the Proc area	OM: avg TM: avg	Select
Application Host Resource	CpuCores	count	field	Number of CPU cores	OM: avg TM: avg	Select
Application Host Resource	CpuWait	pct	field	CPU utilization taken to wait for disk IOs	OM: avg TM: avg	Select

At the bottom of the table, there is a pagination control showing '1' of 1 page and '50 / page'.

3. Select the event level (**Critical**, **Warning**, Info). You can configure multiple levels for a single metric.

❗ **Event Level Behavior**

Even when multiple conditions are satisfied, only the highest-priority level (**Critical** → **Warning** → Info) triggers an event. For details, see the [Level Action Guide](#).

- Event levels are classified as Critical (**Critical**), Warning (**Warning**), and Info (Info).



- It is recommended to configure multiple levels for the same metric with stepwise thresholds.

4. Enter the field, operator, and value for the selected event level.

a. Click + **Add** to add more conditions.

- **&&** : An event is triggered only when **all conditions are satisfied simultaneously**
- **||** : An event is triggered when **at least one condition is satisfied**

b. Click — to remove the condition.



For condition syntax and supported operators, refer to the [Condition Settings Guide](#).

Trigger Count

Configure how many times a condition must be met before an event is triggered. This is useful for detecting persistent issues only.

1. Select the trigger frequency.

- **Consecutive:** An event is triggered when the condition occurs consecutively for the configured number of times.
- **Recent:** An event is triggered when the condition occurs the configured number of times within the selected time window.
 - Available durations: 5 seconds , 10 seconds , 15 seconds , 30 seconds , 1 minute , 2 minutes , 3 minutes , 4 minutes , 5 minutes , 10 minutes , 30 minutes , 1 hour

2. Enter the trigger count.

- The **Interval** value corresponds to the data collection interval of the selected category. (Interval: 10s)

Pause

To prevent excessive event generation, select a pause duration for event execution. After an event occurs, the same event will not be triggered again during the configured time period.

- Available durations: Disabled , 5 minutes , 10 minutes , 15 minutes , 20 minutes , 30 minutes , 1 hour , 2 hours , 3 hours , 6 hours , 12 hours , 1 day

ⓘ When recovery notifications are enabled

Even if the same event condition is met again, no notification is sent until the configured time has passed since the **RECOVERED** notification was received.

Resolved Notification

When the resolved notification feature is enabled, you can track state changes after an event occurs (in progress → resolved). The event is shown as an ongoing event in the **Event History** menu.

- Resolved notifications apply only to **Critical** and **Warning** levels. Info-level events always operate as one-time events.
- When a **Critical** or **Warning** event is resolved, a notification with the RECOVERED (normal) state is sent.
- Click the toggle button to **enable** or **disable** this feature.
 -  **Enabled:** Events operate based on state. They enter an in-progress state when the threshold is exceeded and transition to a resolved state when the value recovers below the threshold.
 -  **Disabled:** A one-time event is generated each time the threshold is exceeded.

Event Operation Time

Configure tags so that events operate only during specific time periods (such as business hours, non-business hours, or maintenance windows). If no tag is configured, events operate continuously for 24 hours while enabled.

1. Click **+ Add**.
2. In Event Operation Time, click **+ Create New Tag**.

3. In **Create Event Operation Time Tag**, select the **tag name, days, time, and color**, then click the [Apply] button.
4. You can view created tags in the tag list, and apply them by selecting the checkbox.
5. To modify or delete a tag, click the  icon and update or delete it in **Edit Event Operation Time Tag**.

 When a tag is deleted, it is removed from all users to whom it was applied. However, tags that are currently in use by event rules cannot be deleted.

Simulation

The simulation feature lets you pretest configured event conditions using historical data. If the expected number of events is too high or too low, you can adjust the conditions to find optimal threshold values.

- **Rule Validation:** Verify how newly created event rules would behave when applied to historical data.
- **Threshold Tuning:** Adjust thresholds to an appropriate level when events occur too frequently or not at all.

Start Simulation

1. After completing event condition settings, click the [Simulation] button.
2. Review the expected number of events and the visualized results.

Understanding Simulation Results

1. Time-Series Data

A graph displays metric changes over time along with threshold lines for each level, allowing you to intuitively identify when events are triggered.

2. Event Counts by Level

The number of events for the Critical, Warning, and Info levels is shown on the right side of the screen. This helps you understand overall event distribution and frequency.

3. Target-Level Details

Hover over a specific point on the chart to see the exact metric value and monitoring target at that moment. When multiple targets exist, they are distinguished by color. If patterns differ by target, you can use filtering to analyze them individually.

4. Peak Interval Identification

Identify periods where metrics spike sharply to understand when issues occur most frequently or to recognize recurring patterns.

✓ Simulation Tips

- **Simulate Different Time Ranges**

Run simulations for weekdays vs. weekends and peak vs. off-peak hours to set optimal thresholds for each time period.

- **Include Actual Incident Periods**

Include past outage periods in simulations to verify whether current rules accurately detect real incidents.

- **Add Target Filtering**

When many targets are monitored, simulation results can become complex. Use filters to focus on specific targets for clearer insights.

- **Adjust Thresholds Gradually**

Instead of making large changes at once, adjust thresholds incrementally, such as in 10 percent steps, to find optimal values.

Select Event Target

If no event target is specified, events are triggered for all targets in the project, which can result in excessive events. It is therefore recommended to clearly define the targets to which the event applies.

ⓘ If you change the event targets, the number of events may change. Click the **Simulation** button and run it again to check the updated expected event count.

Target Selection

You can specify the targets to which the event applies using either **Select Input** or **Direct Input**.

- **Select Input** **Direct Input**

1. Select or enter a tag, operator, and value.
 - a. If it is difficult to find a tag, you can search for it using the **Metrics Finder**.
2. Click + **Add** to add event targets using conditions (`&&` , `||`).

Enter the tag, operator, and value of the target to monitor directly.

- When you choose direct input, a list of available tags (**Tag Select Input**) is provided below the input field.

```
ex. endsWith(okindName, 'example_name') && container == 'prod.billing'  
ex. ${4xxErrorType} == '401'
```

 If you change the input method, any content you have entered will be reset.

Basic Information and Notification Settings

Configure the **event name**, **message**, and **recipients**. Because this information is displayed as is in event history and actual notification messages, it should be clear and easy to identify.

Event Activation

Configure whether the event is triggered when the defined conditions are met.

-  **Enabled:** The event is triggered when the configured conditions are satisfied.
-  **Disabled:** Conditions are not evaluated.

Event Name

Enter the event name to be used as the event title. The configured event name is displayed in the **Event List** and **Event History** menus and can be used as a search keyword.

Message

Enter the event message. The message you enter is used in event history and notifications. Click the  icon to view previously written message history. You can include variables in the message to insert actual values at the time the event occurs.

- **Variable Usage Rules**

- Variable format: `${metric_name}`
- Only metrics within the **same category as the configured category** can be used as variables.
- You can combine multiple variables to create a more detailed message.
- The list of available variables can be found in the **Metrics Query** menu.

- **Variable Usage Example**

	Event Configuration	When Event Occurs
Title	Disk Usage Increase	Disk Usage Increase
Message	Disk = <code>\${disk}%</code>	Disk = 89.9%

Event Recipients

Specify the members who will receive event notifications.

- **Receive All Receive by Tag Selection**

Notifications are sent to all targets (members and channels) in the project that have event reception enabled.

Notifications are sent to project members who have the selected tag.

1. Select **Receive by Tag Selection**.
2. In the **Reception Tag** section, click the [+ Add Tag] or [+] button.
3. Under the **Event Reception Tag** window, click + **Create New Tag**.
4. In the **Create Tag** window, enter a tag name, select a color, and click the [Create Tag] button.
5. You can edit or delete created tags by clicking the  icon in the tag list.

6. In the **Event Reception Tag** window, select the desired tag from the **Tag List** to apply it.

ⓘ Reception Tags

Reception tags are a feature for managing users or channels that receive notifications **at a group level**. When a reception tag is assigned to an event, notifications are delivered only to users associated with that tag. If no reception tag is set, notifications are sent to all users in the project.

Edit/Delete Event

1. Go to **Alert Notifications > Event Settings** and open the **Metrics** tab.
2. In the event list, click the  icon for the event you want to edit or delete.
3. In the **Edit Event Rule** window, update the options and click the [Save] button.
 - a. To delete the selected event, click the [Delete] button in the upper-right corner of the **Edit Event Rule** window.

Event Sharing

You can save metric event settings as a JSON file to share configurations with other users or import settings created by others.

- JSON file name: `event-rules-YYYYMMDD.json`

Export

1. Click the [**JSON** ] button at the top right of the screen.
2. When the JSON editor opens, click the [ Export] button.
 - If you use the export function after searching for events, only the searched list will be downloaded as a JSON file.
3. After the JSON file is downloaded, share it with other users.

Import

1. Click the [] button at the top right of the screen.
2. Select the JSON file downloaded using the **Export** function.
3. When the JSON editor opens, choose either [Add to List] or [Overwrite].

 It is recommended to use this feature between projects of the same product type. You can import event settings from projects of different products, but they may not function correctly.

Edit in JSON Format

1. Click the [**JSON** ] button at the top right of the screen.
2. When the editor opens, modify the content according to the JSON format.
3. After completing the edits, click the [Save] button at the bottom of the screen.

 If the modified content does not conform to the JSON format, an error message will be displayed at the bottom of the screen and the content cannot be saved. The error message may vary depending on the type of formatting issue.

JSON error

JSON Data Structure

```
{
  "version": 2,
  "eventId": "zcef7s7f27rum1",
  "enabled": true,
  "stateful": false,
  "title": "Active Transaction",
  "message": "Active Transaction = ${active_tx_count}",
  "category": "app_counter",
  "alertLabel": [
    "oid"
  ],
  "repeatCount": 1,
  "repeatDuration": 0,
```

```

"silent": 300000,
"receiver": [],
"timeTag": [],
"selectString": "",
"conditions": [
  {
    "level": 20,
    "enabled": true,
    "rule": "active_tx_count > 100"
  }
],
"createTime": 1763632674345,
"lastModifiedTime": 1763632674347,
"lastModifiedUser": "support@whatap.io",
"basic": true,
"metaId": "java004",
"selectCondition": {
  "oid": [
    "-1010758404",
    "-250906941"
  ]
}
}

```

Table | JSON Data Structure

JSON Field	Type	Description	Notes
version	Integer	Version of the event rule	Restricted
eventId	String	Unique identifier of the event rule	
enabled	boolean	Whether the event is enabled	
stateful	boolean	Whether the event is state-based	
title	String	Event rule name defined by the user	
message	String	Event rule message defined by the	

JSON Field	Type	Description	Notes
		user	
category	String	Data category	
alertLabel	List<String>	Default identifier (primary key) value assigned per category for event state management	Restricted
repeatCount	Integer	Number of repetitions	
repeatDuration	Long	Duration for repetition	
silent	Integer	Pause duration	
receiver	List<String>	List of recipient tag keys	
timeTag	List<String>	List of operation-time tag keys	
selectString	String	Target selection	
conditions	List<Object>	List of trigger conditions by level	
◇ level	byte	Level value	Critical : 30, Warning : 20, Info : 10
◇ enabled	boolean	Whether the level is enabled	
◇ rule	String	Trigger condition for the level	
createTime	Long	Initial creation time of the event rule	
lastModifiedTime	Long	Last modification time of the event rule	
lastModifiedUser	String	Account that last modified the event rule	

JSON Field	Type	Description	Notes
basic	boolean	Whether the event was created using Quick Setup	Restricted
metaId	String	Unique identifier of the Quick Setup template	Restricted: This value exists only for events created via Quick Setup (basic=true).
selectCondition	Object	Target selection value for Quick Setup	This value exists only for events created via Quick Setup (basic=true).

 Fields marked as “Restricted” may affect event behavior, so it is recommended not to modify their values.

Appendix

Trigger Conditions and Target Selection Guide

Metric alert trigger conditions and event target selection use the same syntax. However, trigger conditions use the field key as the variable, while target selection uses the tag key as the variable.

Level Behavior Guide

Event levels are categorized into Critical, Warning, and Info.

1. Priority-based Event Triggering

When multiple level conditions are met at the same time, only the event with the highest priority is triggered.

Example

Configuration:

- Warning: CPU > 70%
- Critical: CPU > 90%

Current state: CPU 95%

→ Result: Only the Critical event is triggered (Warning is suppressed)

2. Level Escalation

For events that maintain an in-progress state, when a higher-level condition is met while a lower-level event is already in progress, both levels enter the in-progress state.

Example

Configuration:

- Warning: CPU > 70%
- Critical: CPU > 90%

Scenario:

- 1) CPU 80% → Warning triggered (in progress)
- 2) CPU 95% → Warning (in progress), Critical (in progress)

3. Level De-escalation

For events with an in-progress state, when the level transitions from a higher level to a lower level, the higher-level event is resolved while the lower-level event remains active.

Example

Configuration:

- Warning: CPU > 70%
- Critical: CPU > 90%

Scenario:

- 1) CPU 97% → Critical triggered
- 2) CPU 85% → Critical resolved, Warning remains
- 3) CPU 65% → Warning resolved

4. Execution Flow

- Warning: CPU > 70%

- Critical: CPU > 90%

Metric value: 60% → 75% → 92% → 85% → 60%

State: Info Warning Critical Warning Info
 (triggered) (triggered) (maintained) (resolved)
 Warning
 (maintained)

- ⓘ Events that have an in-progress state remain active as long as their conditions are met and continue until those conditions are cleared.

Condition Configuration Guide

- ⓘ If a metric name starts with a number or contains special characters, you must wrap it in the `${metric_name}` format.

- Brackets `()`, `[]`
- Operators `+`, `-`, `*`, `/`, `%`
- Delimiters `:`, `@`, `#`, `,`, `,`, (space)
- Other special characters `!`, `^`, `&`, `|`, `~`, `\"`, `=`

Examples using `${}` :

```
${cpu(xos)} > 50
${mem[0]} >= 100
${cpu-usage} > 80
${namespace:cpu} > 70
${metric name} > 60
${4xx_error} > 10
```

Usable without `${}` : Metric names that contain only alphabetic characters, underscores (`_`), or dots (`.`)



```
cpu > 80
cpu_usage > 50
CPUUtilization.Average > 0.8
```

1. Comparison Operators

Numeric Comparison

```
Greater than:      cpu > 80
Greater than or equal:  cpu >= 80
Less than:         memory < 1000
Less than or equal:  memory <= 1000
Equal to:          status == 200
Not equal to:       error != 0
```

String Comparison

```
status == 'OK'
region == "us-east-1"
```

2. Arithmetic Operators

- Basic arithmetic operations

```
Addition:      cpu + 10 >= 90
Subtraction:    memory - 100 >= 500
Multiplication: cpu * 2 >= 100
Division:       disk / 1024 >= 100
Modulo:         value % 10 == 0
```

- Controlling precedence with parentheses

```
(cpu + memory) * 2 >= 200
(disk - used) / total >= 0.2
((cpu + memory) / 2) >= 50
```

- Negative values

```
cpu > -100
```

3. Logical Operators

- AND operator (&&)

```
cpu > 80 && memory > 1000  
${cpu(xos)} > 50 && ${mem(xos)} > 60
```

- OR operator (| |)

```
cpu > 90 | | memory > 90  
disk < 10 | | network > 1000
```

- Combined logical operations

```
(cpu > 50 && memory > 50) | | disk < 20  
cpu > 80 && (memory > 1000 | | disk > 500)
```

4. Pattern Matching Operators

- LIKE operator

```
oname like 'prod-*'  
url like '*error*'  
message like 'WARN%'
```

- NOT LIKE operator

```
oname not like 'test-*'  
url not like '*debug*'
```

5. Built-in Functions

- Null check functions

```

isNull(value)      # Checks whether the value is null
isNotNull(value)   # Checks whether the value is not null
nvl(value, 0)      # Returns a default value if null
isEmpty(str)       # Checks whether the string is empty
isNotEmpty(str)    # Checks whether the string is not empty

```

- Aggregate functions

```

sum(cpu, memory, disk)    # Sum
avg(cpu, memory)          # Average
max(cpu, memory, disk)    # Maximum
min(cpu, memory, disk)    # Minimum
count(value1, value2, value3) # Count

```

- Math functions

```

round(cpu, 2)    # Rounds to 2 decimal places

```

- String functions

```

length(str)      # String length
startsWith(str, 'prefix') # Checks prefix
endsWith(str, 'suffix') # Checks suffix
indexOf(str, 'search') # String index
substring(str, 0, 10) # Substring
trim(str)        # Trims whitespace
replace(str, 'old', 'new') # String replacement
hasStr(str, 'search') # Checks string containment

```

- Conditional functions

- **if:** A conditional function that returns the `trueValue` if the condition evaluates to `true`, or the `falseValue` if it evaluates to `false`

Syntax

```

if(condition, trueValue, falseValue)

```

Example

```
if(value >= 90, 'High', 'Medium') == 'High'
```

- **decode:** Compares a value against conditions and returns the result of the matching condition, or a default value if no match is found

Syntax

```
decode(value, condition1, result1, condition2, result2, ..., conditionN, resultN, defaultValue)
```

Example

```
decode(value, 1, 'Low', 2, 'Medium', 3, 'High', 'Unknown') == 'Unknown'
```

- **in:** Returns `true` if the value matches any of the specified candidates; otherwise returns `false`

Syntax

```
in(value, candidate1, candidate2, ..., candidateN)
```

Example

```
in(status, 200, 201, 204) == false
```

! Writing Guidelines

- Recommended
 - Use simple variable names without `${}` (e.g. `cpu`, `memory`)
 - Always use `${}` when special characters are included
 - Use parentheses to clearly separate complex expressions
 - Enter function names exactly as defined (case-sensitive)
- Avoid



- Incomplete expressions (e.g. `cpu > , memory && )`)
- Mismatched parentheses (e.g. `(cpu > 80 && cpu > 70) )`)
- Consecutive or invalid operators (e.g. `cpu >> 80 , cpu >< 80 )`)

Real-time Log Events

Home > Select Project >  **Alert** > **Event Configuration** > **Real-time Log** tab

You can set alert notifications by filtering collected log data against conditions. Log monitoring must be activated to use log events. For more information on **activating log monitoring**, refer to the log documentation for each product.

- **Real-time log event:** An event is triggered when the configured event conditions are met in real-time collected logs.

Default options

- **Search:** Search for events by selecting an event name or event notification tag
- **JSON download/upload:** Event rules can be directly edited through the JSON editor, or downloaded as a file, modified, and re-uploaded
- **Notification message settings:** Select detailed items for the message to receive and edit the content for each item
- **+ Add event:** Add a new log event

Table | Real-time log event list configuration

Item	Description
Event name	Event name used as the notification title - Up to 255 bytes
Category	Log classification name (log folder name)
Search key	The field to search for in log data - Example: To access an HTTP response status code, use the search key <code>status</code>
Search value	The actual value corresponding to the search key. An alert is sent when logs contain the specified value

Item	Description
	- Example: If the search key is <code>status</code> and the search value is <code>200</code>, an alert is triggered when logs containing HTTP status code 200 are collected
Pause event notifications	Pauses the same event notification for a specified time after an event notification is triggered - Available times: Disabled, 5 min, 10 min, 15 min, 20 min, 30 min, 1 hour, 2 hours, 3 hours, 6 hours, 12 hours, 1 day - If resolved notifications are enabled, no notification is sent for the same event during the selected time after receiving a normal (RECOVERED) status notification
Event target filtering	Conditions that define the targets to which the event applies - If no condition is set, alerts are evaluated for all collected data - Select input: Build conditions by selecting tags, operators, and values - Direct input: Manually enter conditions to specify targets
Event notification tags	User or group (notification tag) to receive the event notification
	Edit and delete events
	Whether the event is active - Enabled: Event occurs when conditions are met - Disabled: Event does not operate and conditions are not checked

Event settings

Log events allow you to configure alert notifications when logs are collected in real time or when logs that meet specific conditions accumulate.

Add real-time log event

1. Click the **Real-time Log** tab in **Alert > Event Configuration**.

2. Click the [+ Add event] button on the right side of the screen.
3. Set the log event conditions in the **Add event rule** window.
 - **Event name:** Enter a name to use as the event title.
 - **Event activation:** When the toggle button is activated, an event occurs when conditions are met.
 - **Level:** Select the event level (**Critical**, **Warning**, **Info**).
 - **Message:** Enter the notification message to receive when the event is triggered. You can use variables in the message to include actual values at the time of event occurrence.

✓ **Variable usage rules**

- Variable format: `${metricName}`
 - Only metrics within the **same category as the configured category** can be used as variables.
 - You can combine multiple variables to compose a detailed message.
 - A list of available variables can be checked in **Metrics Query**.
- **Category:** Select the log classification name (log folder name) from the list or enter it directly. You must select a category before entering a search key and search value.
 - **Search key:** Enter the field you want to search for in log data.
 - Example: To access an HTTP response status code, use the search key `status`.
 - **Search value:** Enter the actual value corresponding to the **search key**. An alert is sent when logs contain the specified value.
 - Example: If the **search key** is `status` and the **search value** is `200`, an alert is triggered when logs containing HTTP status code 200 are collected.
 - **Event target filtering:** Enter conditions to define the targets to which the event applies. If no value is entered, alerts are evaluated for all collected data.
 - **Select input:** Select a tag, operator, and value to set the target conditions.

ⓘ For condition expression syntax and supported operators, refer to the [Condition](#)

! settings guide.

- **Direct input:** Enter conditions directly.
- **Pause event notifications:** Select a duration to temporarily suppress alerts. After the event is triggered, the same event will not occur during the selected time.
 - Available times: Disabled , 1 min , 2 min , 3 min , 5 min , 10 min , 20 min , 30 min , 1 hour , 3 hours , 6 hours , 12 hours
- **Event operation time:** Set tags to activate the event during specific time periods (work/non-work/maintenance hours, etc.). If no tag is set, the event operates 24 hours a day when activated.
 1. Click + **Add**.
 2. Click + **Create new tag** in **Event operation time**.
 3. In **Create event operation time tag**, select the **tag name, day of week, time, and color**, then click the [Apply] button.
 4. Created tags can be viewed in the tag list and are applied by selecting the checkbox.
 5. Click the  icon for the tag to modify or delete, and edit or delete the tag in **Edit event operation time tag**.

! Deleting a tag removes the tag for all users to whom it has been applied. However, tags in use by event rules cannot be deleted.

- **Event notification tags:** If notification tags are set for an event, notifications are sent to project members with those tags. If no recipient tags are set, notifications are sent to all users in the project.
 1. Click + **Add tag** or +.
 2. Click + **Create new tag** at the bottom of the **Event notification tags** window.
 3. In the **Create tag** window, enter the tag name, select a color, and click the [Create tag] button to create the tag.

4. Edit or delete created tags by clicking the  icon in the tag list.
5. Select the desired tag from the **tag list** in the **Event notification tags** window to apply it.

✓ When configuring events, you can send notifications to project members and 3rd-party plugins with the selected event notification tags. You can assign tags to project members and 3rd-party plugins respectively in **Alert notifications > Event notification settings**.

4. When configuration is complete in the **Add event rule** window, click the [Save] button.

Edit/delete events

1. Go to the **Real-time log** tab in **Alert > Event Configuration**.
2. Click the  icon for the event to edit or delete in the log event list.
3. Modify options in the **Event settings** screen and click the [Save] button.
 - a. To delete the selected event, click the [Delete] button in the upper right of the **Event settings** screen.

Event sharing

You can save metric event settings as a JSON file to share configurations with other users or import settings created by others.

- JSON file name: **event-rules- YYYYMMDD .json**

Export

1. Click the [**JSON** ] button at the top right of the screen.
2. When the JSON editor opens, click the [ Export] button.
 - If you use the export function after searching for events, only the searched list will be downloaded

as a JSON file.

3. After the JSON file is downloaded, share it with other users.

Import

1. Click the [] button at the top right of the screen.
2. Select the JSON file downloaded using the **Export** function.
3. When the JSON editor opens, choose either [Add to List] or [Overwrite].

 It is recommended to use this feature between projects of the same product type. You can import event settings from projects of different products, but they may not function correctly.

Edit in JSON Format

1. Click the [JSON ] button at the top right of the screen.
2. When the editor opens, modify the content according to the JSON format.
3. After completing the edits, click the [Save] button at the bottom of the screen.

 If the modified content does not conform to the JSON format, an error message will be displayed at the bottom of the screen and the content cannot be saved. The error message may vary depending on the type of formatting issue.

JSON error

Appendix

Trigger Conditions and Target Selection Guide

Metric alert trigger conditions and event target selection use the same syntax. However, trigger conditions use the field key as the variable, while target selection uses the tag key as the variable.

Level Behavior Guide

Event levels are categorized into Critical, Warning, and Info.

1. Priority-based Event Triggering

When multiple level conditions are met at the same time, only the event with the highest priority is triggered.

Example

```
Configuration:  
- Warning: CPU > 70%  
- Critical: CPU > 90%  
  
Current state: CPU 95%  
→ Result: Only the Critical event is triggered (Warning is suppressed)
```

2. Level Escalation

For events that maintain an in-progress state, when a higher-level condition is met while a lower-level event is already in progress, both levels enter the in-progress state.

Example

```
Configuration:  
- Warning: CPU > 70%  
- Critical: CPU > 90%  
  
Scenario:  
1) CPU 80% → Warning triggered (in progress)  
2) CPU 95% → Warning (in progress), Critical (in progress)
```

3. Level De-escalation

For events with an in-progress state, when the level transitions from a higher level to a lower level, the higher-level event is resolved while the lower-level event remains active.

Example

Configuration:

- Warning: CPU > 70%
- Critical: CPU > 90%

Scenario:

- 1) CPU 97% → Critical triggered
- 2) CPU 85% → Critical resolved, Warning remains
- 3) CPU 65% → Warning resolved

4. Execution Flow

- Warning: CPU > 70%
- Critical: CPU > 90%

Metric value: 60% → 75% → 92% → 85% → 60%

State:	Info	Warning	Critical	Warning	Info
		(triggered)	(triggered)	(maintained)	(resolved)
			Warning		
			(maintained)		

- ⓘ Events that have an in-progress state remain active as long as their conditions are met and continue until those conditions are cleared.

Condition Configuration Guide

- ⓘ If a metric name starts with a number or contains special characters, you must wrap it in the `${metric_name}` format.

- Brackets `()`, `[]`
- Operators `+`, `-`, `*`, `/`, `%`
- Delimiters `:`, `@`, `#`, `,`, `.` (space)
- Other special characters `!`, `^`, `&`, `|`, `~`, `"`, `=`

Examples using `${}` :



```

${cpu(xos)} > 50
${mem[0]} >= 100
${cpu-usage} > 80
${namespace:cpu} > 70
${metric name} > 60
${4xx_error} > 10

```

Usable without `${}` : Metric names that contain only alphabetic characters, underscores (`_`), or dots (`.`)

```

cpu > 80
cpu_usage > 50
CPUUtilization.Average > 0.8

```

1. Comparison Operators

Numeric Comparison

```

Greater than:      cpu > 80
Greater than or equal:  cpu >= 80
Less than:         memory < 1000
Less than or equal:  memory <= 1000
Equal to:          status == 200
Not equal to:       error != 0

```

String Comparison

```

status == 'OK'
region == "us-east-1"

```

2. Arithmetic Operators

- Basic arithmetic operations

```

Addition:    cpu + 10 >= 90
Subtraction: memory - 100 >= 500
Multiplication: cpu * 2 >= 100

```

```
Division:  disk / 1024 >= 100
Modulo:    value % 10 == 0
```

- Controlling precedence with parentheses

```
(cpu + memory) * 2 >= 200
(disk - used) / total >= 0.2
((cpu + memory) / 2) >= 50
```

- Negative values

```
cpu > -100
```

3. Logical Operators

- AND operator (&&)

```
cpu > 80 && memory > 1000
${cpu(xos)} > 50 && ${mem(xos)} > 60
```

- OR operator (| |)

```
cpu > 90 | | memory > 90
disk < 10 | | network > 1000
```

- Combined logical operations

```
(cpu > 50 && memory > 50) | | disk < 20
cpu > 80 && (memory > 1000 | | disk > 500)
```

4. Pattern Matching Operators

- LIKE operator

```
oname like 'prod-*'
url like '*error*'
```

```
message like 'WARN%'
```

- NOT LIKE operator

```
oname not like 'test-*'
url not like '*debug*'
```

5. Built-in Functions

- Null check functions

```
isNull(value)      # Checks whether the value is null
isNotNull(value)   # Checks whether the value is not null
nvl(value, 0)      # Returns a default value if null
isEmpty(str)       # Checks whether the string is empty
isNotEmpty(str)    # Checks whether the string is not empty
```

- Aggregate functions

```
sum(cpu, memory, disk)  # Sum
avg(cpu, memory)        # Average
max(cpu, memory, disk)  # Maximum
min(cpu, memory, disk)  # Minimum
count(value1, value2, value3) # Count
```

- Math functions

```
round(cpu, 2)          # Rounds to 2 decimal places
```

- String functions

```
length(str)           # String length
startsWith(str, 'prefix') # Checks prefix
endsWith(str, 'suffix') # Checks suffix
indexOf(str, 'search')  # String index
substring(str, 0, 10)   # Substring
```

```
trim(str)          # Trims whitespace
replace(str, 'old', 'new') # String replacement
hasStr(str, 'search') # Checks string containment
```

- Conditional functions

- **if:** A conditional function that returns the `trueValue` if the condition evaluates to `true`, or the `falseValue` if it evaluates to `false`

Syntax

```
if(condition, trueValue, falseValue)
```

Example

```
if(value >= 90, 'High', 'Medium') == 'High'
```

- **decode:** Compares a value against conditions and returns the result of the matching condition, or a default value if no match is found

Syntax

```
decode(value, condition1, result1, condition2, result2, ..., conditionN, resultN, defaultValue)
```

Example

```
decode(value, 1, 'Low', 2, 'Medium', 3, 'High', 'Unknown') == 'Unknown'
```

- **in:** Returns `true` if the value matches any of the specified candidates; otherwise returns `false`

Syntax

```
in(value, candidate1, candidate2, ..., candidateN)
```

Example

```
in(status, 200, 201, 204) == false
```

Writing Guidelines

- Recommended
 - Use simple variable names without `${}` (e.g. `cpu`, `memory`)
 - Always use `${}` when special characters are included
 - Use parentheses to clearly separate complex expressions
 - Enter function names exactly as defined (case-sensitive)
- Avoid
 - Incomplete expressions (e.g. `cpu >` , `memory &&`)
 - Mismatched parentheses (e.g. `(cpu > 80 && cpu > 70)`)
 - Consecutive or invalid operators (e.g. `cpu >> 80` , `cpu >< 80`)

Composite Log Event Settings

Home > Select Project >  **Alert** > **Event Configuration** > **Composite Log** tab

You can set alert notifications by filtering collected log data against conditions. Log monitoring must be activated to use log events. For more information on **activating log monitoring**, refer to the log documentation for each product.

- **Composite log event:** An event is triggered when the number of logs that meet the log search conditions within the configured **data query range** reaches or exceeds the threshold count.

Permission

Alert configuration permission is required to configure composite metrics events. Alert configuration permission can be set by clicking the  icon for the member to grant permission in **Home Screen** > **Select Project** > **Management** > **Integrated Member Management**.

Default options

- **Search:** Search for events by selecting an event name or event notification tag
- **JSON download/upload:** Event rules can be directly edited through the JSON editor, or downloaded as a file, modified, and re-uploaded
- **Notification message settings:** Select detailed items for the message to receive and edit the content for each item
- **+ Add event:** Add a new log event

Table | Composite log event list configuration

Item	Description
Event name	Event name used as the notification title
Template	Composite log template

Item	Description
Category	Log classification name (log folder name)
Rule	Conditions that trigger the event
Additional notification when event is resolved	Whether to receive notifications in normal (RECOVERED) status when Critical and Warning level events are resolved - Enabled: Sends a RECOVERED notification when transitioning from in-progress → normal status, and the in-progress/resolved status can be distinguished in event history - Disabled: A one-time event occurs each time the threshold is exceeded
Pause event notifications	Pauses the same event notification for a specified time after an event notification is triggered - Available times: Disabled, 5 min, 10 min, 15 min, 20 min, 30 min, 1 hour, 2 hours, 3 hours, 6 hours, 12 hours, 1 day - If resolved notifications are enabled, no notification is sent for the same event during the selected time after receiving a normal (RECOVERED) status notification
Interval	Time range for querying log data
Event notification tags	User or group (notification tag) to receive the event notification
	Edit and delete events
	Whether the event is active - Enabled: Event occurs when conditions are met - Disabled: Event does not operate and conditions are not checked

Event settings

The composite log event feature triggers alert notifications when the number of logs matching the log search conditions within the configured **data query range** reaches or exceeds the threshold count.

Add composite log event

1. Click the **Composite Log** tab in **Alert > Event Configuration**.
2. Click the [+ Add event] button on the right side of the screen.
3. Set the log event conditions in the **Add event rule** window. For details, refer to [Basic settings](#), [Detail settings](#), and [Additional settings](#).

Add composite log event

4. When configuration is complete in the **Add event rule** window, click the [Save] button.

Basic settings

- **Event name:** Enter a name to use as the event title.
- **Event activation:** When the toggle button is activated, an event occurs when conditions are met.
- **Level:** Select the event level (**Critical**, **Warning**, Info).

Detail settings

Configure [Template \(preset\)](#), [Message](#), [Category](#), [Filter](#), and [Event trigger conditions](#).

Template (preset)

Selecting a predefined preset automatically configures the **filter**, **query type**, and **event trigger conditions**. You can modify the settings after selecting a preset.

Composite log template types		
Preset	Query type	Default rule
None	none	rows > 10
2xx status code count	count	count > 10



Preset	Query type	Default rule
3xx status code count	count	count > 10
4xx status code count	count	count > 10
5xx status code count	count	count > 10
Normal status code (2xx, 3xx) count	count_v2	include_minus_exclude_count > 10
Error status code (4xx, 5xx) count	count_v2	include_minus_exclude_count > 10
Error received count	count	count > 10
Exception received count	count	count > 10
Log count per second	rps	rps > 100
Field aggregation (Aggregation)	aggr	avg > 1000
Rate	rate	rate > 10
Unique value count (Cardinality)	cardinality	cardinality > 100
Percentile	percentile	p99 > 3000

Message

Enter the notification message to receive when an event occurs. You can use variables in the message to compose dynamic notifications.

- **Example:** Message `${host}` error `${count}` occurred → Notification message web-01 error 15 occurred

Table | Message variables

Variable	Description	Usage condition
<code>\${count}</code>	Number of logs matching the condition	Always available
<code>\${groupField}</code>	Grouping field name	Grouped notifications only
<code>\${groupValue}</code>	Grouping field value	Grouped notifications only
<code>\${fieldName}</code>	Grouping field value (referenced directly by field name)	Grouped notifications only

❗ The `${count}` variable can be used in all notifications regardless of grouping. The `${groupField}`, `${groupValue}`, and `${fieldName}` variables are only substituted when a grouping field is configured.

Category

Select the log classification name (log folder name) from the list or enter it directly. You can enter the search key and search value only after selecting a category.

Filter

Enter log search conditions using Lucene query syntax. Two input modes are supported.

- **Filter mode** : Build a query by selecting fields and values.
- **Code mode** : Enter Lucene queries directly.

Table | Lucene query examples

Query	Description
<code>level:ERROR</code>	Logs where the level field is ERROR
<code>status_nxx:status_4xx OR status_nxx:status_5xx</code>	4xx or 5xx status codes
<code>appender:accessLog</code>	Logs from the accessLog appender
<code>response_time.n:*</code>	Logs where the response_time.n field exists

Query	Description
host:web-01 AND level:ERROR	Error logs from host web-01

Event trigger conditions

Configure by query type. Additional settings are displayed depending on the selected template (preset) or query type.

Table | Event trigger condition key meanings

Key	Meaning
rows / count	Number of logs matching the filter condition
rps	Log count per second (count / interval in seconds)
avg	Average value of the target field
sum	Sum of the target field values
min	Minimum value of the target field
max	Maximum value of the target field
rate	Numerator count / denominator count × 100 (%)
cardinality	Number of unique values in the target field
p99 / p95 / p90 / p75 / p50	Percentile values of the target field
include_minus_exclude_count	Include filter count – exclude filter count

- **Count (count / none):** Counts the number of logs matching the configured filter conditions.

Table | Count settings

Setting	Description
Filter	Log search condition (Lucene query)
Event trigger condition	<code>count</code> or <code>matching logs</code> <code>{operator}</code> <code>{threshold}</code> (e.g., <code>count > 10</code>)

- **Log count per second (rps):** Calculates the number of logs per second (Requests Per Second) matching the configured filter conditions. The total count during the interval is divided by the interval in seconds.

Table | Log count per second settings

Setting	Description
Filter	Log search condition (Lucene query)
Event trigger condition	<code>RPS</code> <code>{operator}</code> <code>{threshold}</code> (e.g., <code>rps > 0.5</code>)

- **Field aggregation (aggr):** Calculates the aggregated value (average, sum, min, max) of the specified numeric field. The aggregation method can be selected from a dropdown in the event trigger condition.

Table | Field aggregation settings

Setting	Description
Filter	Log search condition (Lucene query)
Target field	Numeric field to aggregate (e.g., <code>response_time.n</code>)
Event trigger condition	<code>avg / sum / min / max</code> <code>{operator}</code> <code>{threshold}</code> (e.g., <code>avg > 1000</code>)

- **Rate:** Calculates the percentage (%) ratio of the numerator filter count to the denominator filter count. Useful for monitoring the ratio of specific conditions against all logs.

Table | Rate settings

Setting	Description
Filter (numerator)	Log condition for the numerator (e.g., <code>level:ERROR</code>)
Denominator filter	Log condition for the denominator (e.g., <code>level:*</code>)
Event trigger condition	<code>Rate (%) {operator} {threshold}</code> (e.g., <code>rate > 10</code>)

- **Example:** Filter `level:ERROR`, denominator filter `level:*`, rule `rate > 10` → Alert triggered when the error rate among all logs exceeds 10%
- **Unique value count (cardinality):** Counts the number of unique values in the specified field. Useful for detecting an abnormally large variety of values.

Table | Unique value count settings

Setting	Description
Filter	Log search condition (Lucene query)
Target field	Field to calculate unique values for (e.g., <code>host</code>)
Event trigger condition	<code>Unique value count {operator} {threshold}</code> (e.g., <code>cardinality > 100</code>)

- **Percentile:** Calculates the percentile value of the specified numeric field. Useful for monitoring tail latency in response times. Changing the percentile selection automatically updates the key in the event trigger condition.

Table | Percentile settings

Setting	Description
Filter	Log search condition (Lucene query)
Target field	Numeric field to calculate percentile for (e.g., <code>response_time.n</code>)
Percentile	Select from P99, P95, P90, P75, P50

Setting	Description
Event trigger condition	P99 / P95 / ... {operator} {threshold} (e.g., p99 > 3000)

Table | Percentile meanings

Percentile	Meaning
P50	Median. 50% of all data is at or below this value
P75	75% of all data is at or below this value
P90	90% of all data is at or below this value
P95	95% of all data is at or below this value
P99	99% of all data is at or below this value

- **Include/exclude count (count_v2):** Calculates the net count by subtracting the exclude filter count from the include filter count.

Table | Include/exclude count settings

Setting	Description
Filter (include)	Log conditions to include (e.g., status_nxx:status_4xx OR status_nxx:status_5xx)
Filter exclude	Log conditions to exclude (e.g., status_nxx:status_4xx)
Denominator filter	Denominator for ratio calculation (optional)
Event trigger condition	Include-exclude count {operator} {threshold} (e.g., include_minus_exclude_count > 10)

- **Example:** Filter status_nxx:status_4xx OR status_nxx:status_5xx , filter exclude status_nxx:status_4xx → Alert triggered when the net 5xx count (4xx+5xx minus 4xx) exceeds 10

Group by field (Group By)

When a group by field is configured, alert conditions are evaluated independently for each value of that field. Select the field to group by from the dropdown. Clicking the dropdown shows a list of available fields for the current category with autocomplete.

- **Example:** Filter `level:ERROR` , group by field `host` , rule `rows > 3`

Table | Grouping comparison

Item	Without grouping	With grouping
Evaluation unit	All logs aggregated as one	Independent aggregation per group value
Alert occurrence	1 alert	One per group meeting the condition
Alert title	Event title	[Group value] Event title
Stateful	State transition based on all logs	Independent state transition per group

- `[web-01] Error alert` — 8 errors from web-01 → alert triggered
- `[web-02] Error alert` — 5 errors from web-02 → alert triggered
- `[api-01] Error alert` — 2 errors from api-01 → condition not met, no alert

❗ When **Additional notification when event is resolved** is enabled, the ABNORMAL → RECOVERED state transition occurs independently per group. For example, when errors from `web-01` are resolved, a `[web-01] RECOVERED` notification is triggered, with no effect on other hosts.

Additional settings

- **Interval:** Check alert conditions at the selected time interval (`unset` , `1 min` , `5 min` , `10 min` , `30 min` , `1 hour` , `12 hours`).
- **Mute:** After an event occurs, the same event is not triggered during the selected time (`unset` , `1 min` , `5 min` , `10 min` , `30 min` , `1 hour` , `12 hours`).

- **Event operation time:** Set tags to activate the event during specific time periods (work/non-work/maintenance hours, etc.). If no tag is set, the event operates 24 hours a day when activated.
 1. Click **+ Add**.
 2. Click **+ Create new tag** in **Event operation time**.
 3. In **Create event operation time tag**, select the **tag name**, **day of week**, **time**, and **color**, then click the **[Apply]** button.
 4. Created tags can be viewed in the tag list and are applied by selecting the checkbox.
 5. Click the  icon for the tag to modify or delete, and edit or delete the tag in **Edit event operation time tag**.

 Deleting a tag removes the tag for all users to whom it has been applied. However, tags in use by event rules cannot be deleted.

- **Event notification tags:** If notification tags are set for an event, notifications are sent to project members with those tags. If no recipient tags are set, notifications are sent to all users in the project.
 1. Click **+ Add tag** or **+**.
 2. Click **+ Create new tag** at the bottom of the **Event notification tags** window.
 3. In the **Create tag** window, enter the tag name, select a color, and click the **[Create tag]** button to create the tag.
 4. Edit or delete created tags by clicking the  icon in the tag list.
 5. Select the desired tag from the **tag list** in the **Event notification tags** window to apply it.

 When configuring events, you can send notifications to project members and 3rd-party plugins with the selected event notification tags. You can assign tags to project members and 3rd-party plugins respectively in **Alert notifications > Event notification settings**.

Simulation

You can simulate using historical data before saving the alert to check the expected number of alert occurrences.

1. After completing the alert configuration, select a time range at the bottom.
2. Click the [Simulation] button.
3. The chart displays data by time period and the number of alert occurrences.

 You can adjust the simulation time range to check alert occurrence patterns across different time periods. Data from the last 1–2 minutes may not be available due to index build delays.

4. When configuration is complete in the **Add event rule** window, click the [Save] button.

Create alert from Log Explorer

You can create alerts directly from Log Explorer.

1. Navigate to **Log > Log Search** menu.
2. Configure the category and search query.
3. Click the [Add event] button in the upper right of the screen.
4. The alert configuration form opens in a drawer, with the current category and search query automatically filled in.
5. Complete additional settings and click the [Save] button.
6. After saving, click the [Go] button to navigate to the alert list.

Edit/delete events

1. Go to the **Composite log** tab in **Alert > Event Configuration**.
2. Click the  icon for the event to edit or delete in the log event list.
3. Modify options in the **Event settings** screen and click the [Save] button.
 - a. To delete the selected event, click the [Delete] button in the upper right of the **Event settings** screen.

Event sharing

You can save metric event settings as a JSON file to share configurations with other users or import settings created by others.

- JSON file name: `event-rules-YYYYMMDD.json`

Export

1. Click the [**JSON** ] button at the top right of the screen.
2. When the JSON editor opens, click the [ Export] button.
 - If you use the export function after searching for events, only the searched list will be downloaded as a JSON file.
3. After the JSON file is downloaded, share it with other users.

Import

1. Click the [] button at the top right of the screen.
2. Select the JSON file downloaded using the **Export** function.
3. When the JSON editor opens, choose either [Add to List] or [Overwrite].

 It is recommended to use this feature between projects of the same product type. You can import event settings from projects of different products, but they may not function correctly.

Edit in JSON Format

1. Click the [**JSON** ] button at the top right of the screen.
2. When the editor opens, modify the content according to the JSON format.
3. After completing the edits, click the [Save] button at the bottom of the screen.

- ❗ If the modified content does not conform to the JSON format, an error message will be displayed at the bottom of the screen and the content cannot be saved. The error message may vary depending on the type of formatting issue.

JSON error

Event Notification Settings

Home > Select Project >  **Alert** > **Notifications**

Event Notification Settings per User

 Information about other members in **Event notification settings per-user** can only be viewed by members with the **Member Management** permission. For all other members, required information is masked.

- Mobile device management for bulk event notification receipt settings and access control can be managed in **Account Information**.
- For detailed information about member permissions, refer to the [following document](#).

Notification Methods

Alerts are available via [Email](#), [SMS](#), [WhatsApp](#), and [Mobile](#).

In the **Event notification settings per-user** section, select the checkbox for the desired notification method to receive alerts. If you clear the checkbox for a notification method, alerts will not be sent through that channel.

 All users except the project super administrator can change only their own notification settings.

Email Alerts

Email alerts are sent to the email address provided during sign-up.

SMS Alerts

You can receive alerts via SMS.

ⓘ The SMS alert feature is available only for paid projects.

1. Click your **profile icon** at the top right of the screen, then go to **Manage Accounts**.
2. In **Account Information**, click the [General Mobile Phone] button in the **User Phone Number** section.
3. Enter the **phone number** that will receive the verification code.
 - Only **Korean mobile phone numbers** can be registered to receive SMS alerts.
4. Click the [Save] button.

✔ Change or delete the registered phone number

- To change the phone number, click the [Change] button. SMS verification is required.
- To delete the phone number, click the [Delete] button. After deletion, you will no longer receive alerts via SMS.

WhatsApp Alerts

You can receive alerts via WhatsApp.

ⓘ The WhatsApp alert feature is available only for paid projects.

1. Click your **profile icon** at the top right of the screen and go to **Manage Accounts**.
2. In **Account Information**, click the [WhatsApp] button in the **User Phone Number** section.
3. Enter the **phone number** that will receive the verification code.
4. Click the [Send Verification Code] button.
5. Enter the 6-digit verification code sent to your WhatsApp application.
 - If you do not receive the verification code within 10 minutes, click the [Resend Verification Code] button.
6. Click the [Verify] button.

Mobile Alerts

Users who need to receive mobile alerts must install the WhaTap application on their mobile device and log in.

- Download the WhaTap application – [Android](#)
- Download the WhaTap application – [iOS](#)

Repeated Alerts (Escalation) Settings

If the alert condition is not resolved after an alert is triggered, you can configure repeated alert intervals based on the time of the initial alert. This setting applies only to **Critical** level alerts.

For example, to repeat alerts immediately (0 minutes), 1 hour later, and 1 day later after the initial alert, enter `0,1H,1D` in the **Repeated Alerts (Escalation)** column.

1. Go to **Alert > Notifications** and navigate to the **Event notification settings per-user** section.
2. Configure the settings in the **Repeated Alerts (Escalation)** field.
 - **M**: minutes, **H**: hours, **D**: days. If the unit is omitted, the value is interpreted as minutes.
 - Enter a number or a number with a unit (**M, H, D**). If the input format is invalid, an error message is displayed.
3. Click the [Save] button. The settings will not be applied unless they are saved.

Event Reception Tags

When configuring events, you can select **event reception tags** to send alerts to project members and 3rd-party plugins associated with those tags. If no event reception tag is set, alert notifications are sent to **all project members**.

Add / Edit / Delete Event Reception Tags

 Event reception tags that are **currently applied to an event** cannot be deleted.

1. Go to **Alert > Notifications**.

2. In the user list under the **Event notification settings per-user** section, click + **Add Tag** or +.
3. In the **Event reception tag** window, click + **Create New Tag**.
4. In the **Create Tag** dialog, enter a tag name, select a color, and click the [Create Tag] button.
5. To edit or delete a created tag, click the  icon in the tag list.
6. Select the desired tag from the **Tag List** in the **Event reception tag** window to apply it.

Remove Event Reception Tags

1. Go to **Alert > Notifications** and open the **Event notification settings per-user** section.
2. Click + in the **Event reception tag** field for the user whose event reception tag you want to remove.
3. In the **Event reception tag** window, click the **X** on the tag you want to remove. The tag will be removed.

Receive Alerts Without Reception Tags

When configuring events, you can select event reception tags to send alerts to project members and 3rd-party plugins associated with those tags.

- To receive alerts **without event reception tags**, enable the **Receive alerts without reception tags** option.
- If you want to receive **only alerts that have event reception tags**, disable this option.

 To stop receiving all alerts, disable this option and make sure that no **Event reception tag** is selected.

Selecting the event reception mode

In the **Event notification settings per-user** section, you can select the event alert reception scope per member or third-party channel from the following three modes. The reception mode is selected in the **Event reception tags** column of each member or channel row.

Mode	Description
Receive all alerts	Receive all alerts regardless of the event reception tags
Receive alerts for selected tags	Receive only alerts that have the selected event reception tags. Select at least

Mode	Description
only	one tag
Receive no alerts	Receive no event alerts. The selected tags are preserved without being deleted

- In the **Receive alerts for selected tags only** mode, select at least one tag to receive alerts in the tag selection area.
- To also receive alerts that have no event reception tag specified, select the **Unclassified** virtual tag in the tag selection area. The **Unclassified** tag can be selected or deselected like a normal tag, but the tag itself cannot be edited or deleted.

ⓘ In the **Receive no alerts** mode, the corresponding member or channel receives no event alerts at all. The selected tags are preserved without being deleted, so if you switch back to the **Receive alerts for selected tags only** mode, you can use the previously selected tag settings as is.

Select Notification Language

You can choose the language in which you receive alert notifications.

1. Go to **Alert > Notifications**.
2. In the user list under the **Event notification settings per-user** section, click the ► button at the far left of the user whose language you want to configure.
3. Select the **Notification Language**.

Day and Time-based Notification Settings

You can configure whether a user receives notifications by day of the week and time of day.

1. Go to **Alert > Notifications**.
2. In the user list under the **Event notification settings per-user** section, click the ► button at the far left of the user whose day and time settings you want to configure.

3. Select the days and/or enter the time ranges during which you want to receive alerts.

- Settings can be configured separately for each notification method.

3rd-party Plugins

You can receive alert notifications through external applications.

- **Supported applications:** Slack, Telegram, Teams, Teams Workflows, Jandi, AlertNow, Webhook, Webhook JSON, Opsgenie, PagerDuty, LINE, ilert

ⓘ Internal messengers that are not included in WhaTap Labs' supported scope can be integrated via standard **Webhook** or **Webhook JSON**.

Table | 3rd-party Plugin List

Item	Description
Plugin Name	Name of the plugin
Auth Key	Authentication key
Auth Value	Authentication value
Receive Level	Notification level: Warning, All, or Critical
Repeated Alerts (Escalation)	Re-notification interval from the initial alert time if the issue is not resolved - Applies only to Critical alerts
Smart Message (AI)	AI-powered notifications available when adding the Telegram plugin - Supported languages: Korean, English, Japanese
Event Reception Tag	Whether to receive alerts for specific members or teams in the project
Delete	Delete the 3rd-party plugin

Add a 3rd-party Plugin

1. Go to **Alert > Notifications**.
2. In the **3rd-party Plugins** section, click **Add**.
3. In the **Add 3rd-party Plugin** window, select the desired service tab.
4. Follow the instructions for the selected service (see [How to add 3rd-party plugin notifications](#)).
5. Click **Test Receive** to verify the configuration.
6. When finished, click **Add** or **Register**.

How to Add 3rd-party Plugin Notifications

Slack

- ⓘ In the **Add 3rd-party Plugin** window, click the [**Add To Slack**] button on the **Slack** tab to register **WhaTapNotiBot** to your Slack team and channel.

Telegram

- ⓘ
1. In **Telegram**, search for **@WhaTapApmNotiBot**.
 2. Select the bot and click the [**Start**] button.
 3. In the chat window, enter the `/license` command followed by the access key of the project for which you want to receive alerts.

Telegram

ⓘ Available commands

- `/help` : View the list of commands
- `/license` : Enter the access key of the project to receive alerts




- `/remove` : Enter the project code for which you want to stop receiving alerts
- `/list` : View the list of project codes currently receiving alerts


Teams



1. Go to the **connector** of the channel where you want to receive alerts.



2. Add a **Webhook** and click the [Configure] button.



3. Enter a webhook name and click the [Create] button.



Incoming Webhook

[Send feedback](#)

The Incoming Webhook connector enables external services to notify you about activities that you want to track. To use this connector, you'll need to create certain settings on the other service, which needs to support a webhook that's compatible with the [Office 365 connector format](#).

Fields marked with * are mandatory

To set up an Incoming Webhook, provide a name and select Create. *

Customize the image to associate with the data from this Incoming Webhook.

Upload Image



Default Image

Create

Cancel

4. Copy the generated **Webhook URL**.

Copy the URL below to save it to the clipboard, then select Save. You'll need this URL when you go to the service that you want to send data to your group.

<https://whatap.webhook.office.com/webf>



Url is up-to-date.

Done

Remove


5. Return to the **Add a third-party plug-in** window in the WhaTap Monitoring service, and enter values for **Channel Name** and **Webhook URL** at the bottom of the **Teams** tab.
6. Click the [**Registration**] button.

Teams Workflows



- ⓘ To use workflows, you need an account that can access **Microsoft Teams** and a **Business plan** subscription.

1. In the **Teams** menu, click **More actions (...)** > **Add channel**.



2. In the **Create a channel** window, enter the **Channel name** and **Description**. For **Channel type**, select **Standard**, then click the [Create] button.



Create a channel

Channel name *

WhaTap Event Channel

Description

Help others find the right channel by providing a description

Choose a channel type * ⓘ


Select

- Standard**
Everyone on the team has access.
- Shared**
People or teams in your org or outside your org can have access.
- Private**
Specific people on the team have access.


Cancel Create

ⓘ If you select **Shared** as the channel type, workflows cannot be used. If you select **Private**, workflows can be added but notifications cannot be received.


3. Click **More actions (...)** of the created channel, then select **Workflows**.



4. In the **Workflows** dialog, select **Post to a channel when a webhook request is received** from the workflow list, then click the [Add workflow] button.



5. Check the **Name** and **Connection** status, then click the [Next] button.
6. Review the **Details**, then click the [Add workflow] button.
7. The workflow has been added. Copy the generated URL.



8. Return to the **Add a third-party plug-in** window in the WhaTap Monitoring service, and enter values for **Channel Name** and **Webhook URL** at the bottom of the **Teams Workflows** tab.




9. Click the [**Registration**] button.



Jandi



1. Go to **Jandi Connect settings** in the topic or chat where you want to receive notifications.
2. In the **Connect list**, find **Webhook Receiver (Incoming Webhook)** and click the [**Connect**] button.



3. Click the [**Add Webhook Receiver**] button.
4. Copy the Webhook URL from the **Register Webhook URL** section.



Connect settings

Profile settings

Change the name or profile image that appears for messages from this integration

 Incoming Webhooks

Topic / JANDI selection

Please select a topic or DM to receive integration messages. You change it after saving it.

JANDI

Webhook URL Registration

Webhook URL

Copy the generated Webhook URL and register with the service.


<https://wh.jandi.com/connect-api/webhook/31257583/d3...>
Copy

- Return to the **Add a third-party plug-in** window in the WhaTap Monitoring service, and enter values for **Channel Name** and **Webhook URL** at the bottom of the **Jandi** tab.
- Click the [**Registration**] button.

AlertNow



- Access the **AlertNow** service, then go to **Settings > Integration**.



- On the **Integration** screen, click the [Create Integration] button.
- From the list of integration types, select **WhaTap**.
- On the **Create Integration** screen, configure the **Name** and **Naming Rules**.

Integration

Entire

Integration name *

Enter integration name

Integration type *

WhaTap

Set naming convention *

Define and manage incident naming rules.

naming convention

[[level]][title] {projectName} {oname} {oid}

Preview rule settings

Incident names are limited to 3500 bytes in length.

[Warning]CPU Used > 10 % OTE-SERVER ote-front-330494549

sample data

```
{
  "metricName": "cpu",
  "pcode": 188,
  "level": "Warning",
  "metricValue": "26.888296",
  "oid": "-330494549",
  "title": "CPU Used > 10 %",
  "message": "MEMORY is too high 98%",
  "uid": "4dfe0b7686-c4335f0ef4c1cfa45754",
  "metricThreshold": "18.0",
  "oname": "ote-front",
  "time": 1592887274971,
  "projectName": "OTE-SERVER",
  "status": "on"
}
```

56/1600 byte

- 5. Click the [Confirm] button to create the integration.
- 6. Check the list of created integrations and select the newly created integration.

Integration

Create integration

Enter search term

Last update: 2024-02-21 11:31:17

Integration name	Integration type	Integration key	service name
whatap-test	WhaTap		whatap-test

- 7. On the details screen, click the copy icon on the right side of the URL to copy it.

Integration

whatap-test

Basic Information

Integration type

WhaTap

Integration key

fda527a5140

URL

https://alertnowitgr.opsnow.com/integration/whatap/v1/

constructor

Lisa

Copy to clipboard

- 8. Return to the **Add a third-party plug-in** window in the WhaTap Monitoring service, and enter values for **Channel Name** and **Webhook URL** at the bottom of the **AlertNow** tab.
- 9. Click the [**Registration**] button.

Webhook



1. Register the Webhook that will receive notifications. The alert `title` and `text` are sent to the registered URL via the **POST** method.

```
{
  "title": "Alert Title",
  "text": "Sample warning message"
}
```

2. Enter values for **Webhook Name** and **Webhook URL**.
3. Click the [**Registration**] button.

Webhook JSON



1. Register the Webhook that will receive notifications. Alert data is sent to the registered URL in **JSON format** via the **POST** method. For substitution variables that can be included in the JSON payload, see **Available Variables**.
 - `${something}` represents a substitution value provided by WhaTap. Refer to **Available Variables** for the full list and descriptions.
 - The event status (`status`) is categorized as **event triggered** (`on`) or **event recovered** (`off`).
 - The `level` is categorized as `None` , `Info` , `Warning` , or `Critical` .

```
{
  "pcode": "${pcode}",
  "projectName": "${projectName}",
  "time": "${time}",
  "oid": "${oid}",
  "oname": "${oname}",
  "title": "${title}",
  "message": "${message}",
  "level": "${level}",
  "status": "${status}",
}
```




```
"metricName": "${metricName}",
"metricValue": "${metricValue}",
"metricThreshold": "${metricThreshold}",
"uuid": "${uuid}",
"okind": "${okind}",
"onode": "${onode}"
}
```

2. Enter values for the **Webhook Name** and **Webhook URL** fields.
3. To add a header value, click + **Add** in **Add user header value**, then enter the key and value.
4. Click the [**Registration**] button.


Opsgenie



1. In the **Teams** menu, select the team that will receive notifications.
2. Go to the **Integration** menu and click the [Add Integration] button.



3. In the integration list, search for and select **API**.





4. Set the API name, assign a team, configure the required permissions, and then click the [Save Integration] button.

5. Return to the **Add a third-party plug-in** window in the WhaTap monitoring service, and enter values for **Name** and **API Key** at the bottom of the **Opsgenie** tab.
6. Click the [**Registration**] button.

PagerDuty



1. Sign up for a [PagerDuty](#) account and log in.
2. In PagerDuty, go to **Services > Service Directory**, then click the [+ New Service] button.

PagerDuty

3. Enter the service information (**Name**) and click the [Next] button.

PagerDuty

4. Configure the **Escalation Policy**, then click the [Next] button.



PagerDuty

5. Configure **Alert Grouping** and **Transient Alerts**, then click the [Next] button.

PagerDuty

6. In **Integrations**, select **Event API V2**, then click the [Create Service] button.

PagerDuty

7. On the created service page, copy the **Integration Name** and **Integration Key** values.

PagerDuty

8. Return to the **Add a third-party plug-in** screen in the WhaTap monitoring service and enter the **Integration Name** and **Integration Key** values.
9. Click the [Register] button.

ⓘ If you do not receive messages or the plugin is not registered, check whether the **Integration Name** or **Integration Key** values are correct.

LINE



1. Install the LINE app.
2. Scan the QR code below to add the official **WhaTapLabs** channel as a friend.

Add LINE 3rd-party plugin

3. Use the following commands to link your project with LINE and configure notifications.

Commands

```
/help # Show available commands
/list # View the list of linked projects
/license [ACCESS_KEY] # Add a new project linkage
/remove [PROJECT_CODE] # Remove a linked project
/language [ko|ja|en] # Set notification language (Korean/Japanese/English)
```



- You can find the project access key under **Management > Project Management**.
- Using the `/language` command in LINE applies **only to the current LINE chat room**. It does not change the notification language for other 3rd-party plugins or projects.

Examples

```
/license XXXXXXXXXX # Link a project
/language ko # Set language to Korean
/list # Check linked projects
```



- **Project access key:** Available under **Management > Project Management** in the WhaTap project.
- **Language settings:** The `/language` command applies only within the LINE chat room and does not affect other notification channels.

ilert



1. In the ilert console, click [+ Create alert source] at the top right of the **Alert sources** page.
2. In the **Integration type** step, enter `WhaTap` in the **Filter integrations** search field at the bottom to search.
3. From the search results, select the **WhaTap** tile and click the [Next] button.
4. Enter the alert source name and assign a team, then click the [Next] button.
 - You can create a new escalation policy or select an existing one.
5. Select the alert grouping method and click the [Continue setup] button.
 - The initially selected **Do not group alerts** option can be changed later.
6. On the next screen, you can configure additional settings such as **alert templates and priorities**. For now, click the [Finish setup] button to complete the setup.
7. When the setup is complete, an **API key** or **Webhook URL** will be generated on the completion screen.



8. Copy the generated **Webhook URL** from the **Webhook JSON** tab and paste it into the **Webhook URL** input field below.
9. Click the [Register] button.

Preventing Alert Storms

If a large number of alerts occur within a short period, alert notifications are temporarily paused for a specified duration. For example, if 20 events occur within 5 minutes, alert notifications will be suspended for the next 5 minutes. Once the configured **Pause Duration** has elapsed, the alert-storm prevention feature is automatically released.

1. Go to **Alert > Notifications > the Avoids a large number of alerts** section.
2. Click the toggle button to **Enable** or **Disable** the feature.
 -  **Enabled:** Temporarily pause alerts when an alert storm is detected
 -  **Disabled:** Resume alert notifications
3. Configure the **Detection Time**, **Detection Count**, and **Pause Duration**.
 - If the number of events exceeds the **Detection Count** within the **Detection Time**, alert notifications are paused for the configured **Pause Duration**.



When alerts are blocked due to an alert storm, a notification banner is displayed. To immediately release the alert-storm block, click the **Stop Alert Storm Prevention** button in the banner.

Event Notification Format

Common Alert Template

The common template can be applied in most cases. It allows you to receive alerts in a consistent format across various environments.

- Product Type: Application, Database, Kubernetes
- Event Type: Application Alerts, Database Alerts, Metrics Alerts
- Event Channel: SMS, Mobile, 3rd-party plugins, Plugins

Event Title

Event title format

```
[Level][Platform][ProjectName][ApplicationName][EventTitle]
```

Event title example

```
[Info][JAVA][Application Project][TC-0-1-8081][CRITICAL_HIGH_MEMORY]
```

❗ If ApplicationName is not set, it is omitted from the event title.

- Platform is displayed as one of the following.
 - JAVA
 - NODEJS
 - PYTHON
 - PHP
 - DOTNET
 - GO
 - POSTGRESQL



- ORACLE
- MYSQL
- MSSQL
- BSM_JAVA
- CLOUDWATCH
- TIBERO
- KUBERNETES
- KUBE_NS
- URLCHECK
- URLCHECK_ADMIN
- CUBRID
- ALTIBASE
- CLUSTER
- REDIS
- MONGODB
- VR
- RUM

Event Message

This section describes the information that can be included in the event message. If `optional` is set to `false`, the item is always included in the message. If `optional` is set to `true`, the item is displayed only when the data is available.

Event message example

```
Project Name : Application Project
Project Code : 3
Agent Name : TC-0-1-8081
Message : RECOVERED: Memory is too high. less than 10%
```

Event Time : 2022-04-12 18:53:24 +0900
 Event OFF Time : 2022-04-12 18:53:24 +0900
 Alert Type : APPLICATION_MEMORY
 Metric Name : memory
 Metric Value : 20
 Metric Threshold : 10
 Stateful : true

Table | Event Message Components

En	Ko	Supported Alert Types	Description
Project Name	Project name	All	-
Project Code	Project code	All	-
Agent Name	Agent name	All (optional)	oname
Message	Event message	All	-
Alert Type	Alert type	All	See the AlertType table below
Event Time	Event occurrence time	All	In 2022-04-13 10:40:49 +0900, +0900 indicates the GMT offset
Event Off Time	Event recovery time	All (optional)	In 2022-04-13 10:40:49 +0900, +0900 indicates the GMT offset
Metric Name	Metric name	All (optional)	Name of the metric used to evaluate the event condition
Metric Value	Metric value	All (optional)	Included when the metric value exceeds the threshold
Metric Threshold	Metric threshold	All	Threshold value used to determine the event

En	Ko	Supported Alert Types	Description
		(optional)	condition
Stateful	Recovered-event notification	All (optional)	<code>true</code> if recovered-event notification is enabled; otherwise <code>false</code>
Event Rule	Event trigger condition	Metrics alerts	-
Target Filter	Event target filter	Metrics alerts	Evaluates event conditions only for metrics collected from specific targets
Repeat Count	Event repeat count	Metrics alerts	The event is triggered when the condition is met the specified number of times within the repeat duration
Repeat Duration	Event repeat duration	Metrics alerts	Time window used to evaluate the repeat count
Receiver	Receiver	Metrics alerts	-
Query	MXQL query	Composite metrics alerts	-
Rule	Event trigger condition	Composite metrics alerts	-
Query Period	Query period	Composite metrics alerts	-
Query Interval	Query interval	Composite	-

En	Ko	Supported Alert Types	Description
		metrics alerts	
Silent Time	Silent time	Composite metrics alerts	-
Query	URL	Exception alerts	URL of the request that caused the exception
TXID	Transaction ID	Exception alerts	-
Class	Error class name	Exception alerts	-
Log Message	Log message	Server – file log alerts	-
Log File	Log file path	Server – file log alerts	-
IP	IP address	All server alerts	-
CPU	CPU usage	All server alerts	Snapshot at the time the event occurred
CPU_load1	CPU load (1 min)	All server alerts	Snapshot at the time the event occurred
CPU_loadPerCore	CPU load per core	All server	Snapshot at the time the event occurred

En	Ko	Supported Alert Types	Description
		alerts	
Memory	Memory usage	All server alerts	Snapshot at the time the event occurred
Swap	Swap usage	All server alerts	Snapshot at the time the event occurred
Disk Name Used Percent Free Size IO Percent	Disk performance	All server alerts	Snapshot at the time the event occurred
Name Bps Pps	Network traffic performance	All server alerts	Snapshot at the time the event occurred
Message Time Name	Acknowledgement information	All server alerts	-

ⓘ This section displays as much information as possible that can be provided for the event.

- **AlertType** is displayed as one of the following values.

AlertType	Description
APPLICATION_CPU	Application CPU alert
APPLICATION_MEMORY	Application memory alert
APPLICATION_DISK	Application disk alert
APPLICATION_ACTIVE_TRANSACTION	Application active transaction alert
APPLICATION_ERROR_TRANSACTION	Application error transaction alert



AlertType	Description
APPLICATION_SLOW_TRANSACTION	Application slow transaction (response time) alert
METRICS	Metrics alert
COMPOSITE_METRICS	Composite metrics alert
ANOMALY	Anomaly detection alert
LOG_REALTIME	Real-time log alert
COMPOSITE_LOG	Composite log alert
SERVER_REBOOT	Server reboot alert
SERVER_NO_DATA	Server no-data alert
SERVER_PORT	Server port alert
SERVER_NETWORK_IOPS	Server network IOPS alert
SERVER_NETWORK_BPS	Server network BPS alert
SERVER_DISK_IO	Server disk I/O alert
SERVER_DISK_QUOTA	Server disk usage alert
SERVER_DISK_INODE	Server inode alert
SERVER_CPU	Server CPU alert
SERVER_MEMORY	Server memory alert
SERVER_CPU_STEAL	Server CPU steal alert



AlertType	Description
SERVER_MEMORY_SWAP	Server swap memory alert
SERVER_LOG_FILE	Server log file alert
SERVER_WINDOW_EVENT	Server Windows event alert
SERVER_OFF	Server alert OFF notification
SERVER_ACKNOWLEDGE	Server acknowledge notification
SERVER_PROCESS_COUNT	Server process count alert
SERVER_PROCESS_CPU	Server process CPU alert
SERVER_PROCESS_MEMORY	Server process memory alert
SERVER_PROCESS_OFF	Server process alert OFF notification
AGENT_ACTIVE	Agent activated alert
AGENT_INACTIVE	Agent deactivated alert
AGENT_REACTIVATED	Agent reactivated alert
URL	URL alert
TOO_MANY_EVENT	Too many events alert
CLOUD_WATCH	CloudWatch alert
EXCEPTION	Exception alert

Event History And Monitoring

Home > Select Project >  **Alert** > **Event History**

Event History allows you to view events that have occurred in the project and analyze their detailed information. It supports Elasticsearch DSL queries, enabling you to find events that match specific conditions quickly and accurately.

Viewing Event History

Click the **event title** in the search results to move to a screen where you can analyze detailed data for the time period when the event occurred. This allows you to comprehensively understand the situation before and after the event.

1. Go to **Alert > Event History**.
2. In the **Event History** screen, use **Time Selection** to set the period for viewing event records.
3. Enter a query in the filter input field and click the [Magnifying Glass] button.
 - If no query is entered, all events are displayed.

Column Settings

You can choose which columns to display or hide.

1. In the **Event History** screen, click the [ Column Settings] button at the top left.
2. In the **Column Settings** window, select the columns you want to show or hide.
 - You can change the column order in **Displayed Items**.
 - You can search for specific columns using the **Column Name Search** field.

Column Item	Description
Event Name	The title of the event

Column Item	Description
Message	The event message or snapshot information
Event Status	The current status of the event
Agent Name	The name (Oname) of the agent where the event occurred - Displayed as blank if the event is not associated with a specific agent
Agent Group	The agent type (OkindName) where the event occurred - Displayed as blank if the event is not associated with a specific agent
Agent Node	The server (OnodeName) of the agent where the event occurred - Displayed as blank if the event is not associated with a specific agent
Event Occurred Time	The time when the event occurred
Event Resolved Time	The time when the event was resolved - Displayed as - for events without a resolved state

- After completing the settings, click the [Confirm] button to apply the column configuration.

↓ CSV Download

You can download the query results as a CSV file.

- Go to **Alert > Event History** and click the [CSV] button.
- The list of events that match the current filter conditions is saved as a CSV file.

View Only Events in Progress

You can view only events that are currently in progress.

- Go to **Alert > Event History** and click the [View Only Events in Progress] button.
- The following condition is automatically added to the filter query and applied to the search.

Stateful: **true** and Status : **"ON"**

Event Search

- **Precise Search:** Specify exact conditions to retrieve only the events you need
- **Composite Conditions:** Search across multiple fields at once to handle complex conditions in a single query
- **Flexible Patterns:** Supports various search methods such as wildcards and partial matches
- **High Performance:** Fast searches powered by indexed data

Searchable Fields

Field names and values are not case sensitive.

- **Field Name Tag:** Event types that support the field
- **Field Type:** The data type of the field
 - #: Number
 - T: String
 - : Boolean
 - : Date

Event Information Fields

Field Name	Type	Description	Example Value
Title	String	Event title	"Database Connection Error"
OffTitle	String	Event recovery title	"RECOVERED: Database Connection Error"
Message	String	Event message	"Connection timeout occurred"
OffMessage	String	Event recovery message	"RECOVERED: Connection timeout occurred"

Field Name	Type	Description	Example Value
Level	String	Current event level	Critical, Warning, Info
OriginLevel	String	Original event level	Critical, Warning, Info
Status	String	Event status	ON, OFF, CANCEL, ACKNOWLEDGE, MAINTENANCE, DISABLED
ActCount	Number	Number of occurrences	2
MetricName	String	Metric name	"memory"
MetricValue	String	Metric value	"85.5"
OffValue	String	Recovery value	"72.8"
MetricThreshold	String	Threshold value	"80"
alertType	String	Event type	"BASIC", "METRICS", "TRANSACTION", etc.
alertId	String	Unique identifier of the event rule	"zf3uojer0fv4v7"

Type-specific Supported Fields

ⓘ These fields are supported only for specific event types. If the event type does not match, the value may be empty.

Field Name	Type	Event Type	Description	Example Value
eventRule	String	Basic, Metrics	Event trigger rule	"memory ≥ 80"
field	String	Real-	Log search	"content"

Field Name	Type	Event Type	Description	Example Value
		time Log	key	
keyword	String	Real-time Log	Log search value	"Error"
logCategory	String	Real-time Log	Log category	"AppLog"
logContent	String	Real-time Log	Log content	"00:00000:00009:2025/04/06 23:03:52.55 server Error: 1601, Severity: 17, State: 3\n00:00000:00009:2025/04/06 23:03:52.55 server There are not enough 'user connections' available to start a new process."

Target Fields

Field Name	Type	Description	Example Value
Oid	Number	Agent unique identifier	-98765432
Oname	String	Agent name	"web-server-01"
Okind	Number	Agent type unique identifier	867318026
OkindName	String	Agent type name	"web-server"
Onode	Number	Agent node unique identifier	334634079
OnodeName	String	Agent node name	"production-node-1"

Unique Event Fields

Field Name	Type	Description	Example Value
Id	Number	Unique identifier of the event	5768121
UUID	String	Unique event identifier	"550e8400-e29b-41d4-a716"

State / Flag Fields

Field Name	Type	Description	Example Value
Stateful	Boolean	Whether the event is state based	true, false
Disabled	Boolean	Whether the event is disabled	true, false
Escalation	Boolean	Whether escalation is applied	true, false
SystemEvent	Boolean	Whether it is a system event	true, false

Query Syntax by Search Type

ⓘ For the basic search syntax, refer to the [WhaTap log search query syntax](#) documentation.

1. Keyword Search

Events that contain the string "Database Connection"

```
title: "Database Connection"
```

2. Multiple Value Search

Retrieves events that include any one of the specified values.

```
title: Database Connection
```

Events whose current event level is Info or Warning

```
level: info warning
```

3. Pattern Search

You can search using patterns with wildcards.

Find events that occurred on specific agents using a name pattern

```
oname: web-*-prod
```

タイトルが Connection で終わるイベントを検索

```
title: *Connection
```

4. Complex Condition Search

You can combine multiple conditions to perform precise searches.

Events with a Critical level and a title containing Database

```
level: critical and title: Database
```

Warning or Critical events from a specific agent

```
oname: web-server-01 and level: Warning Critical
```

5. OR Condition Search

Retrieves events that satisfy **at least one** of the specified conditions.

Events where the title or message contains a specific keyword

```
title: Connection or message: Connection
```


6. Exclusion Conditions

You can exclude specific conditions from the search.

All events except those with an Info level

```
not level: info
```

Real-time Notifications

The **Alert status** feature displays pop-up notifications when events configured in **Event Configuration** occur. You can receive alerts immediately when an event is triggered, allowing you to quickly recognize issues and respond. You can also review notifications that have already occurred or were missed, helping you understand and resolve situations where real-time response was difficult.


Viewing Real-time Notifications

Real-time notifications appear as pop-ups in the upper-right corner of the screen. You can check the time, title, and message of each notification.

- To close an individual notification, click the  button in the upper-right corner of the message.
- To close all pop-up notifications displayed on the screen, click **Close All Notifications**.

Viewing Notification History

1. Click the  icon at the top of the screen after selecting a project.
2. The **Alert status** window appears.
 - You can view up to 100 notifications from the past 24 hours for the current project.
 - Selecting an individual event item navigates you to the related menu.
 - In the time-based chart within the **Alert status** window, selecting a specific time range takes you to the corresponding event items.



- To view the full notification history, click **View All Notification History**. This navigates to **Alert > Event History**.

ⓘ If the **Alert status** list is empty, the **View All Notification History** button is displayed. Disable the **Receive notifications without reception tags** option.

Resolving Events in Bulk

You can resolve all ongoing events at once from the **Alert status** list. The bulk resolution process may take approximately 30 seconds to 1 minute to complete.

- When using the **Add Acknowledge** feature, a large number of recovery (resolved) notifications may be generated for currently ongoing events.
- Even after using the **Add Acknowledge** feature, events may occur again if the event rules remain enabled and the data still exceeds the threshold.
- Recovery notifications can be disabled by editing the event in **Alert > Event Configuration**. Disable **Send additional notification when the event is resolved** or **Recovered notification**.

How to use:

1. Click the  icon at the top of the screen after selecting a project.
2. When the **Real-time Notifications** window appears, click the  icon in the upper-right corner.
3. When the message `Resolve all ongoing events in the real-time notification list.` appears, click the **Resolve** button.

Configuring Real-time Notifications

1. Click the  icon at the top of the screen after selecting a project.
2. When the **Alert status** window appears, click the  icon in the upper-right corner.
3. Configure each option in the real-time notification settings window.
 - **Popup notifications:** Displays notifications as pop-ups even when the **Alert status** window is closed.
 - **Show resolved events:** Allows resolved events to be displayed in the **Alert status** list.
 - **Preview notification sound:** Lets you preview browser notification sounds. Click the  button next to **Critical** or **Warning**. If you cannot hear the sound, check your browser settings.

System Events

WhaTap Monitoring provides common system events that are automatically detected and generated across the platform. This document explains the types of events generated by the system, their trigger conditions, and recommended response actions.

EVENT_BLOCKED_BY_SYSTEM

If an excessive number of events occurs within a short period of time in a project, the WhaTap system detects this situation and blocks further events. When events are blocked, the system delivers the following message to the user. If you encounter this event, review and adjust your event configuration settings.

 **Critical**



Event Message

Excessive events detected, event blocked for stability. name: server agent

Project No	8
Project Name	new server
Target	server agent
Event Title	EVENT_BLOCKED_BY_SYSTEM
Event Time	(GMT+0900) 2024-11-18 17:58:58

Log In to WhaTap

Events are blocked as a result of excessive events in the server agent of new server project. If no events occur in server agent for a period of time, it is automatically unblocked.

Blocking Criteria

Events are blocked based on the following criteria:

1. Agent-level blocking

If **10 events** occur within **1 minute** from a single agent, all subsequent events from that agent are blocked.

2. Project-wide blocking triggered by multiple agents

If **3 or more agents** are blocked within **1 minute** according to condition 1, events from the entire project are blocked.

3. Project-level blocking

If **100 events** occur within **1 minute** in a single project, all events from that project are blocked.

How to Release the Block

When events are blocked by the system, no events must occur from the affected agent or project for **1 minute**.

- Disable all configured events from **Alert > Event Configuration**. **Review your event settings and adjust thresholds to prevent excessive event generation.**
- Before configuring events, use the **Simulation** feature to estimate how many events may occur.

실험실

사용자에게 새로운 기능 또는 실험적인 기능을 제공합니다. 사이드 메뉴의  실험실 또는 **사이트맵**을 통해 접근할 수 있습니다. 제품에 따라 제공하는 기능이 다를 수 있습니다.

✔ 사이트맵으로 실험실 메뉴 접근하기

사이드 메뉴에서  실험실이 보이지 않는다면  **사이트맵**으로 접근할 수 있습니다.

1. 전체 화면 왼쪽 아래에  **사이트맵** 아이콘을 클릭하세요.
2. **사이트맵** 창의 오른쪽 아래에 **실험실** 섹션에서 원하는 메뉴를 선택하세요.

⚠ 주의

실험실의 기능은 현재 개발 중인 베타 버전으로 예기치 않은 오류가 발생할 수 있습니다.

해당 기능 사용 시 주의가 필요합니다. 중요한 데이터나 운영 환경에서 사용을 권장하지 않습니다. 피드백이나 문제점이 발생하면 지원팀(support@whatap.io)으로 문의하시기 바랍니다.

MXQL 데이터 조회

MXQL은 와탭의 성능 데이터(메트릭스)를 유연하게 조회하기 위한 쿼리 언어입니다. 하나의 프로젝트에 포함된 여러 에이전트에서 수집한 메트릭스들을 종합적으로 조회하고 활용하기 위해서 사용합니다. **MXQL**에 대한 자세한 내용은 [다음 문서](#)를 참고하세요.


알림 메시지

프로젝트 멤버에게 메일 또는 서드 파티 플러그인을 통해 전달하는 경고 알림 메시지를 수정할 수 있는 사용자 정의 기능을 제공합니다. 자세한 내용은 [다음 문서](#)를 참고하세요.

메트릭스 익스플로러

홈 화면 > 프로젝트 선택 > 실험실 > 메트릭스 익스플로러 Beta

메트릭스 익스플로러 Beta 메뉴는 모니터링 대상으로부터 수집한 데이터를 시계열 차트로 시각화하여 제공합니다. 이를 통해 사용자는 시스템 성능을 효과적으로 모니터링하고 이상 징후를 빠르게 식별하여 문제 해결에 필요한 인사이트를 얻을 수 있습니다.



- 다양한 지표를 시계열 차트를 통해서 시각적으로 확인할 수 있으며, 장애 발생 시점의 여러 지표를 한번에 분석할 수 있습니다.
- 특정 시간 범위를 선택하여 과거 데이터를 조회하고, 해당 기간 동안의 성능을 분석할 수 있습니다.
- 특정 시간대나 조건에 맞는 데이터를 필터링하여 보다 정확한 분석을 수행할 수 있습니다.
- 사용자는 원하는 차트를 화면에 배치하고 레이아웃을 자유롭게 조정할 수 있습니다.
- 사용자가 원하는 지표들을 배치하고 프리셋으로 저장하여 필요 시 불러올 수 있습니다.

차트 위젯 추가하기

대시보드에 지표를 위젯으로 추가하려면 화면 왼쪽 상단의 **차트 추가** 버튼을 선택하세요.

조회 시간 설정하기

시간에서 지표를 조회할 날짜와 시간을 설정하세요.

대상 선택하기

프로젝트에 포함된 에이전트 중 원하는 대상을 선택하세요. 개별 에이전트를 선택하거나 에이전트 설정을 통해 분류한 그룹 단위로 선택할 수 있습니다.

대상

전체

에이전트

종류별

서버별

대상 검색

Q

전체 선택

0/9

전체 접기

↺

▼ Project

▼ demo-okind-0

demo-8104

demo-8102

demo-8100

▼ demo-okind-1

demo-8103

demo-8105


demo-8101

- **에이전트:** 개별 에이전트를 선택하거나 모두 선택할 수 있습니다.
- **종류별:** 에이전트 설정에서 `whatap.okind` 로 분류된 그룹 단위로 에이전트를 모니터링합니다.
- **서버별:** 에이전트 설정에서 `whatap.onode` 로 분류된 그룹 단위로 에이전트를 모니터링합니다.

- ❗ • 설정한 시간과 날짜에 따라 선택할 수 있는 대상이 변경될 수 있습니다.
- 에이전트 목록에서 마우스 클릭해 대상을 선택하거나 해제할 수 있습니다. **Shift** 키를 이용해 여러 개의 에이전트를 한번에 선택할 수도 있습니다.
- **전체 선택** 또는 **전체 해제** 버튼을 선택하면 관련 하위 대상들을 한번에 선택 또는 해제할 수 있습니다.
- 원하는 에이전트가 목록에 표시되지 않는다면 **새로고침** 버튼을 선택하세요.
- **전체 접기** 또는 **전체 펼치기** 버튼을 선택해 그룹 단위로 목록을 감추거나 펼칠 수 있습니다.

지표 선택하기

모니터링할 대상 에이전트를 선택한 후 **지표** 패널에 선택할 수 있는 목록이 표시됩니다. 목록에서 원하는 지표를 선택하세요.



선택한 지표에 해당하는 차트 위젯이 오른쪽의 대시보드에 배치됩니다. 다른 지표를 추가로 선택해 차트 위젯을 추가할 수 있습니다. 차트 영역을 확대해 넓게 보려면 **시간** 옵션 오른쪽의 **×** 버튼을 선택하세요. **대상**과 **지표** 패널이 사라지고 차트 영역만 화면에 표시됩니다.

- ❗ • **지표 검색** 입력란에 문자열을 입력하면 원하는 지표를 빠르게 검색할 수 있습니다.
- 지표의 카테고리 및 필드 이름을 기준으로 검색할 수 있습니다. 카테고리 및 필드에 대한 자세한 내용은 [다음 문서](#)를 참조하세요.
- **전체 접기** 또는 **전체 펼치기** 버튼을 선택해 카테고리 단위로 지표 목록을 감추거나 펼칠 수 있습니다.



- 원하는 지표 항목이 목록에 표시되지 않는다면 **새로고침** 버튼을 선택하세요.

NUMA Node 지표


지표 목록에 **NUMA Node** 카테고리를 추가했습니다. NUMA 노드별 CPU 사용률, 메모리 사용량, 노드 간 메모리 접근 비율 등을 조회할 수 있습니다. 노드별 지표를 비교하면 특정 노드에 부하가 쏠렸는지, 원격 노드 메모리 접근이 과도한지 확인할 수 있습니다.

차트 범례는 다음과 같이 표시합니다.

Item	Description
NUMA Node 카테고리	Node 0 형태로 노드 번호 표시
CPU Core 카테고리	Core cpu1 (Node 0) 형태로 코어 이름과 소속 노드를 함께 표시

CPU Core 범례의 노드 표기는 NUMA 연결 정보를 수집한 Linux 환경에서만 나타납니다. 정보가 없는 코어와 Windows·Unix 환경에서는 Core cpu1 형태를 유지합니다.


인터벌 설정하기



차트마다 인터벌을 설정할 수 있어 데이터의 집계 단위를 조정하여 다양한 시간 간격의 데이터를 시각화할 수 있습니다. **지표** 패널의 오른쪽 상단에 ⌚ 버튼을 선택하면 인터벌 옵션을 선택할 수 있습니다.

지표마다 기본 설정된 인터벌 시간이 다르며 조회 시간 범위에 따라 인터벌 시간이 변경됩니다.


예를 들어, **5분** 인터벌은 5분 단위의 평균 값을 표시하고, **1시간** 인터벌은 1시간 단위의 평균 값을 표시합니다.




- ✓ 인터벌 옵션을 조정하여 짧은 시간 간격의 세부 데이터부터 긴 시간 간격의 전체 추세까지 다양한 관점에서 데이터를 분석할 수 있습니다.

조회 시간 변경하기


개별 위젯의 차트 조회 시간 범위를 변경하려면 위젯 상단의 < 또는 > 버튼을 선택하세요. 시간 변경 범위는 시간 옵션에 설정된 시간 기준과 같습니다. 예를 들어, 시간 옵션에서 1일로 설정되어 있을 때, < 또는 > 버튼을 선택하면 시간을 1일 단위로 변경할 수 있습니다.



또는 차트의 특정 영역을 드래그하면, 드래그한 범위만큼 시간을 변경해 조회할 수도 있습니다.



- ✓ 화면 상단에 **차트 시간 일괄 적용** 옵션을 활성화한 상태에서 개별 위젯의 시간을 변경하면 모든 위젯의 시간을 변경할 수 있습니다.




- 조회 시간 범위를 변경한 다음 최근 시간으로 다시 변경하려면  버튼을 선택하세요.

만약 화면 상단에 **차트 시간 일괄 적용** 옵션을 활성화한 상태라면  **시간 일괄 초기화** 버튼이 표시됩니다. 이 버튼을 선택하면 모든 위젯의 시간을 초기화할 수 있습니다.

- ❗ **일괄 변경** 기능을 이용해 전체 위젯의 시간을 변경하고, 모니터링 대상과 인터벌도 수정할 수 있습니다. 자세한 내용은 [다음 문서](#)를 참조하세요.

차트 위젯 상세 안내

차트 위젯에 표시되는 각 항목에 대한 자세한 내용을 확인하세요.



- **1 차트 위젯 이름:** 차트 위젯 상단에서 선택한 지표 이름, 지표 그룹명, 대상을 표시합니다.
 - **a:** 마우스 포인터를 위젯 이름 위로 가져가면 모니터링 대상을 확인할 수 있습니다.
- **2 위젯 옵션:** 차트 위젯의 옵션 메뉴입니다. 자세한 내용은 [다음 문서](#)를 참조하세요. 인터벌 옵션에 대한 자세한 내용은 [다음 문서](#)를 참조하세요.
- **3 시간:** 차트의 데이터를 조회하는 시간 범위를 표시합니다. < 또는 > 버튼을 선택해 조회 시간 범위를 조정할 수 있습니다.
- **4 데이터:** 모니터링 대상 중 조회한 데이터의 최댓값(Max)과 데이터 병합 방식(Value Type)을 확인할 수 있습니다.

- **5 범례:** 모니터링 대상을 범례로 표시합니다. 범례 표시에 대한 자세한 내용은 [다음 문서](#)를 참조하세요.
- **6 차트:** 조회한 데이터를 시계열 그래프로 표시합니다. 마우스 포인터를 그래프 위로 가져가면 툴팁을 통해 상세 정보를 확인할 수 있습니다. 차트의 특정 영역을 드래그하면, 드래그한 범위만큼 시간을 변경해 조회할 수 있습니다.
- **7 크기 조절:** 차트 위젯의 가운데 하단 영역을 드래그하면 위젯의 크기를 조절할 수 있습니다.

차트 위젯 옵션 이용하기

차트 위젯의 오른쪽 상단에 위치한 옵션 버튼을 통해 다양한 기능을 이용할 수 있습니다.

▽ 대상 필터링하기

화면에 추가한 차트 위젯에 모니터링 대상을 추가하거나 제외할 수 있습니다.

1. ▾ 버튼을 선택하면 모니터링 대상을 선택할 수 있는 팝업 메뉴가 나타납니다.




2. 화면의 안내에 따라 원하는 대상을 선택하세요.
3. 적용 버튼을 선택하세요.

상세 조회하기

차트의 데이터를 더 넓은 화면에서 자세한 정보를 확인하려면  버튼을 선택하세요. 상세 조회 창이 나타납니다. 개별 대상에 대한 차트와 평균 및 최댓값을 조회할 수 있습니다.



- **1일전:** 하루 전의 데이터와 비교해 확인할 수 있습니다.
- **1주전:** 한 주 전의 데이터와 비교해 확인할 수 있습니다.
- **↓ CSV:** 조회한 데이터를 CSV 형식의 파일로 다운로드할 수 있습니다.
- 그래프 차트 하단의 아이콘을 드래그하면 차트의 크기를 조절할 수 있습니다.




🗑 차트 위젯 삭제하기


화면에 배치한 차트 위젯을 삭제하려면 🗑 버튼을 선택하세요.

범례 활용하기


범례를 통해서 차트에 표시되는 모니터링 대상을 확인하고, 원하는 기준으로 정렬하거나 필터링할 수 있습니다. 화면 상단의 **위젯 범례 표시** 옵션을 켜거나 끄면 위젯의 차트 오른쪽에 범례를 표시하거나 숨길 수 있습니다. 옵션을 끈 상태에서는 **범례**를 클릭해 모니터링 대상을 확인하고, 선택할 수 있습니다.




- 전체 선택/전체 해제 버튼을 선택해 모니터링 대상을 선택하거나 해제할 수 있습니다.
- 범례 목록에서는 모니터링 대상을 이름순, 최댓값, 평균값 기준으로 정렬할 수 있습니다.



- 범례 목록에서 각 항목에 마우스를 오버하면 해당 모니터링 대상에 대한 그래프를 강조해서 확인할 수 있고, Max 값이 기록된 특정 시간을 톨팁으로 표시합니다. 또는 범례 목록의 대상을 선택해 필터링할 수도 있습니다.



범례 색상 변경하기




모니터링 대상에 해당하는 그래프의 색상을 변경할 수 있습니다. 범례 목록에서 색상을 변경하려는 대상의 왼쪽에 네모 상자를 선택하세요. 원하는 색상을 선택한 다음 **적용** 버튼을 클릭하세요.

❗ 설정한 색상은 프리셋 기능을 통해서 저장되지 않습니다. 이 기능은 향후 업데이트를 통해 제공할 예정입니다.

추가 옵션 이용하기

차트 위젯 오른쪽 상단에  버튼을 선택하면 다음의 추가 옵션을 이용할 수 있습니다.



• 설정

- **자동 Y축 사용:** 조회한 데이터 범위 내의 Y축의 크기를 자동 조정할 수 있습니다.
- **최대값 표시:** 차트의 그래프 영역에 최댓값을 표시할 수 있습니다.
- **이미지 다운로드:** 차트를 이미지 형식의 파일로 다운로드할 수 있습니다.
- **이미지 복사:** 차트를 이미지 형식으로 클립보드에 복사할 수 있습니다.
- **CSV 다운로드:** 차트에서 조회한 데이터를 CSV 형식의 파일로 다운로드할 수 있습니다.
- **CSV 복사:** 차트에서 조회한 데이터를 CSV 형식으로 클립보드에 복사할 수 있습니다.


프리셋 설정하기

대시보드에서 사용자가 설정한 위젯의 설정과 레이아웃 상태를 저장하고 불러올 수 있습니다. 위젯의 크기를 조절하고, 원하는 위치에 배치해 새로운 프리셋을 만들 수 있습니다.

❗ **프로젝트 수정** 권한이 있는 멤버만 프리셋을 저장하거나 수정할 수 있습니다. 멤버 권한에 대한 자세한 내용은 [다음 문서](#)를 참조하세요.

새로운 프리셋 만들기

1. 대시보드에서 원하는 형식으로 위젯을 배치해 보세요. 크기를 조절하고 자주 확인하는 위젯만 배치할 수도 있습니다.
2. 화면 오른쪽 위에  버튼을 선택하세요.
3. 새로운 프리셋 이름을 입력하세요.



4. **저장** 버튼을 선택하세요.

프리셋 목록에서 새로 저장한 프리셋을 확인할 수 있습니다.

- ❗
- 새로 만든 프리셋에 변경 사항이 생겼다면 다시 프리셋을 저장해야 합니다.  버튼을 선택한 다음 같은 이름으로 프리셋을 저장하세요. 기존의 프리셋에 변경 사항을 덮어쓰기합니다.
 - 대시보드의 변경 사항을 저장하지 않고 다른 메뉴로 이동하면 변경 사항은 저장되지 않습니다.
 - 프리셋은 프로젝트 단위로 저장되어 다른 사용자와 공유할 수 있습니다.
 - 프리셋에 저장되는 대상은 대시보드의 레이아웃과 대시보드에 배치한 위젯, 위젯에서 수정한 옵션(**대상**, **인터벌**)입니다. 화면 상단의 **시간** 옵션은 저장되지 않습니다.
 - 프리셋을 불러온 다음 위젯의 조회 시간을 변경하려면 **일괄 변경** 기능을 이용하세요. 자세한 내용은 [다음 문서](#)를 참조하세요.


프리셋 삭제하기

사용하지 않는 프리셋이 있다면 프리셋 목록에서 삭제할 수 있습니다. 프리셋 목록에서 삭제하려는 항목의 오른쪽에 ✕ 버튼을 선택하세요.

대시보드 위젯 편집하기

대시보드에 배치한 위젯을 사용자가 원하는 위치에 배치하거나 크기를 조절할 수 있습니다.


위젯 이동하기



위젯의 상단 부분을 마우스 포인터로 이동하면 + 모양으로 변경됩니다. 이때 마우스로 클릭한 상태로 원하는 위치로 드래그하여 위젯을 이동할 수 있습니다.

위젯 크기 조절하기

대시보드 옵션의 레이아웃이 사용자 정의로 설정되어 있으면 위젯의 크기를 조절할 수 있습니다.



위젯의 오른쪽 아래에  요소를 마우스로 클릭한 상태에서 원하는 크기로 드래그하세요.

ⓘ 레이아웃이 사용자 정의로 설정되어 있지 않다면 위젯의 높이만 조절할 수 있습니다. 위젯 하단 가운데 영역을 마우스로 클릭한 상태로 크기를 조절하세요.

대시보드 옵션 이용하기

화면 상단에 위치한 옵션 메뉴를 이용해 차트 위젯에 적용된 다양한 설정을 한번에 적용할 수 있습니다. 또한 대시보드의 레이아웃을 원하는 형태로 설정할 수도 있습니다.



차트 위젯 옵션 일괄 변경하기

화면에 배치된 차트 위젯에 세부 옵션을 한번에 적용할 수 있는 기능을 제공합니다. 화면 왼쪽 상단에 **일괄 변경** 버튼을 선택하세요. **일괄 변경** 창이 나타나면 원하는 옵션을 설정한 다음 **적용** 버튼을 선택하세요.

일괄 변경

×

시간

<

2025/01/21 00:00 ~ 2025/01/21 23:59

23시간

□

>

대상

전체

대상 검색

Q

전체 해제

2/3

▼ 전체 접기

↻

▼ Project

qa-node-demo

✓

qa-php-demo

✓

인터벌

변경 안함

5분

1시간

적용 위젯

전체 선택

☐ User+Sys (%)

2025/01/21, 00:00:00 ~ 23:59:59

sm-oid

5분

☐ Used (%)

2025/01/21, 00:00:00 ~ 23:59:59

sm-oid

5분

☐ Rx (bps)

2025/01/21, 00:00:00 ~ 23:59:59

sm-oid

5분

☐ Total Thread#

2025/01/21, 00:00:00 ~ 23:59:59

sm-oid

5분

취소

적용

- **시간:** 차트의 데이터 조회 범위를 설정할 수 있습니다.
- **대상:** 모니터링 대상을 선택하세요. 개별 에이전트를 선택하거나 개별 에이전트를 선택하거나 에이전트 설정을 통해 분류한 그룹 단위로 선택할 수 있습니다.
 - **에이전트:** 개별 에이전트를 선택하거나 모두 선택할 수 있습니다.
 - **종류별:** 에이전트 설정에서 `whatap.okind` 로 분류된 그룹 단위로 에이전트를 모니터링합니다.
 - **서버별:** 에이전트 설정에서 `whatap.onode` 로 분류된 그룹 단위로 에이전트를 모니터링합니다.

607

- **인터벌:** 다양한 시간 간격으로 데이터의 집계 단위를 조정할 수 있습니다.
- **적용 위젯:** 변경할 옵션을 적용할 차트 위젯을 선택하세요. 선택한 위젯에 변경한 옵션을 한번에 적용할 수 있습니다.

대시보드 자동 정렬하기

대시보드에 배치한 위젯들의 위치를 자동 정렬하려면 **자동 정렬** 버튼을 선택하세요. 위젯의 너비, 높이가 작은 크기 순으로 위젯의 배치를 자동 정렬합니다. 너비가 같으면 높이가 작은 순으로 정렬합니다.

여러 개의 차트 비교하기

대시보드에 배치한 여러 개의 차트 데이터를 하나의 차트에서 확인할 수 있는 기능입니다. 이 기능을 통해서 동시간대의 여러 지표 정보를 다른 지표와 비교해가며 확인할 수 있습니다. **차트 오버레이** 버튼을 선택하세요. 화면에 오버레이 차트를 배치할 수 있습니다.



- 오버레이 차트 왼쪽 상단에 **건수**를 선택하면 x축 차트의 데이터 기준을 변경할 수 있는 목록이 나타납니다. 원하는 항목을 선택해 차트의 데이터를 비교할 수 있습니다.
- : 현재 조회 중인 데이터를 기준으로 차트를 이미지 형식의 파일로 다운로드할 수 있습니다.
- : 현재 조회 중인 데이터를 기준으로 차트를 CSV 형식의 파일로 다운로드할 수 있습니다.
- 오른쪽 섹션에서는 각 지표별 에이전트 목록을 확인할 수 있습니다. 개별 에이전트의 색상은 차트의 그래프 색상과 일치합니다.

배치된 위젯 일괄 삭제하기

대시보드에 배치한 여러 개의 위젯을 한번에 삭제할 수 있습니다. **전체 위젯 삭제** 버튼을 선택하세요. 확인 메시지가 나타나면 **삭제** 버튼을 선택하세요.

CSV 파일로 다운로드하기

대시보드에 배치한 위젯의 데이터를 CSV 형식의 파일로 다운로드할 수 있습니다. **다운로드 > CSV 다운로드**를 선택하세요. 각 시간대별 지표의 데이터를 확인할 수 있습니다.

이미지 파일로 다운로드하기

대시보드에 배치한 모든 위젯을 하나의 이미지 형식의 파일로 다운로드할 수 있습니다. **다운로드 > 이미지 다운로드**를 선택하세요. 대시보드 화면을 그대로 캡처한 이미지를 저장할 수 있습니다.

기타 옵션 이용하기

그 밖의 다음 기타 옵션을 통해 대시보드의 레이아웃을 조정하거나 위젯의 적용된 옵션을 쉽고 간편하게 일괄 설정할 수 있습니다.



- **레이아웃**: 대시보드에서 열의 수, 즉 한 줄에 배치할 수 있는 위젯의 개수를 설정할 수 있습니다. 최대 4열까지 설정할 수 있습니다. **사용자 정의**를 선택하면 열의 수에 상관없이 위젯을 배치할 수 있습니다.

① **사용자 정의**에서는 위젯의 너비와 높이를 자유롭게 조절할 수 있습니다. 다른 옵션을 선택하면 위젯의 높이만 조절할 수 있습니다.

- **위젯 범례 표시**: 옵션을 켜거나 끄면 위젯의 차트 오른쪽에 범례를 표시하거나 숨길 수 있습니다. 옵션을 끈 상태에서는 **범례**를 클릭해 모니터링 대상을 확인하고, 선택할 수 있습니다. 범례 활용에 대한 자세한 내용은 [다음 문서](#)를 참조하세요.
- **툴팁에 대상 모두 표시**: 차트에 마우스 포인터를 가져가면 해당 그래프의 위치에 해당하는 값을 툴팁을 통해 확인할 수 있습니다. 마우스 포인터가 위치한 시간대의 모든 대상의 데이터를 툴팁을 통해서 확인하려면 옵션을 켜세요.
- **위젯 높이 일괄 적용**: 옵션을 켜면 모든 위젯의 높이를 동일하게 설정할 수 있습니다.

- **차트 시간 일괄 적용:** 개별 차트 위젯의 데이터 조회 시간을 변경할 때, 다른 위젯의 차트 시간도 같이 변경하려면 옵션을 켜세요.

튜토리얼 이용하기

화면 오른쪽 상단에 **튜토리얼** 버튼 선택 시 튜토리얼을 사용할 수 있습니다.