

Alibaba Cloud Cloud Monitor

User Guide

Issue: 20190430

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of the Alibaba Cloud website, product programs, or content shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).
6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK .
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid <i>Instance_ID</i></code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Visual reports.....	1
1.1 Use dashboards.....	1
1.1.1 Dashboard overview.....	1
1.1.2 Manage dashboards.....	2
1.1.3 Add charts.....	5
1.2 Connect CloudMonitor to Grafana.....	10
2 Host monitoring.....	21
2.1 Host monitoring overview.....	21
2.2 Process monitoring.....	22
2.3 GPU monitoring.....	26
2.4 Host monitoring metrics.....	31
2.5 Alarm service.....	39
2.6 CloudMonitor Java agent introduction.....	40
2.7 Install CloudMonitor Java agent.....	42
2.8 Introduction to the CloudMonitor GoLang agent.....	53
2.9 Install CloudMonitor GoLang agent.....	54
2.10 Agent release notes.....	66
3 Site Monitoring.....	70
4 Alarm service.....	71
4.1 Alarm service overview.....	71
4.2 Use alarm templates.....	72
4.3 Alarm rules.....	73
4.3.1 Create a threshold alert rule.....	74
4.3.2 Create an event alert rule.....	77
4.3.3 Manage alarm rules.....	80
4.3.4 Create an alert callback.....	82
4.4 Alarm contacts.....	85
4.4.1 Create an alert contact and an alert contact group.....	85
4.4.2 Manage alarm contacts and alarm contact groups.....	87
4.5 Use one-click alert.....	90
5 Availability monitoring.....	97
5.1 Create an availability monitoring task.....	97
5.2 Manage availability monitoring.....	100
5.3 Local service availability monitoring.....	102
5.4 Status codes.....	106
6 Cloud service monitoring.....	107
6.1 ApsaraDB for RDS.....	107

6.2 SLB.....	112
6.3 OSS.....	123
6.4 CDN.....	124
6.5 EIP.....	129
6.6 ApsaraDB for Memcache.....	133
6.7 ApsaraDB for Redis.....	138
6.8 ApsaraDB for MongoDB.....	142
6.9 Message Service monitoring.....	149
6.10 AnalyticDB monitoring.....	152
6.11 Log Service.....	156
6.12 Container Service monitoring.....	161
6.13 Shared Bandwidth monitoring.....	162
6.14 Global Acceleration.....	164
6.15 HiTSDB.....	169
6.16 VPN.....	173
6.17 API Gateway.....	177
6.18 DirectMail.....	181
6.19 Elasticsearch.....	184
6.20 Auto Scaling.....	188
6.21 E-MapReduce.....	192
6.22 Express Connect.....	202
6.23 Function Compute.....	206
6.24 StreamCompute.....	210
6.25 ApsaraDB for HybridDB.....	212
6.26 NAT Gateway.....	214
6.27 Open Ad.....	218
6.28 HybridDB for MySQL.....	220
7 RAM for CloudMonitor.....	224
8 Application groups.....	227
8.1 Application group overview.....	227
8.2 Create application groups.....	227
8.3 Check application group details.....	230
8.4 Modify an application group.....	235
8.5 Add resources to an application group.....	242
8.6 Apply an alert template to an application group.....	245
8.7 Manage alarm rules.....	246
9 Event monitoring.....	249
9.1 Event monitoring overview.....	249
9.2 Cloud product events.....	251
9.2.1 Cloud service events.....	251
9.2.2 View cloud service events.....	261
9.2.3 Use the event alert function for Alibaba Cloud services.....	264
9.3 Custom events.....	269
9.3.1 Report custom event data.....	269

- 9.3.2 View custom events.....279
- 9.3.3 Use the custom event alarm function.....279
- 9.3.4 Event monitoring best practices..... 281
- 10 Custom monitoring..... 288**
 - 10.1 Custom monitoring overview..... 288
 - 10.2 Report monitoring data..... 291
 - 10.3 View custom monitoring charts..... 308

1 Visual reports

1.1 Use dashboards

1.1.1 Dashboard overview

The CloudMonitor dashboard provides you with a real-time metric visualization solution for a comprehensive overview of your applications and services, enabling you to quickly troubleshoot problems and monitor resource usage.

Display metric trends for multiple instances

The dashboard provides detailed metrics and trends for multiple instances. For example, you can view the metrics of all the ECS instances on which your application is deployed all on one metric chart. This can help you see trends across multiple instances all in one area. Similarly, you can also view the CPU usage of multiple ECS instances over time in one chart.

Display multiple metrics per instance

With dashboards, you can also view several metrics of an ECS instance, such as CPU usage, memory usage, and disk usage all displayed on one metric chart. This visualization solution can help you find exceptions and monitor resource usage efficiently.

Display and sort instance resource usage

Instances can be sorted based on resource usage levels, allowing you to quickly gain insight into resource usage per instance and how usage levels differ between instances. With this information, you can make informed decisions and avoid unnecessary costs.

Display metrics distribution of multiple instances

The CPU usage distribution of an ECS instance group can be visualized with a heat map, allowing you to quickly and accurately discover the real time usage levels of different machines and compare them with each other. These heat maps are not only powerful visualization tools but are also interactive. You can click any one of the color blocks on the heat map to view the metrics and trends of the corresponding machine for a specified period of time.

Display aggregated metrics of multiple instances

With dashboards, you can view the average aggregation value of a particular metric, such as CPU usage of multiple ECS instances, all in one chart. With this capability, you quickly estimate overall CPU usage capacity and check whether the resource usage of different instances is balanced.

Provides full-screen visualization solution

The dashboard supports a full-screen mode that automatically refreshes. In this mode, you can easily add several application and product metrics to the full-screen display, allowing you to have a quick visual overview of all monitored data.

1.1.2 Manage dashboards

You can easily view, create, and delete dashboards. The procedure for these actions is as follows.

View a dashboard

You can view a dashboard to view and monitor metrics from several different products and instances all within one area.



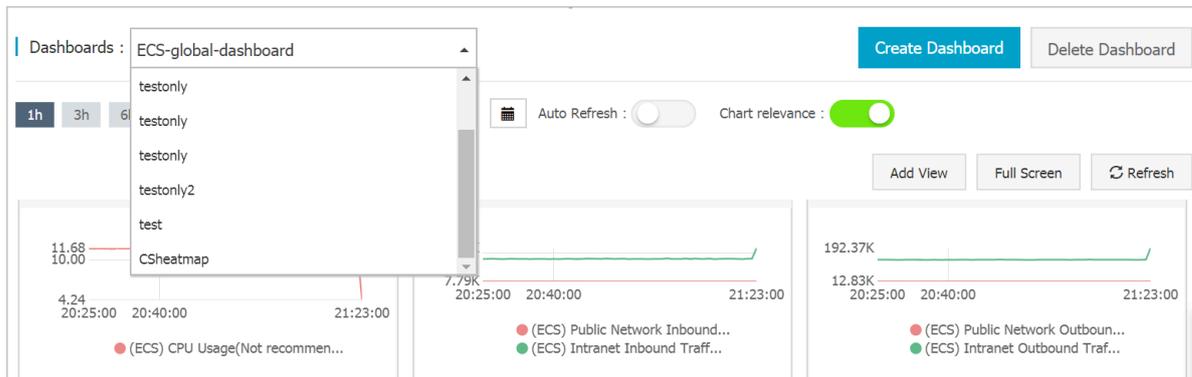
Note:

- CloudMonitor automatically initializes an ECS dashboard and displays ECS metrics.
- CloudMonitor refreshes data measured in one-hour, three-hour, and six-hour periods automatically. However, data measured for more than six hours cannot be refreshed automatically.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Dashboard > Custom Dashboard.

- By default, ECS-global-dashboard is displayed. You can select another dashboard from the drop-down list.



- To view the dashboard in full screen, click Full Screen in the upper-right corner of the page.
- Select a time range. Click the time range button at the top of the page. From there, you can quickly select the time range shown in the charts of the dashboard. The time range you select apply to all the charts on the dashboard.
- Automatic refresh. After you turn on the Auto Refresh switch, whenever you select a query time span of 1 hour, 3 hours, or 6 hours, automatic refresh is performed every minute.
- The units of the metrics measured are displayed in parentheses for the chart name.
- When you rest the pointer over some point on a chart, values at that time point are displayed across all charts.

Create a dashboard

You can create a dashboard and customize the charts for when your business operations grow complex and the default ECS dashboard does not meet your monitoring requirements.



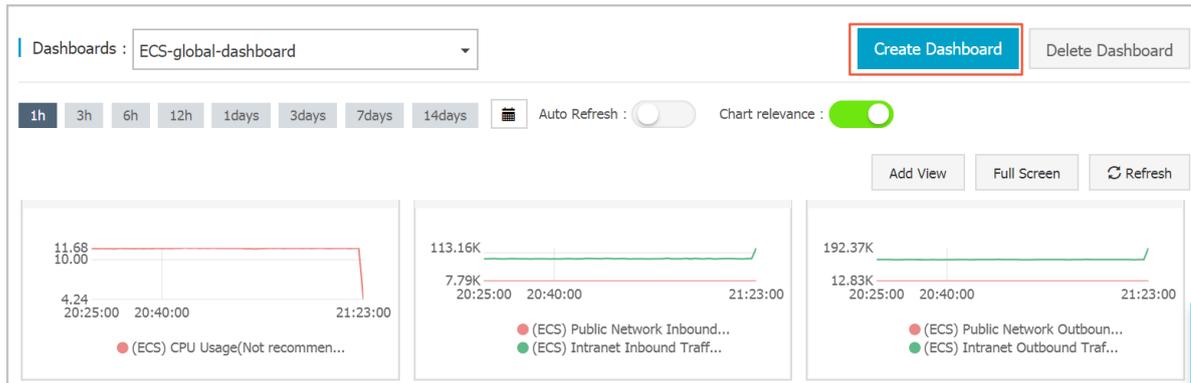
Note:

Up to 20 charts can be created on one dashboard.

Procedure

- Log on to the [CloudMonitor console](#).
- In the left-side navigation pane, choose Dashboard > Custom Dashboard.

3. In the upper-right corner of the page, click Create Dashboard.



4. Enter the name of the dashboard.

5. Click Create. The page is automatically redirected to the new dashboard page where you can add various metric charts as needed.

6. When you rest the pointer over the dashboard name, the Edit option appears on the right hand side. To modify the dashboard name, click Edit.

Delete a dashboard

You can delete a dashboard if you do not need it given changes in your business operations.



Notice:

When you delete a dashboard, all charts that are added to the dashboards are also deleted.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Dashboard > Custom Dashboard.
3. Select the target dashboard from the Dashboards drop-down list.
4. In the upper-right corner of the page, click Delete Dashboard to delete the dashboard.

1.1.3 Add charts

This topic describes several types of charts common in the CloudMonitor dashboard and how to add a chart.

Scenarios

By default, CloudMonitor creates an initialized ECS dashboard. You can add more charts and tables to the dashboard to view even more data related to your ECS instances.

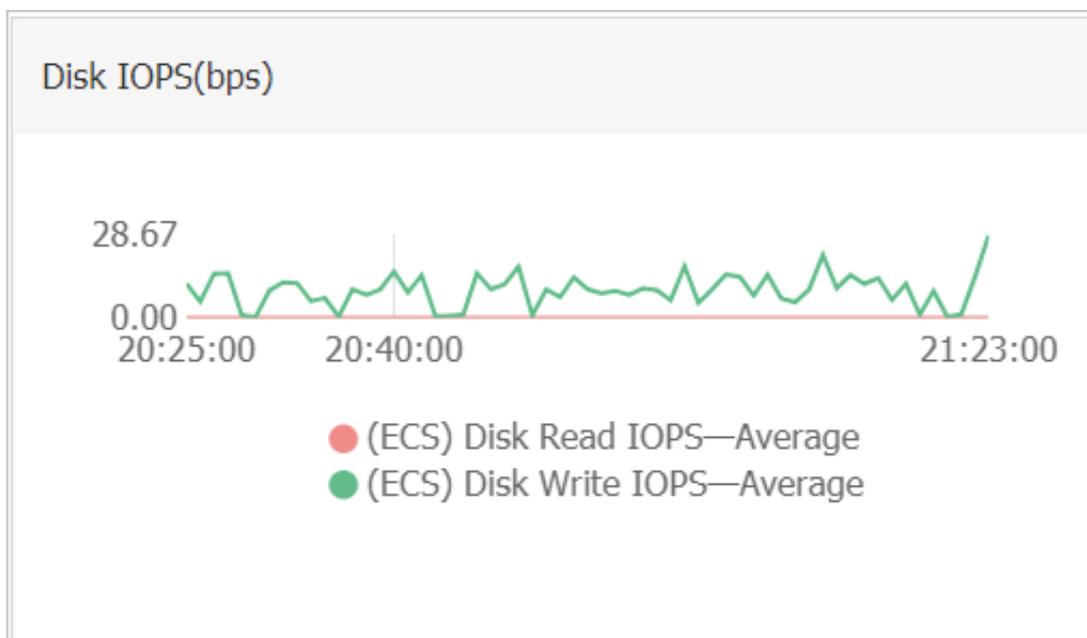
In the case that the ECS dashboard does not meet your monitoring needs, we recommend that you create an additional dashboard to which you can add charts to display custom monitoring data.

Before you begin

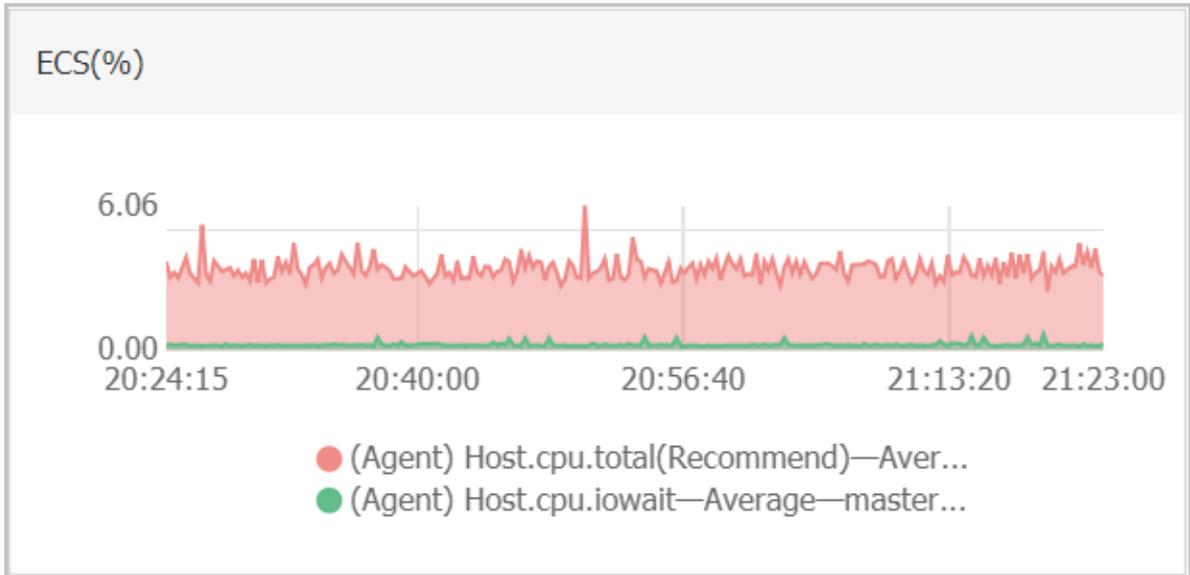
Before you can add a chart, you need to create a dashboard.

Chart types

- **Line chart:** Displays monitoring data on a basis of time series. Multiple metrics can be added.



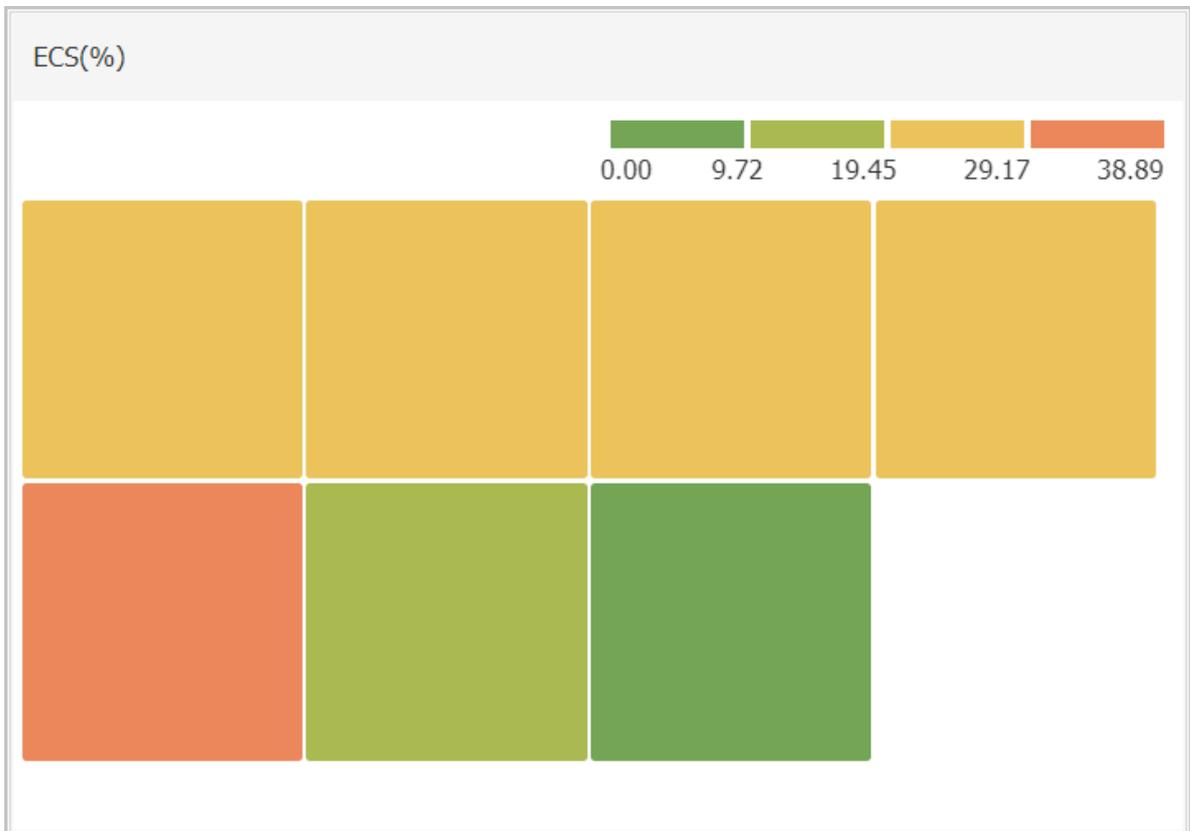
- **Area chart:** Displays monitoring data on a basis of time series. Multiple metrics can be added.



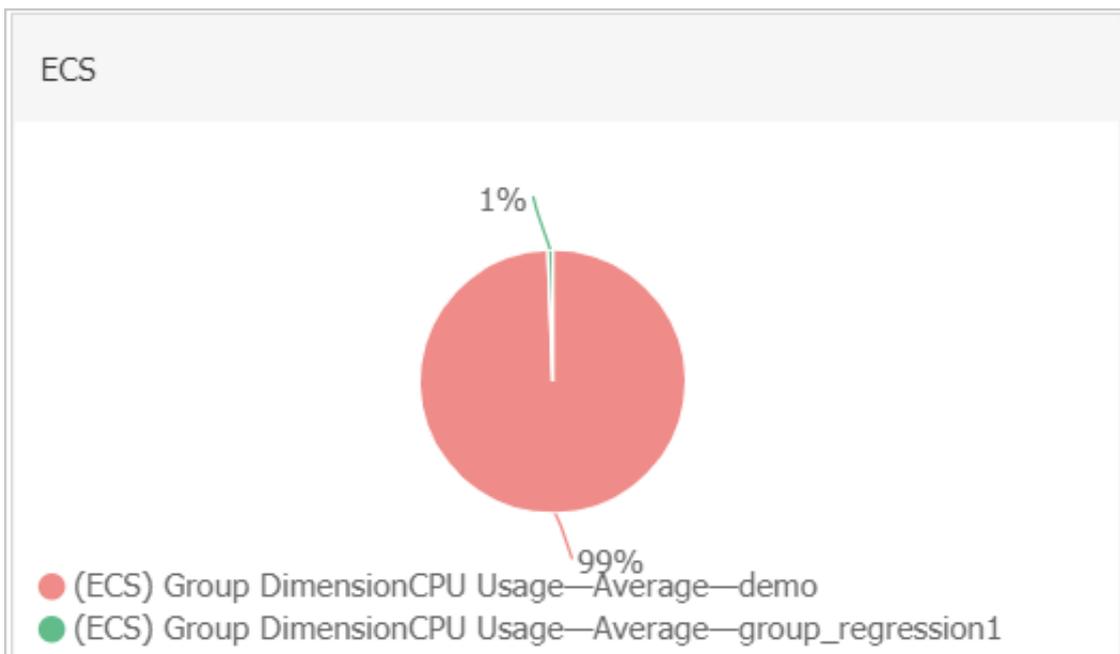
- **Table:** Displays real-time metric data in descending order. Each table displays up to 1,000 data records, which are either the first 1,000 records or the last 1,000 records. Only one metric can be added.

ECS(%)		
Time	Dimensions	Maximum Value
2018-12-06 21:25:00	ESS-asg-yinna_test	100
2018-12-06 21:20:00	node-0003-k8s-for-cs-c9ebd45a41dd645a498a5c06af2b88c53	55.56
2018-12-06 21:25:00	master-02-k8s-for-cs-c9ebd45a41dd645a498a5c06af2b88c53	38.89
2018-12-06 21:25:00	master-03-k8s-for-cs-c9ebd45a41dd645a498a5c06af2b88c53	38.1
2018-12-06 21:00:00	master-01-k8s-for-cs-c9ebd45a41dd645a498a5c06af2b88c53	37.5
2018-12-06 21:00:00	node-0001-k8s-for-cs-c9ebd45a41dd645a498a5c06af2b88c53	35.29
2018-12-06 21:20:00	node-0002-k8s-for-cs-c9ebd45a41dd645a498a5c06af2b88c53	29.41

- Heat map: Displays real-time metric data. Heat maps show the distribution and comparison of real-time data of a specific metric for multiple instances. Only one metric can be added.



- Pie chart: Displays real-time metric data and can be used for data comparisons. Only one metric can be added.



Add a chart



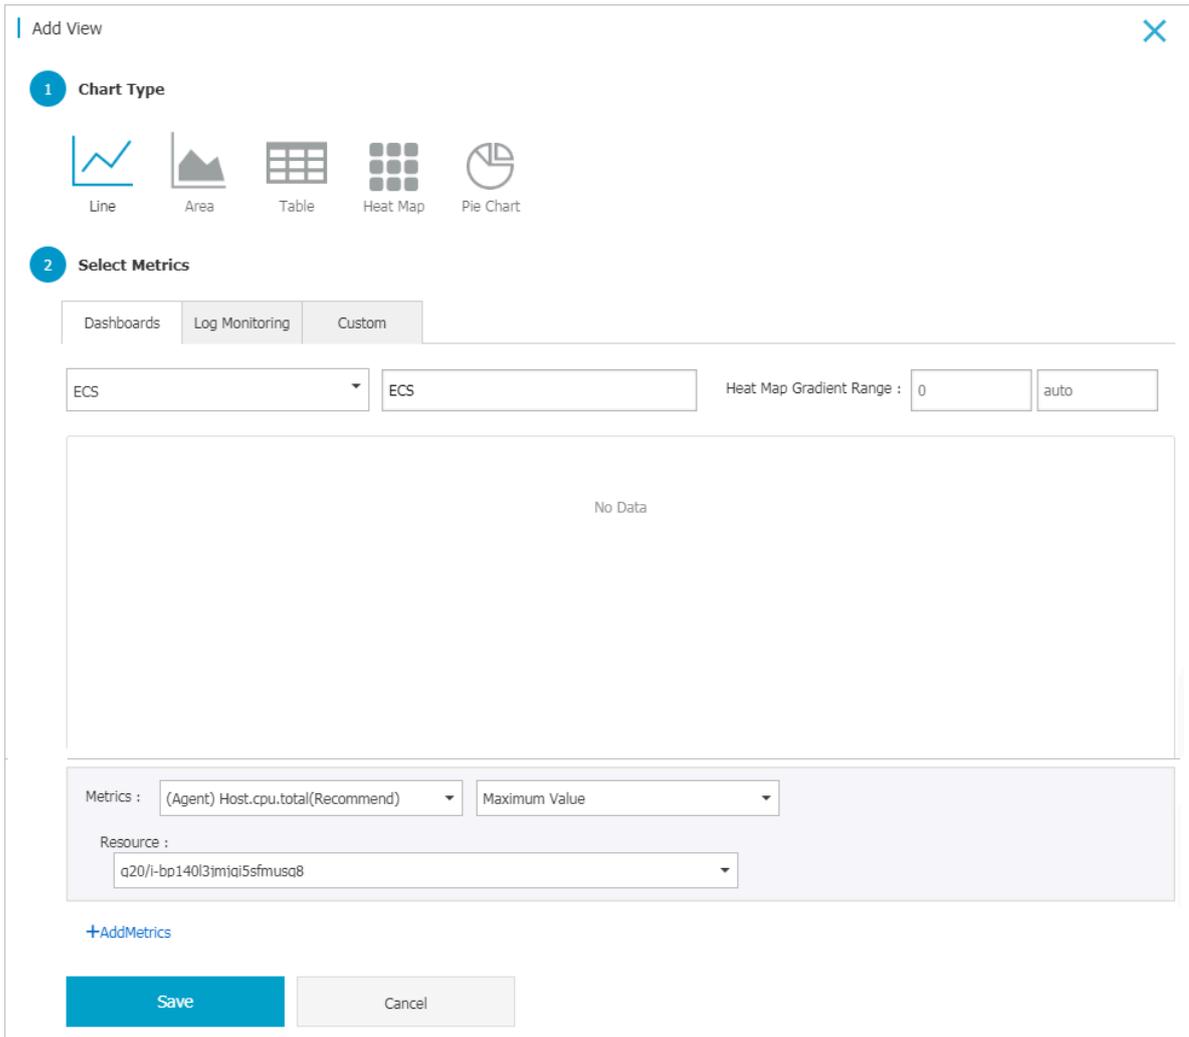
Note:

- The default ECS dashboard provides the following seven charts: CPU Usage, Network Inbound Bandwidth, Network Outbound Bandwidth, Disk BPS, Disk IOPS, Network Inbound Traffic, and Network Outbound Traffic.
- Up to 20 charts can be added in a dashboard.
- Each line chart can display up to 10 lines.
- Each area chart can display up to 10 areas.
- Each table can display up to 1,000 sorted data records.
- A heat map can display up to 1,000 color blocks.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Dashboard > Custom Dashboard.

3. In the upper-right corner of the displayed page, click Add View.



4. Select a chart type.

5. Choose from Dashboards, Log Monitoring, and Custom tab pages. In this example, click the Dashboards tab.

6. Select the target Alibaba Cloud product and enter a name for the chart.

7. Select the metric, the statistical method, and the resources.

- Select the metric you want to view.
- Select the statistical method by which the metric data is aggregated. You can choose maximum, minimum, or average.
- Select the resources that you want to monitor.

8. To add a metric, click AddMetrics and repeat the preceding steps.

9. Click Save. The chart is displayed on the dashboard.

10. If you want to resize the chart, drag the right border, lower border, or lower-right corner of the chart.

Metrics

- **Dashboards:** Displays the monitoring data of Alibaba Cloud products.
- **Log monitoring:** metrics added through log monitoring.
- **Custom:** metrics added through custom monitoring.
- **Metrics:** monitoring indicators, such as CPU usage and memory usage.
- **Statistical method:** means by which metric values are aggregated during a statistical period. Some common statistical methods are maximum, minimum, and average.
- **Resource:** You can use an application group or instance to filter resources and view the monitoring data of these resources.

1.2 Connect CloudMonitor to Grafana

This topic describes how to import monitoring data from CloudMonitor to Grafana for data visualization.

Background information

CloudMonitor stores both custom monitoring data and the system monitoring data of the core products of Alibaba Cloud. In addition to using the built-in charts, graphs, and dashboards provided by CloudMonitor to display the data, you can also use the third-party tool Grafana for further data visualization options. To use Grafana, complete the instructions in the following sections.

Preparations

1. Download and install Grafana.

You can install Grafana on CentOS by using the following two commands:

Command 1:

```
yum install https://s3-us-west-2.amazonaws.com/grafana-releases/release/grafana-5.3.0-1.x86_64.rpm
```

Command 2:

```
wget https://s3-us-west-2.amazonaws.com/grafana-releases/release/grafana-5.3.0-1.x86_64.rpm
```

```
sudo yum localinstall grafana - 5.3.0-1.x86_64.rpm
```

For more information, see [Officially recommended installation methods](#).

2. Start Grafana.

Run the `service grafana - server start` command to start Grafana.

Procedure

1. Install the CloudMonitor data source agent.

Confirm the directory in which the Grafana agent is to be installed, install the agent, and then restart grafana-server.



Note:

For example, the agent is installed in the `/var/lib/grafana/plugins/` directory on CentOS.

On CentOS, the installation command is as follows:

```
cd /var/lib/grafana/plugins/  
git clone https://github.com/aliyun/aliyun-cms-grafana.git  
service grafana - server restart
```

Alternatively, you can download `aliyun-cms-grafana.zip`, decompress it, upload it to the plugins directory of the Grafana on the server, and then restart grafana-server.



Note:

You cannot set alarms for monitoring data in the current version of Grafana.

2. Configure the CloudMonitor data source agent.

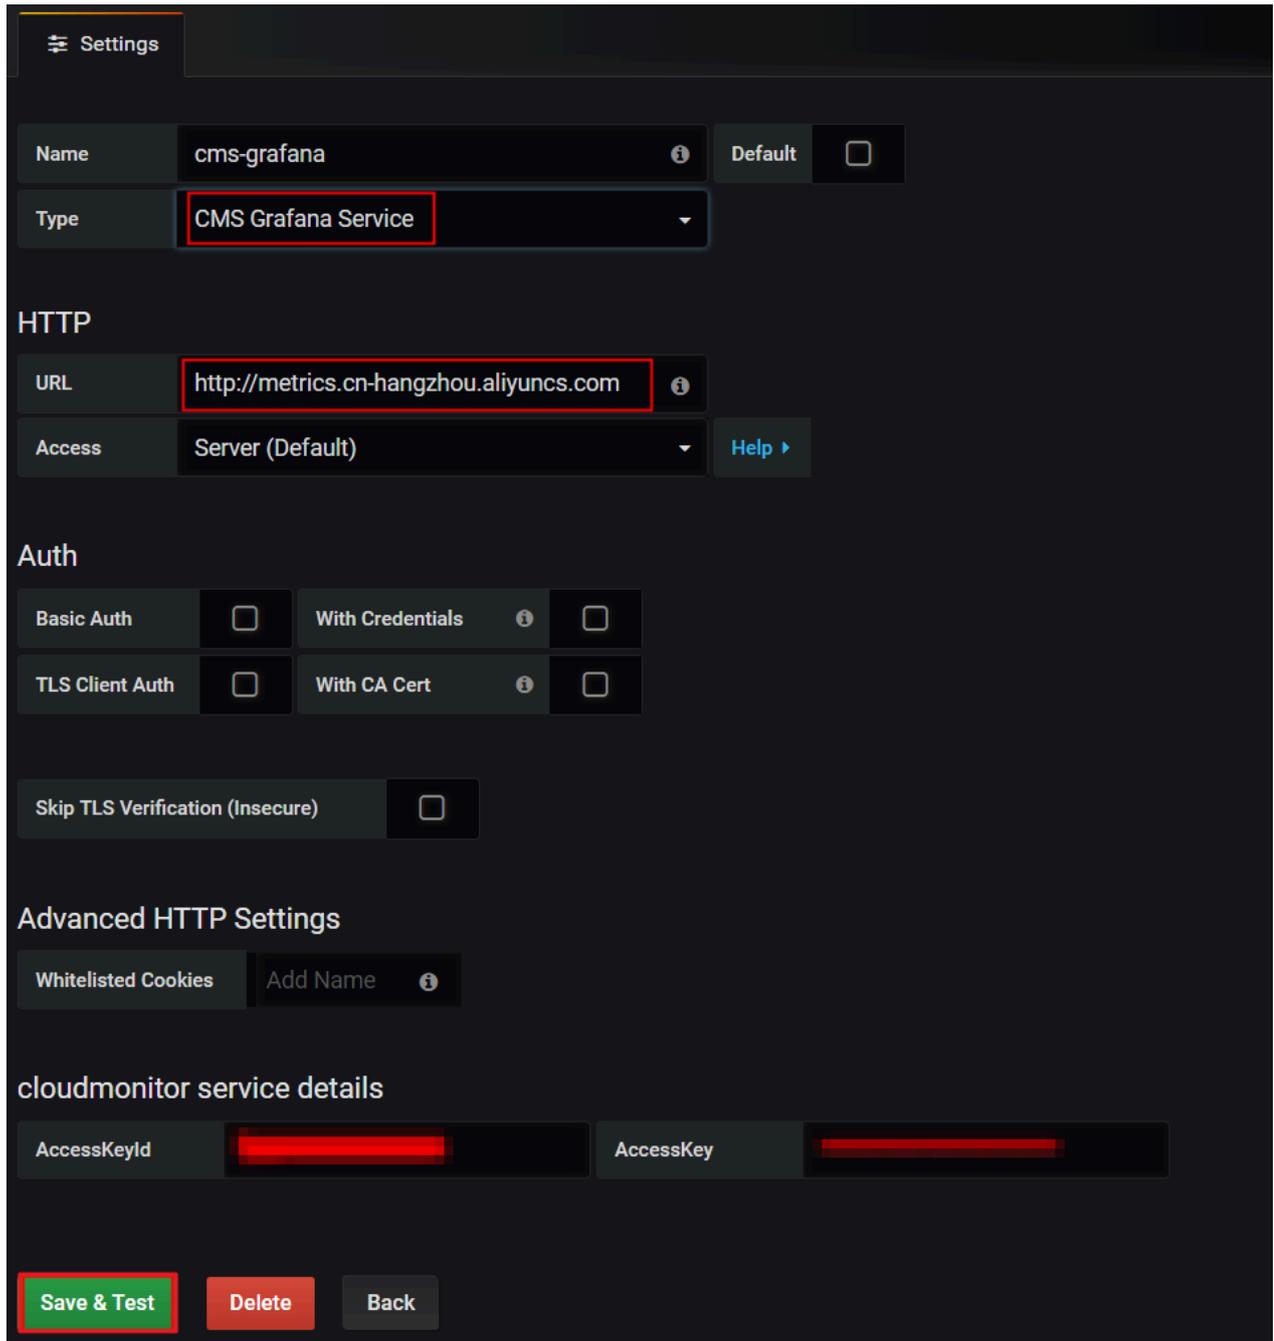
After Grafana is successfully installed, its default access port number is 3000. The user name and password are both set as admin.

- a. On the Grafana homepage, choose Configuration > Data Sources.
- b. On the Data Sources page, click Add data source in the upper-right corner.
- c. Set parameters for the data source.

Configuration item	Description
Data source	Name: Enter a name for the data source. Type: Select CMS Grafana Service.
HTTP	URL: <code>http://metrics.cn-shanghai.aliyuncs.com</code> is used as an example. For more information, see Endpoints . Access: Retain the default option.
Auth	Retain the default settings.

Configuration item	Description
CloudMonitor service details	Enter an AccessKey (AK) of an account that has the appropriate read and write permissions. The AK of your RAM user account is recommended.

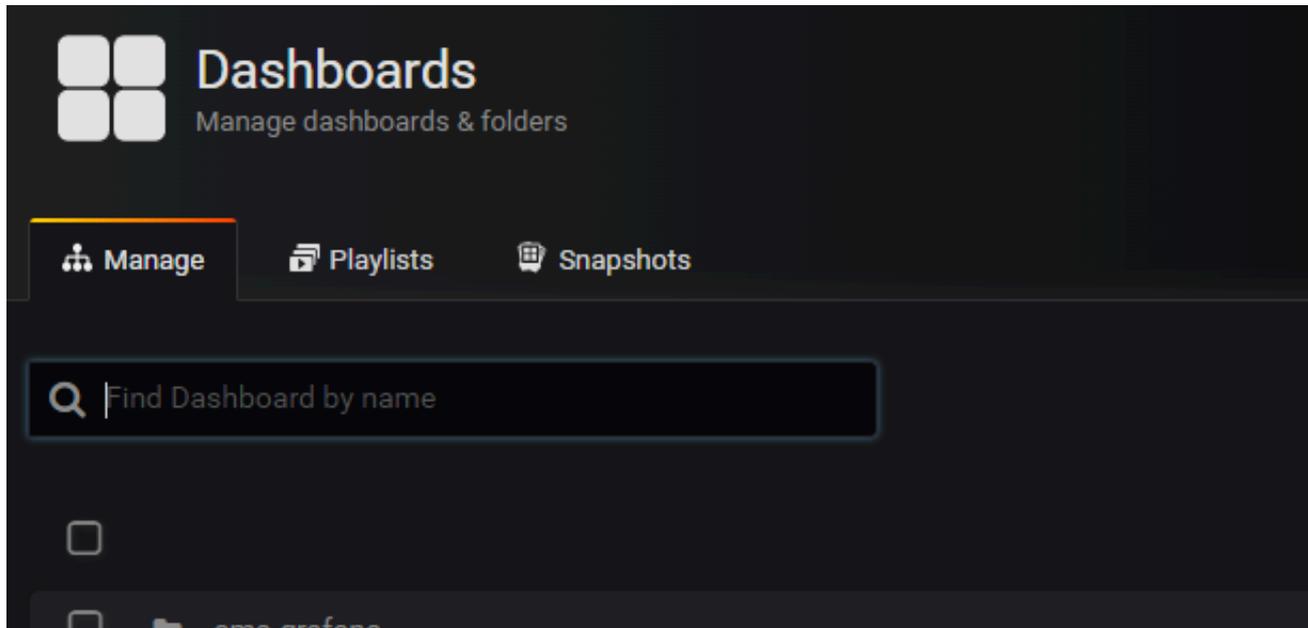
The following figure shows the configuration items.



d. Click Save & Test.

3. Create a dashboard.

- a. On the Grafana homepage, choose Dashboards > Manage.



- b. Create a dashboard by using any of the following conditions:

- Click +Dashboard.
- Click +Folder to create a folder, and then click +Dashboard.
- Click +Import to import a dashboard.

4. Configure a graph.

- a. Choose New Panel > Add > Graph and click Panel Title. In the displayed dialog box, click Edit.
- b. In the Metrics area, set datasource to cms-grafana and set Project , Metric , Period , Y - column , and X - column , as shown in the following figure.



For more information, see [QueryMetricList](#).

The following describes some of the other parameters in detail:

Group : Indicates the CloudMonitor application group to which your Alibaba Cloud account belongs.

Dimensions : Indicates the latest set of the instance monitoring data that relates to the configuration item of **Project** and **Metric** . If you set this

parameter to `Group` , monitoring data for instances in this group will be displayed.

`Y - column` : You can select more than one option.

`X - column` : Set to `timestamp` .

`Y - column describe` : Indicates what is each option displayed in `Y - column` .

For more information about the graph, click [here](#).



Note:

- You can set all the parameters manually by following the instructions in [QueryMetricList](#).
- You can enter null for a parameter to cancel it. This can be done for any of the parameters.
- You can refresh the page to view the full list or enter the `InstanceID` in the search bar in the case of incomplete information relating to the instances (previously set as dimensions).

For custom monitoring data, you need to manually enter the following parameters:

- `Project` : Enter `acs_custom Metric` and your Alibaba Cloud account ID.
- `Metric` : Indicates the `metricName` for reporting monitoring data.
- `Period` : Indicates the period of time for reporting monitoring data.
- `Group` : Indicates the group ID corresponding to `Metric` .
- `Dimensions` : Indicates the dimension for reporting monitoring data.

Currently, no drop-down list is available that can provide multiple options. Moreover, only one dimension can be selected at a time. Selecting more than one dimension is currently not supported. Therefore, if you enter multiple dimensions, only the first one will be valid by default.



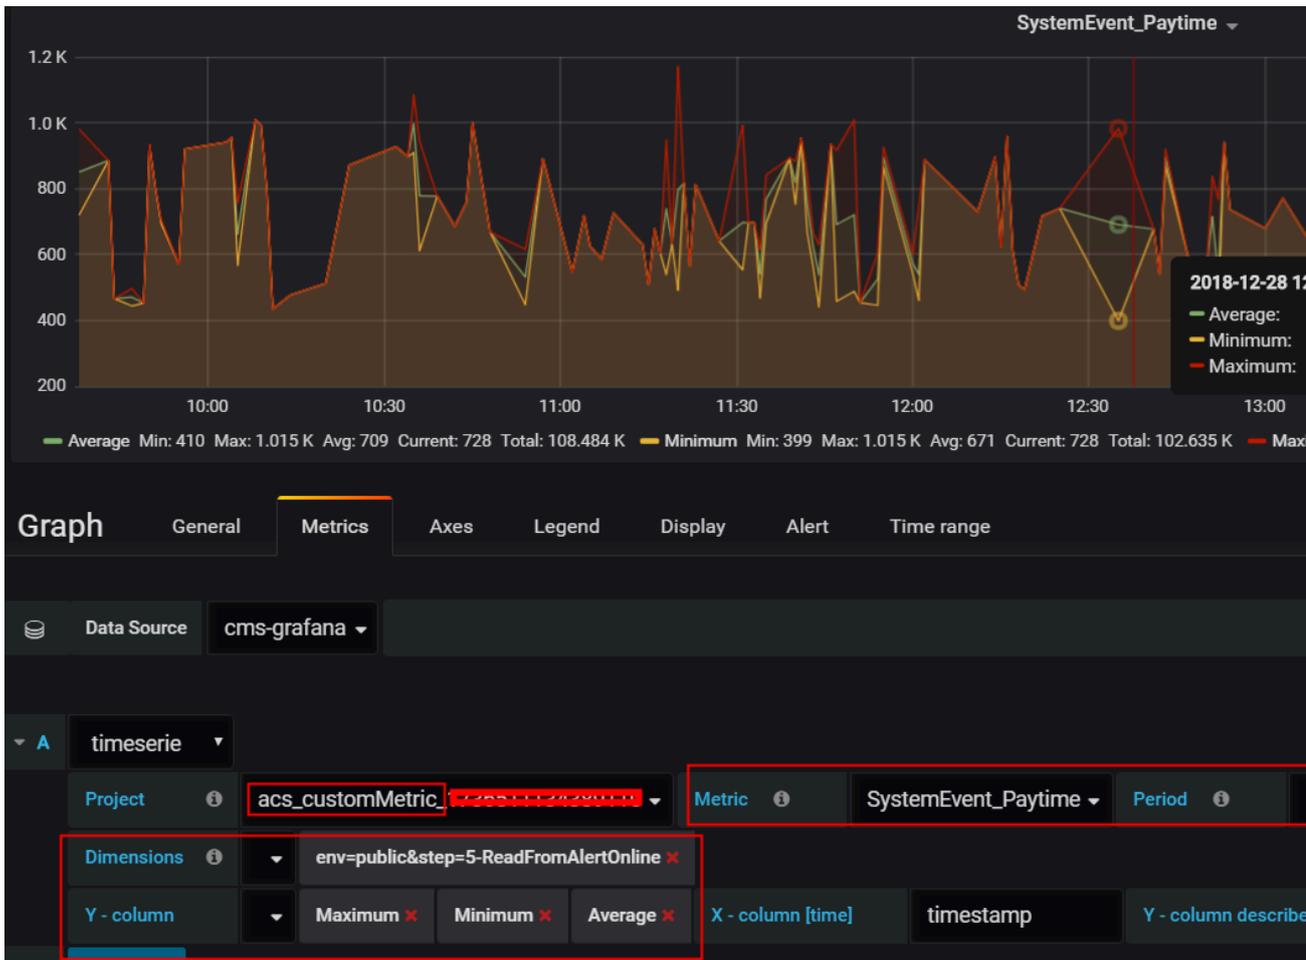
Note:

If the `dimensions` provided by the CloudMonitor console are found in the following format `env : public , step : 5 - ReadFromAl`

ertOnline , then you will need to replace the commas (,) with ampersands (&).

- Y - column :Includes Average , Maximum , Minimum , Sum , SampleCount , P10 , P20 , P99 , along with other options for reporting monitoring data.
- X - column :Set to timestamp .

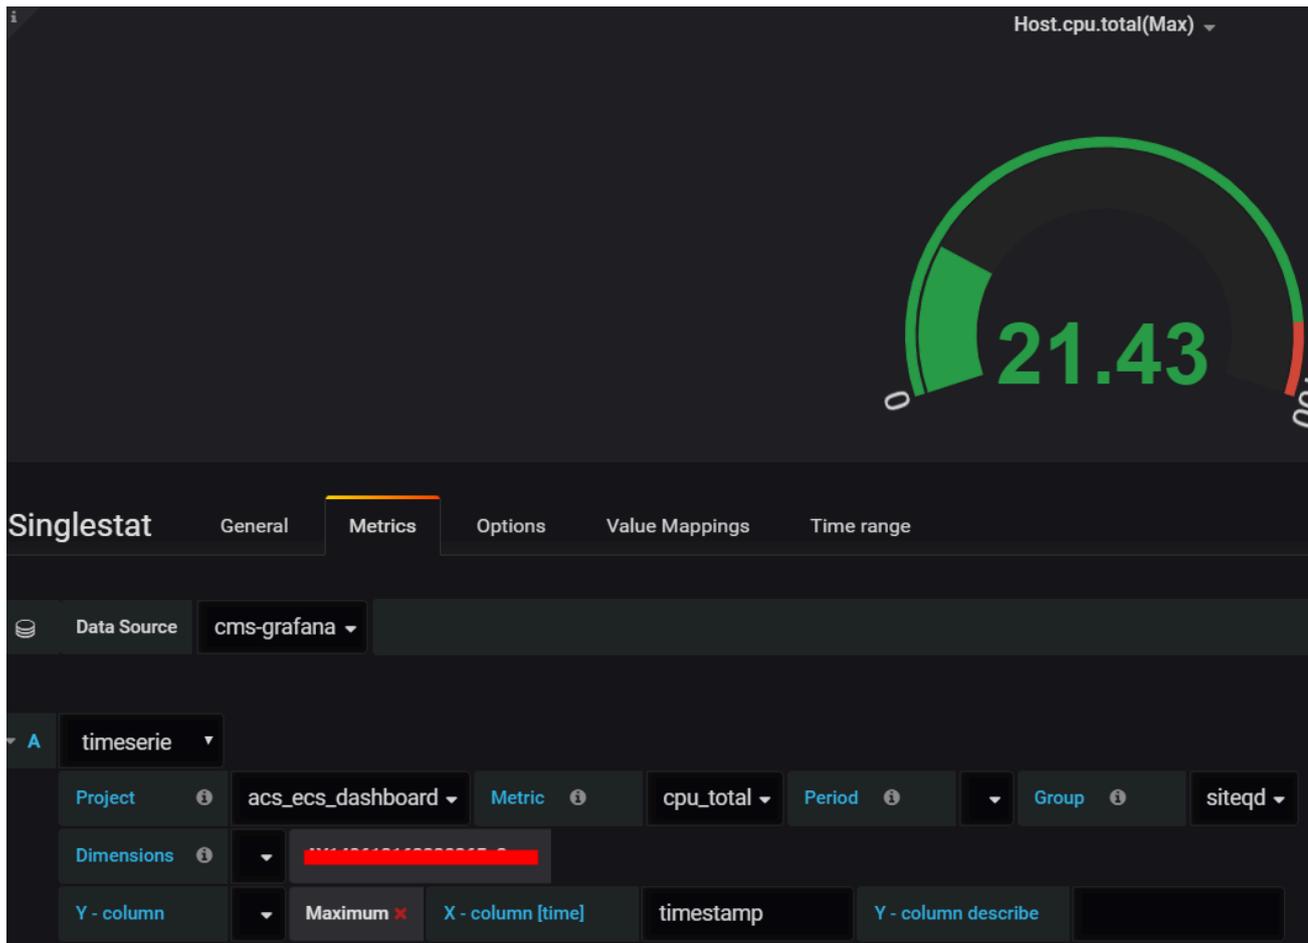
The following figure shows an example visualization for custom monitoring data



5. Configure the Singlestat panel.

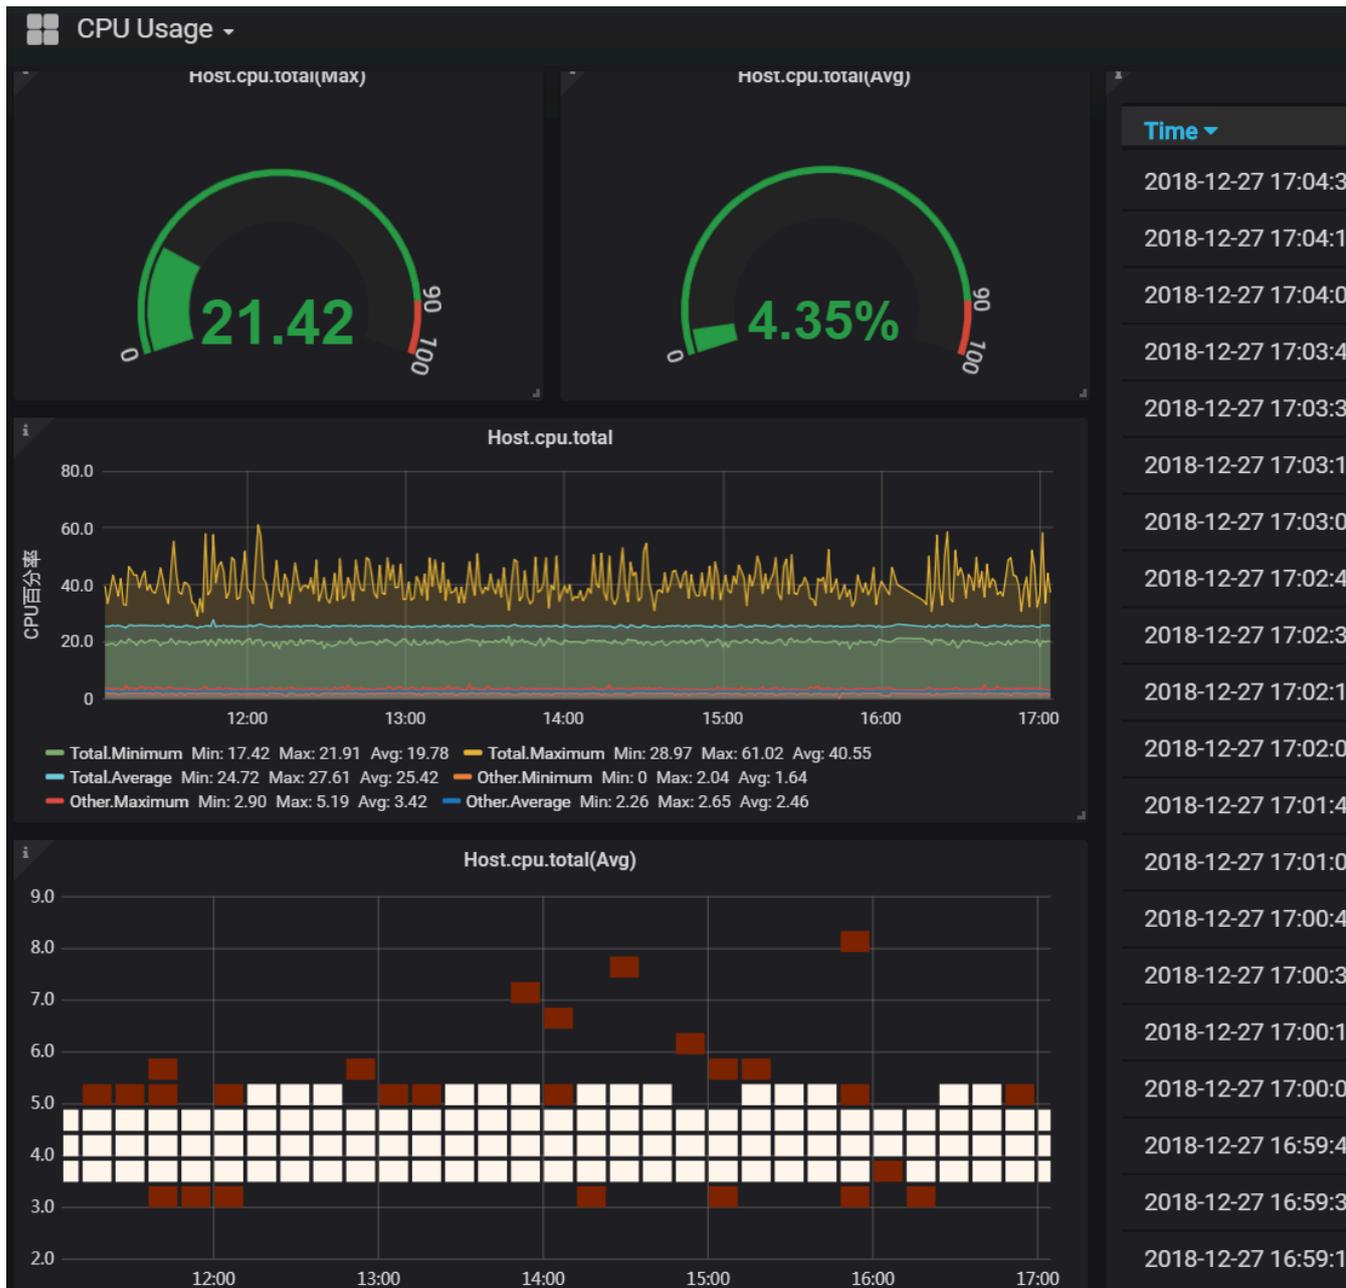
- a. Choose New Panel > Add > Singlestat and click Panel Title. In the displayed dialog box, click Edit.
- b. In the Metric area, set parameters by following the instructions provided in step 4.

The following figure shows an example of a configured Singlestat panel.



For more information, see [Singlestat](#).

6. View monitoring results.





2 Host monitoring

2.1 Host monitoring overview

The host monitoring service of CloudMonitor allows you to monitor your servers in a systematic manner by installing an agent on the servers. Host monitoring currently supports Linux and Windows Operating Systems (OSs).

Scenarios

Host monitoring is available for both Alibaba Cloud ECS servers, and virtual and physical machines provided by other vendors.

Host monitoring collects statistics of a diverse range of OS-related metrics by using the agent, allowing you to retrieve the server resource usage and obtain metrics for troubleshooting.

Hybrid cloud monitoring solution

Host monitoring uses the agent to collect server metrics. You can install the agent on an ECS server or a non-ECS server for monitoring on and off the cloud.

Enterprise-level monitoring solution

Host monitoring also provides an application group function, which allows you to allocate servers from different regions of Alibaba Cloud to the same group for more efficient server management from a business operations perspective. Host monitoring supports group-based alarm management, meaning that you only need to configure one alarm rule for the entire group, which can improve O&M efficiency and the overall management experience.



Note:

- Host monitoring supports Linux and Windows, but does not support Unix.
- Root permissions are required for the agent installation on a Linux OS and administrator permissions are required for that on a Windows OS.

- The TCP status statistics function is similar to the Linux `netstat -anp` command. This function is disabled by default because a large portion of CPU time is consumed when many TCP connections exist.
 - To enable this function in Linux, set `netstat . tcp . disable` in the `cloudmonit` or `/config/conf.properties` configuration file to `false`. Restart the agent after you modify the configuration.
 - To enable this function in Windows, set `netstat . tcp . disable` in the `C:\Program Files\Alibaba\cloudmonit` or `\config` configuration file to `false`. Restart the agent after you modify the configuration.

Monitoring capability

Host monitoring provides more than 30 metrics covering CPU, memory, disk, and network to meet your monitoring and O&M requirements. Click [here](#) to view the full list of the metrics.

Alarm capability

Host monitoring provides an alarm service for all metrics, allowing you to set alarm rules for instances, application groups, and all resources. You can use the alarm service according to your business requirements.

You can use the alarm service directly in the host monitoring list or apply the alarm rules to your application groups after you add servers into the groups.

2.2 Process monitoring

By default, process monitoring allows you to collect information about CPU usage, memory usage, and the number of files recently opened by active processes during some period of time. If you add a process keyword, the number of processes containing the keyword is collected.

View the resource consumption of active processes

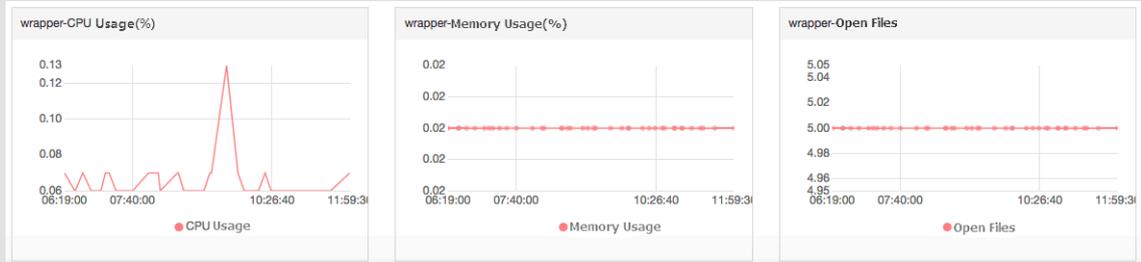
- The CloudMonitor agent filters out the top five processes with the most CPU usage every minute, and records the respective CPU usage, memory usage, and number of files opened by these processes.
- For the CPU and memory usage of a process, see the Linux `top` command.

- For the number of files opened by an active process, see the Linux `lsof` command.

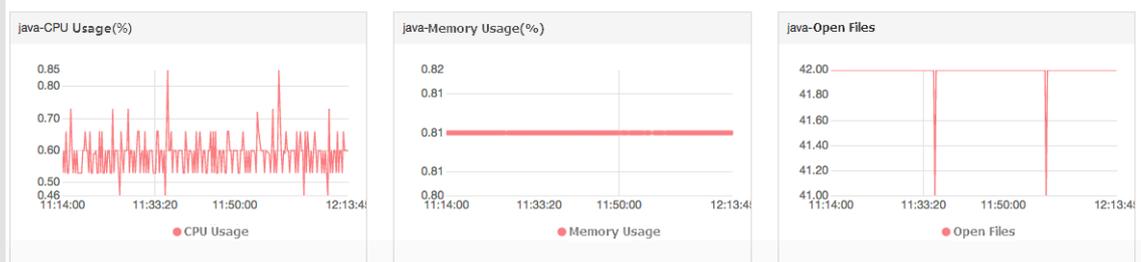
**Note:**

- If your process occupies multiple CPU cores, the percentage shown for CPU usage may exceed 100% because the collected result indicates the total usage of the multiple CPU cores.
- If, during the time period specified for your query, the top five processes have changed, the process list will display all processes that have ever ranked as top five over the specified time period. The times in the list indicate when the processes last ranked in the top five.
- The CPU usage and memory usage, and the number of opened files are collected only for the top five processes. Therefore, if a process has not ranked top five continuously over the time period specified for your query, its data points will appear discontinuous in the charts. The density of the data points for a process indicates its degree of activity on the server.
 - As shown in the following figure, the wrapper process has not continuously ranked in the top five processes each time measured. Therefore, the data points in the charts are sparse and discontinuous. The data points in the

following charts mean that the process has ranked top five for the particular time measured.



- The following figure shows the charts of the java process. The data points in the charts are dense and continuous. This means that the process continuously ranks in the top five processes with the most CPU usage.



Monitor the number of specified processes

You can learn the number and viability status of key processes by monitoring the number of processes. Specifically, you can add process keywords to the Number of Processes(Count) chart to monitor the number of related processes.

- Add processes for monitoring

For example, assume the following processes run on your server: `/usr/bin/java -Xmx2300m -Xms2300m org.apache.catalina.startup.Bootstrap`, `/usr/bin/ruby`, and `nginx -c /etc/nginx/nginx.conf`. You then add the following six keywords (the keywords can

be process names, file paths, parameter names, or other related words), and the corresponding number of processes for each target keyword is output as follows:

- Keyword: `ruby` , number of processes collected: 1
- Keyword: `nginx` , number of processes collected: 1
- Keyword: `/usr/bin` , number of processes collected: 2
- Keyword: `apache . catalina` , number of processes collected: 1
- Keyword: `nginx . conf` , number of processes collected: 1
- Keyword: `- c` , number of processes collected: 1

Procedure

1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Host Monitoring.
 3. Click the name of the target host, or click Monitoring Charts in the Actions column to access the host monitoring details page.
 4. On the displayed page, click the Process Monitoring tab.
 5. Rest the pointer over the Number of Processes(Count) chart, and then click Add Process.
 6. On the displayed Add Process Monitor page, add the name or keyword of the process you want to monitor and click Add.
- Delete a monitored process
1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Host Monitoring.
 3. Click the name of the target host, or click Monitoring Charts in the Actions column to access the host monitoring details page.
 4. On the displayed page, click the Process Monitoring tab.
 5. Rest the pointer over the Number of Processes(Count) chart, and then click Add Process.
 6. On the displayed page, find the target process name or keyword and click Delete.

- Set alarm rules

After you configure monitoring for the specified process, you can configure alarm rules for the process. After that, you can receive an alarm notification when the number of the processes changes.

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. Find the host for which you want to set process monitoring alarm rules, and then click Alarm Rules in the actions column.
4. Click Create Alarm Rule in the upper-right corner of the page.
5. In the Set Alarm Rules area, select (Agent)Host.process.number from the Rule Describe drop-down list, set an appropriate alarm threshold, and then select the process you want to monitor from the processName drop-down list. If multiple processes are configured on the host, the number of processes varies. You can click Add Alarm Rule to configure alarm rules for multiple processes at a time.

The screenshot shows the 'Set Alarm Rules' configuration page. It features two rows for defining alarm rules. Each row includes an 'Alarm Rule' input field, a 'Rule Describe' dropdown menu, a time interval selector, an aggregation function, a comparison operator, a threshold value, and a unit selector. The first rule is configured with '(Agent) Host.process.number', '1mins', 'Average', '<', '1', and 'Count/Min'. The second rule uses '(Agent) Host.process.number', '5mins', 'Average', '>', '6', and 'Count/Min'. A '+Add Alarm Rule' button is located at the bottom left. On the right, a line chart displays the metric values over time, with a horizontal line at 6.00. The chart title is '(Agent) Host.process.number—Average—emr_C-7AF9E7BFD87B0EDF_2_RWJW—dfas' and the y-axis ranges from 0.00 to 6.00.

2.3 GPU monitoring

You can query GPU monitoring data either by using the CloudMonitor console or by calling APIs.

Metrics

The metrics for GPU monitoring are based on three dimensions: GPU, instance, and application group.

- GPU-dimension metrics

GPU-dimension metrics measure monitoring data on a per GPU basis. The following table lists GPU-dimension metrics.

Metric	Unit	Description	Dimensions
gpu_memory_freespace	Byte	The free memory of a GPU	instanceId, gpuId
gpu_memory_totalspace	Byte	The total memory of a GPU	instanceId, gpuId
gpu_memory_usedspace	Byte	The memory in use of a CPU	instanceId, gpuId
gpu_gpu_utilization	%	The usage of a GPU	instanceId, gpuId
gpu_encoder_utilization	%	The usage of an encoder with GPU support	instanceId, gpuId
gpu_decoder_utilization	%	The usage of an decoder with GPU support	instanceId, gpuId
gpu_gpu_temperature	°C	The temperature of a GPU	instanceId, gpuId
gpu_power_readings_power_draw	W	The power of a GPU	instanceId, gpuId
gpu_memory_freeutilization	%	The percentage of the free memory of a GPU	instanceId, gpuId
gpu_memory_useutilization	%	The percentage of the memory in use of a GPU	instanceId, gpuId

- Instance-dimension metrics

Instance-dimension metrics measure the maximum, minimum, or average value of multiple GPUs on a per instance basis, so that you can query the overall resource usage at the instance level.

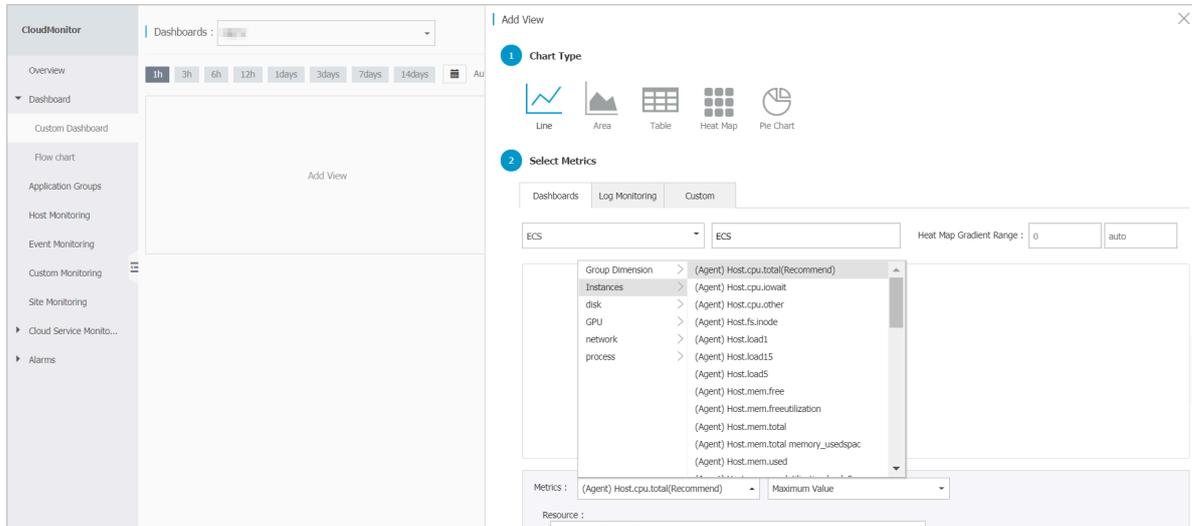
Metric	Unit	Description	Dimension
instance_gpu_decoder_utilization	%	GPU decoder usage at the instance level	instanceId
instance_gpu_encoder_utilization	%	GPU encoder usage at the instance level	instanceId
instance_gpu_gpu_temperature	°C	GPU temperature at the instance level	instanceId
instance_gpu_gpu_utilization	%	GPU usage at the instance level	instanceId
instance_gpu_memory_freespace	Byte	Free GPU memory at the instance level	instanceId
instance_gpu_memory_freeutilization	%	The percentage of free GPU memory at the instance level	instanceId
instance_gpu_memory_totalspace	Byte	GPU memory at the instance level	instanceId
instance_gpu_memory_usedspace	Byte	GPU memory in use at the instance level	instanceId
instance_gpu_memory_utilization	%	GPU memory usage at the instance level	instanceId
instance_gpu_power_readings_power_draw	W	GPU power at the instance level	instanceId

- Group-dimension metrics

Group-dimension metrics measure the maximum, minimum, or average value of multiple instances on a per group basis, so that you can query the overall resource usage at the group level.

Metric	Unit	Description	Dimension
group_gpu_decoder_utilization	%	GPU decoder usage at the application group level	groupId
group_gpu_encoder_utilization	%	GPU encoder usage at the application group level	groupId
group_gpu_gpu_temperature	°C	GPU temperature at the application group level	groupId
group_gpu_gpu_utilization	%	GPU usage at the application group level	groupId
group_gpu_memory_freespace	Byte	Free GPU memory at the application group level	groupId
group_gpu_memory_freeutilization	%	The percentage of free GPU memory at the application group level	groupId
group_gpu_memory_totalspace	Byte	GPU memory at the application group level	groupId
group_gpu_memory_usedspace	Byte	GPU memory in use at the application group level	groupId
group_gpu_memory_utilization	%	GPU memory usage at the application group level	groupId
group_gpu_power_readings_power_draw	W	GPU power at the application group level	groupId

5. On the displayed page of the created dashboard, click Add View.
6. On the Add View page, select the chart type, and then select the metrics.



7. Click Save.

Set alarm rules

We recommend that you use alarm templates to set alarm rules for new GPU metrics in batches. You can create alarm templates for the GPU metrics and then apply the templates to related application groups. For more information, see [Create an alarm template](#).

Query GPU monitoring data through APIs

- For more information about how to call APIs to query GPU monitoring data, see [QueryMetricList](#).
- **Parameter description:** The `Project` parameter should be set to `acs_ecs_dashboard`. For the values of `Metric` and `Dimensions`, see the GPU metrics in the preceding tables.

2.4 Host monitoring metrics

Host monitoring metrics include agent-collected metrics and ECS basic metrics. Agent-collected metrics are monitored by the CloudMonitor agent and monitoring data is collected every 15 seconds. The monitoring data of ECS basic metrics is collected every minute.



Note:

The ECS basic monitoring data may be inconsistent with the agent-collected monitoring data mainly because of the following reasons:

- Different monitoring frequencies

The monitoring data displayed on monitoring charts is the average value of the data collected during one statistical period. The statistical period of ECS basic monitoring data is one minute, whereas the statistical period of agent-collected monitoring data is 15 seconds. In the case of large monitoring data fluctuations, the value of ECS basic monitoring data is smaller than that of agent-collected data.

- Different monitoring perspectives

The network traffic data collected by monitoring ECS basic metrics is used for billing. It does not include the traffic between ECS and SLB because such traffic is not billed. However, the network traffic data collected through the CloudMonitor agent records the actual network traffic of each NIC. Therefore, the network traffic data collected through the agent is greater than that collected by monitoring ECS basic metrics (that is, the agent-collected data value is greater than the actually purchased bandwidth or traffic quota).

Agent-collected metrics

- CPU metrics

You can refer to the Linux top command to understand the meaning of the metrics listed in the following table.

Metric	Definition	Unit	Description
Host.cpu.idle	The percentage of CPU currently not utilized	%	The percentage of CPU currently in the idle state.
Host.cpu.system	The percentage of CPU currently occupied by the kernel space	%	Measures the CPU occupied by system context switchover. A great value indicates that many processes or threads are running on the server.

Metric	Definition	Unit	Description
Host.cpu.user	The percentage of CPU currently occupied by user processes	%	Measures the CPU occupied by user processes.
Host.cpu.iowait	The percentage of CPU currently waiting for I/O operations	%	A high value indicates frequent I/O operations.
Host.cpu.other	The percentage of CPU occupied by other operations	%	Calculation method: CPU usage of Nice + CPU usage of SoftIrq + CPU usage of Irq + CPU usage of Stolen.
Host.cpu.totalUsed	The percentage of CPU currently occupied	%	The sum of the preceding CPU consumption. It is usually used for alarm purposes.

- Memory metrics

You can refer to the free command to understand the meaning of the metrics listed in the following table.

Metric	Definition	Unit	Description
Host.mem.total	Total memory	Byte	The total memory of the server.
Host.mem.used	The amount of memory in use	Byte	Calculation method: the memory used by user programs + buffers + cached . "buffers" is the memory space occupied by the buffer. "cached" is the memory space occupied by system cache.

Metric	Definition	Unit	Description
Host.mem.actualused	The memory actually used by the user	Byte	<ul style="list-style-type: none"> - Calculation method 1: the memory in use - buffers - cached. - Calculation method 2: total memory - available memory. CentOS 7.2, Ubuntu 16.04, and later versions use the new Linux kernel, which is more accurate in memory estimation. For the specific meaning of the column of available, refer to Commit.
Host.mem.free	The amount of the memory not in use	Byte	Calculation method: total memory - memory in use.
Host.mem.freeutilization	The percentage of available memory	%	Calculation method: available memory/total memory \times 100%.
Host.mem.usedutilization	The memory usage	%	Calculation method: actually used memory/total memory \times 100%.

- Metrics of average system loads

You can refer to the Linux TOP command to understand the meaning of the metrics listed in the following table. A higher value of a metric indicates a busier system.

Metric	Definition	Unit
Host.load1	The average system loads over the past one minute. This metric is not available for Windows operating systems.	None
Host.load5	The average system loads over the past five minutes. This metric is not available for Windows operating systems.	None
Host.load15	The average system loads over the past 15 minutes. This metric is not available for Windows operating systems.	None

- Disk metrics

- You can refer to the Linux df command to understand the disk usage and inode usage metrics.
- You can refer to the Linux iostat command to understand the disk read/write metrics.

Metric	Definition	Unit
Host.diskusage.used	The space of the disk in use	Byte
Host.disk.utilization	The disk usage	%
Host.diskusage.free	The remaining storage space of the disk	Byte
Host.diskusage.total	The total disk storage	Byte
Host.disk.readbytes	The number of bytes read per second on the disk	Byte/s
Host.disk.writebytes	The number of bytes written per second on the disk	Byte/s

Metric	Definition	Unit
Host.disk.readiops	The number of read requests received by the disk per second	requests/s
Host.disk.writeiops	The number of write requests received by the disk per second	requests/s

- File system metrics

Metric	Definition	Unit	Description
Host.fs.inode	Inode usage	%	This metric is not available for Windows operating systems . Linux and UNIX systems use inode numbers, instead of file names, to identify files . When inode numbers are used up, new files cannot be created even if the disk space has not been filled up. Therefore, the inode usage must be monitored. The number of inode numbers indicates the number of files . A large number of small files can cause a high inode usage.

- Network metrics

- You can refer to the Linux `iftop` command to understand the network related metrics. You can refer to the Linux `ss` command for the collection of TCP connection data.
- The following TCP connection data is collected by default: `TCP_TOTAL` (the total number of connections), `ESTABLISHED` (the number of established connections), and `NON_ESTABLISHED` (the number of connections not in the established state). If you want to obtain such data, follow these steps:

- Linux

Change the value of `netstat . tcp . disable` in the configuration file `cloudmonit` or `/ config / conf . properties` to `false` to collect the data. Then, restart the agent.

- Windows

Change the value of `netstat . tcp . disable` in the configuration file `C : \ " Program Files " \ Alibaba \ cloudmonit` or `\ config` to `false` to collect the data. Then, restart the agent.

Metric	Definition	Unit
Host.netin.rate	The number of bits received by the NIC per second, that is, the upstream bandwidth of the NIC	bit/s
Host.netout.rate	The number of bits sent by the NIC per second, that is, the downstream bandwidth of the NIC	bit/s
Host.netin.packages	The number of packets received by the NIC per second	packets/s
Host.netout.packages	The number of packets sent by the NIC per second	packets/s
Host.netin.errorpackage	The number of incoming error packets detected by the drive	packets/s

Metric	Definition	Unit
Host.netout.errorpackages	The number of outgoing error packets detected by the drive	packets/s
Host.tcpconnection	The number of TCP connections in various states, including LISTEN, SYN_SENT, ESTABLISHED, SYN_RECV, FIN_WAIT1, CLOSE_WAIT, FIN_WAIT2, LAST_ACK, TIME_WAIT, CLOSING, and CLOSED.	None

- Process metrics

- For the CPU usage and memory usage of processes, refer to the Linux top command. The CPU usage indicates the consumption of multi-core CPUs.
- For details about Host.process.openfile, refer to the Linux lsof command.
- For details about Host.process.number, refer to the Linux ps aux |grep 'keyword' command.

Metric	Definition	Unit
Host.process.cpu	The CPU usage of a process	%
Host.process.memory	The memory usage of a process	%
Host.process.openfile	The number of files opened by the current process	Count
Host.process.number	The number of processes that match the specified keyword	Count

ECS basic metrics

If your host is an ECS server, the metrics listed in the following table are monitored automatically after you purchase the ECS instance. You do not need to install the agent. The monitoring frequency is one minute.

Metric	Definition	Unit
ECS.CPUUtilization	CPU usage	%

Metric	Definition	Unit
ECS.InternetInRate	The average rate of inbound Internet traffic	bit/s
ECS.IntranetInRate	The average rate of inbound intranet traffic	bit/s
ECS.InternetOutRate	The average rate of outbound Internet traffic	bit/s
ECS.IntranetOutRate	The average rate of outbound intranet traffic	bit/s
ECS.SystemDiskReadbps	The number of bytes read on the system disk per second	Byte/s
ECS.SystemDiskWritebps	The number of bytes written on the system disk per second	Byte/s
ECS.SystemDiskReadOps	The read times of the system disk per second	packets/s
ECS.SystemDiskWriteOps	The write times of the system disk per second	times/s
ECS.InternetIn	Internet inbound traffic	Byte
ECS.InternetOut	Internet outbound traffic	Byte
ECS.IntranetIn	Intranet inbound traffic	Byte
ECS.IntranetOut	Intranet outbound traffic	Byte

2.5 Alarm service

Host monitoring provides the alarm service so that you can set alarm rules for a target server, or add servers to an application group and then set alarm rules at the group level. For more information about setting alarm rules for an application group, see [Manage alarm rules](#).

Create an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. Click the Alarm Rules tab.
4. Click Create Alarm Rule in the upper-right corner.

5. In the displayed dialog box, set the parameters. For more information, see [Manage alarm rules](#).
6. Click Confirm to save your alarm rule settings.

Delete an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. Click the Alarm Rules tab.
4. Find the target alarm rule and click Delete in the Actions column. If you want to delete multiple rules at a time, select the target rules and click Delete under the alarm rule list.

Modify an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. Click the Alarm Rules tab.
4. Find the target alarm rule and click Modify.

View alarm rules

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. Click the Instances tab. Then, find the target host and click Alarm Rules in the Actions column to view the alarm rules of the host.
4. To view all the alarm rules, go to the Alarm Rules tab page.

2.6 CloudMonitor Java agent introduction

CloudMonitor provides you with a powerful host monitoring agent that allows you to monitor your servers systematically. The following is a brief introduction to this service, including its installation and resource usage.

Installation path

- **Linux:** `/usr/local/cloudmonit` or
- **Windows:** `C:\Program Files\Alibaba\cloudmonit` or

Process information

After an agent is installed, the following two processes run on your server:

- `/usr/local/cloudmonit` or `/jre/bin/java`
- `/usr/local/cloudmonit` or `/wrapper/bin/wrapper`

Port description

- TCP port 32000 of the local host is accessed and listened to for daemons.
- TCP port 3128, 8080, or 443 of remote servers is accessed for heartbeat monitoring and monitoring data reporting. Port 3128 or 8080 is used for Alibaba Cloud hosts, and port 443 is used for other hosts.
- HTTP port 80 of remote servers is accessed for CloudMonitor agent upgrades.

Agent logs

- Logs of monitoring data are located at `/usr/local/cloudmonit` or `/logs`.
- Logs of startup, shutdown, and daemons are located at `/usr/local/cloudmonit` or `/wrapper/logs`.
- You can modify `/usr/local/cloudmonit` or `/config/log4j.properties` to adjust the log level.

Resource usage

- The process `/usr/local/cloudmonitor/wrapper/bin/wrapper` occupies about 1 MB of memory with little to no CPU usage.
- The process `/usr/local/cloudmonitor/jre/bin/java` occupies about 70 MB of memory and 1% to 2% of one core's CPU usage.
- The installation package is 70 MB and occupies about 200 MB of disk space after the installation is complete.
- Logs use a maximum space of 40 MB and are cleared if they use more than 40 MB.
- Monitoring data is sent every 15 seconds, occupying about 10 KB intranet bandwidth.
- Heartbeat data is sent every three minutes, occupying about 2 KB intranet bandwidth.

External dependencies

- The Java agent of CloudMonitor is built in with JRE 1.8.

- Java service wrapper is used for daemons, start up at boot, and Windows service registration.
- The `ss -s` command is used to capture a TCP connection, and if you do not have this command in the current system, you must install iproute yourself.

Installation instructions

See [Install CloudMonitor Java agent](#).

Install an agent on a host not provided by Alibaba Cloud

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. At the top of the displayed page, click Not Aliyun ecs install. In the Monitor Install Guide dialog box that is displayed, select the agent type and host type to view the corresponding installation method.

2.7 Install CloudMonitor Java agent

Install a CloudMonitor Java agent on Linux

Frequently used commands

```
# Running status
/ usr / local / cloudmonit or / wrapper / bin / cloudmonit or . sh
status

# Start
/ usr / local / cloudmonit or / wrapper / bin / cloudmonit or . sh
start

# Stop
/ usr / local / cloudmonit or / wrapper / bin / cloudmonit or . sh
stop

# Restart
/ usr / local / cloudmonit or / wrapper / bin / cloudmonit or . sh
restart

# Uninstall
/ usr / local / cloudmonit or / wrapper / bin / cloudmonit or . sh
remove && \
rm - rf / usr / local / cloudmonit or
```

Installation command

This command varies by region. Copy the corresponding command and then run it on your server as a root user.

China North 1 (Qingdao) cn-qingdao

```
REGION_ID = cn - qingdao  VERSION = 1 . 3 . 7 \
bash - c "$( curl https :// cms - agent - cn - qingdao . oss
- cn - qingdao - internal . aliyuncs . com / release / cms_instal
l_for_linu x . sh )"
```

China North 2 (Beijing) cn-beijing

```
REGION_ID = cn - beijing  VERSION = 1 . 3 . 7 \
bash - c "$( curl https :// cms - agent - cn - beijing . oss
- cn - beijing - internal . aliyuncs . com / release / cms_instal
l_for_linu x . sh )"
```

China North 3 (Zhangjiakou) cn-zhangjiakou

```
REGION_ID = cn - zhangjiako u  VERSION = 1 . 3 . 7 \
bash - c "$( curl https :// cms - agent - cn - zhangjiako u
. oss - cn - zhangjiako u - internal . aliyuncs . com / release /
cms_instal l_for_linu x . sh )"
```

China North 5 (Hohhot) cn-huhehaote

```
REGION_ID = cn - huhehaote  VERSION = 1 . 3 . 7 \
bash - c "$( curl https :// cms - agent - cn - huhehaote . oss
- cn - huhehaote - internal . aliyuncs . com / release / cms_instal
l_for_linu x . sh )"
```

China East 1 (Hangzhou) cn-hangzhou

```
REGION_ID = cn - hangzhou  VERSION = 1 . 3 . 7 \
bash - c "$( curl https :// cms - agent - cn - hangzhou . oss
- cn - hangzhou - internal . aliyuncs . com / release / cms_instal
l_for_linu x . sh )"
```

China East 2 (Shanghai) cn-shanghai

```
REGION_ID = cn - shanghai  VERSION = 1 . 3 . 7 \
bash - c "$( curl https :// cms - agent - cn - shanghai . oss
- cn - shanghai - internal . aliyuncs . com / release / cms_instal
l_for_linu x . sh )"
```

China South 1 (Shenzhen) cn-shenzhen

```
REGION_ID = cn - shenzhen  VERSION = 1 . 3 . 7 \
```

```
bash -c "$( curl https://cms-agent-cn-shenzhen.oss-cn-shenzhen-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

Hong Kong (China) cn-hongkong

```
REGION_ID = cn-hongkong VERSION = 1.3.7 \
bash -c "$( curl https://cms-agent-cn-hongkong.oss-cn-hongkong-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

US West 1 (Silicon Valley) us-west-1

```
REGION_ID = us-west-1 VERSION = 1.3.7 \
bash -c "$( curl https://cms-agent-us-west-1.oss-us-west-1-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

US East 1 (Virginia) us-east-1

```
REGION_ID = us-east-1 VERSION = 1.3.7 \
bash -c "$( curl https://cms-agent-us-east-1.oss-us-east-1-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

Asia Pacific SE 1 (Singapore) ap-southeast-1

```
REGION_ID = ap-southeast-1 VERSION = 1.3.7 \
bash -c "$( curl https://cms-agent-ap-southeast-1.oss-ap-southeast-1-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

Asia Pacific SE 2 (Sydney) ap-southeast-2

```
REGION_ID = ap-southeast-2 VERSION = 1.3.7 \
bash -c "$( curl https://cms-agent-ap-southeast-2.oss-ap-southeast-2-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

Asia Pacific SE 3 (Kuala Lumpur) ap-southeast-3

```
REGION_ID = ap-southeast-3 VERSION = 1.3.7 \
bash -c "$( curl https://cms-agent-ap-southeast-3.oss-ap-southeast-3-internal.aliyuncs.com/release/cms_instal_l_for_linux.sh )"
```

Asia Pacific SE 5 (Jakarta) ap-southeast-5

```
REGION_ID = ap-southeast-5 VERSION = 1.3.7 \
```

```
bash - c "$( curl https://cms-agent-ap-southeast-5
.oss-ap-southeast-5-internal.aliyuncs.com/release/
cms_instal_l_for_linu_x.sh )"
```

Asia Pacific NE 1 (Tokyo) ap-northeast-1

```
REGION_ID = ap-northeast-1 VERSION = 1.3.7 \
bash - c "$( curl https://cms-agent-ap-northeast-1
.oss-ap-northeast-1-internal.aliyuncs.com/release/
cms_instal_l_for_linu_x.sh )"
```

Asia Pacific SOU 1 (Mumbai) ap-south-1

```
REGION_ID = ap-south-1 VERSION = 1.3.7 \
bash - c "$( curl https://cms-agent-ap-south-1.oss
-ap-south-1-internal.aliyuncs.com/release/cms_instal
l_for_linu_x.sh )"
```

EU Central 1 (Frankfurt) eu-central-1

```
REGION_ID = eu-central-1 VERSION = 1.3.7 \
bash - c "$( curl https://cms-agent-eu-central-1
.oss-eu-central-1-internal.aliyuncs.com/release/
cms_instal_l_for_linu_x.sh )"
```

UK (London) eu-west-1

```
REGION_ID = eu-west-1 VERSION = 1.3.7 \ bash - c "$(
curl https://cms-agent-eu-west-1.oss-eu-west-1-
internal.aliyuncs.com/release/cms_instal_l_for_linu_x
.sh )"
```

Middle East 1 (Dubai) me-east-1

```
REGION_ID = me-east-1 VERSION = 1.3.7 \
bash - c "$( curl https://cms-agent-me-east-1.oss
-me-east-1-internal.aliyuncs.com/release/cms_instal
l_for_linu_x.sh )"
```

China East 1 Finance Cloud (Hangzhou) cn-hangzhou

```
REGION_ID = cn-hangzhou VERSION = 1.3.7 \
bash - c "$( curl https://cms-agent-cn-hangzhou.oss
-cn-hangzhou-internal.aliyuncs.com/release/cms_instal
l_for_linu_x.sh )"
```

China East 2 Finance Cloud (Shanghai) cn-shanghai-finance-1

```
REGION_ID = cn-shanghai-finance-1 VERSION = 1.3.7 \
```

```
bash -c "$( curl https://cms-agent-cn-shanghai-finance-1.oss-cn-shanghai-finance-1-pub-internal.aliyuncs.com/release/cms_instal_l_for_linu_x.sh )"
```

China South 1 Finance Cloud (Shenzhen) cn-shenzen-finance-1

```
REGION_ID = cn-shenzen-finance-1 VERSION = 1.3.7 \
bash -c "$( curl http://cms-agent-cn-shenzhen-finance-1.oss-cn-shenzhen-finance-1-internal.aliyuncs.com/release/cms_instal_l_for_linu_x.sh )"
```

Install a CloudMonitor Java agent on Windows

Installation procedure

1. Download [64-bit agent version](#) or [32-bit agent version](#) based on your operating system version.
2. Create a folder in the path `C:/Program Files/Alibaba` and name it `cloudmonit` or `.`
3. Decompress the installation package to `C:/Program Files/Alibaba/cloudmonit` or `.`
4. Double-click `C:/Program Files/Alibaba/cloudmonit` or `/wrapper/bin/InstallApp-NT.bat` as an administrator to install CloudMonitor.
5. Double-click `C:/Program Files/Alibaba/cloudmonit` or `/wrapper/bin/StartApp-NT.bat` as an administrator to start CloudMonitor.
6. After the installation is complete, you can view, start, and stop CloudMonitor through the service panel of Windows.

Uninstall procedure

1. Stop CloudMonitor through the service panel of Windows.
2. Run `C:/Program Files/Alibaba/cloudmonit` or `/wrapper/bin/UninstallApp-NT.bat` as an administrator to delete CloudMonitor.
3. In the installation directory, delete the entire directory `C:/Program Files/Alibaba/cloudmonit` or `.`

Download the agent with no Internet connection

If you do not have an Internet connection, you can download the installation package from the intranet. For example, if the region of your host is Qingdao and the host uses a 64-bit system, then the intranet download address is as follows: <http://cms->

agent-cn-qingdao.oss-cn-qingdao.aliyuncs.com/release/1.3.7/windows64/agent-windows64-1.3.7-package.zip.

- For a host in another region, change `cn - qingdao` to the corresponding region ID.
- For a host that uses a 32-bit system, change `windows64` to `windows32`.
- For another version, change `1 . 3 . 7` to the corresponding version number.

Security configuration instructions

The following table lists the ports that the CloudMonitor agent uses to interact with your server. If the security software disables these ports, monitoring data may fail to be collected. If your ECS server requires a high level of security, you can add one of the following IP addresses to the whitelist.



Note:

Future version updates and maintenance of CloudMonitor may cause changes to the following IP addresses. To simplify the configuration of your firewall rules, we recommend that you directly allow the 100.100 network segment in the egress direction. This network segment is reserved for the intranet of Alibaba Cloud with no security issues.

Region	IP	Direction	Description
China East 1 (Hangzhou) cn-hangzhou	100.100.19.43:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.45.73:80	Egress	Used to collect monitoring data to CloudMonitor
China North 2 (Beijing) cn-beijing	100.100.18.22:3128	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.18.50:80	Egress	Used to collect monitoring data to CloudMonitor
China North 1 (Qingdao) cn-qingdao	100.100.36.102:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.15.23:80	Egress	Used to collect monitoring data to CloudMonitor
China South 1 (Shenzhen) cn-shenzhen	100.100.0.13:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.0.31:80	Egress	Used to collect monitoring data to CloudMonitor
Hong Kong (China) cn-hongkong	100.103.0.47:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.103.0.45:80	Egress	Used to collect monitoring data to CloudMonitor
China North 5 (Hohhot) cn-huhehaote	100.100.80.135:8080	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.80.12:80	Egress	Used to collect monitoring data to CloudMonitor
China North 3 (Zhangjiakou) cn-zhangjiakou	100.100.80.92:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.0.19:80	Egress	Used to collect monitoring data to CloudMonitor
China East 2 (Shanghai) cn-shanghai	100.100.36.11:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.36.6:80	Egress	Used to collect monitoring data to CloudMonitor
China SW 1 (Chengdu) cn-chengdu	100.100.80.229:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.14:80	Egress	Used to collect monitoring data to CloudMonitor
US East 1 (Virginia) us-east-1	100.103.0.95:3128	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.103.0.94:80	Egress	Used to collect monitoring data to CloudMonitor
US West 1 (Silicon Valley) us-west-1	100.103.0.95:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.29.7:80	Egress	Used to collect monitoring data to CloudMonitor
EU Central 1 (Frankfurt) eu-central-1	100.100.80.241:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.72:80	Egress	Used to collect monitoring data to CloudMonitor
UK (London) eu-west-1	100.100.0.3:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.0.2:80	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 1 (Singapore) ap-southeast-1	100.100.30.20:3128	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.103.7:80	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 2 (Sydney) ap-southeast-2	100.100.80.92:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.13:80 [47.91.39.6:443]	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 3 (Kuala Lumpur) ap-southeast-3	100.100.80.153:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.140:80	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 5 (Jakarta) ap-southeast-5	100.100.80.160:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.180:80	Egress	Used to collect monitoring data to CloudMonitor
Middle East 1 (Dubai) me-east-1	100.100.80.142:8080	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.80.151:80 [47.91.99.5:443]	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific NE 1 (Tokyo) ap-northeast-1	100.100.80.184:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.137:80 [47.91.8.7:443]	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SOU 1 (Mumbai) ap-south-1	100.100.80.152:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.66:80	Egress	Used to collect monitoring data to CloudMonitor

Resource consumption

- Installation package size: 75 MB
- Space occupied after installation: 200 MB
- Memory: 64 MB
- CPU: less than 1%
- Network: intranet, with no Internet bandwidth consumption

FAQ

- Where are CloudMonitor logs saved?
 - **Linux:** `/usr/local/cloudmonit` or `/logs`
 - **Windows:** `C:\Program Files\Alibaba/cloudmonit` or `/logs`

- What should I do if there is a conflict between the port occupied by the agent and the port used by my service?
 1. Change the port range by modifying the CloudMonitor configuration, with the file location: `/usr/local/cloudmonit` or `/wrapper/conf/wrapper.conf`.
 2. Restart CloudMonitor.

```
wrapper . port . min = 40000
wrapper . port . max = 41000
wrapper . jvm . port . min = 41001
wrapper . jvm . port . max = 42000
```

2.8 Introduction to the CloudMonitor GoLang agent

This topic provides a brief introduction to the CloudMonitor GoLang agent and its installation and resource usage. The GoLang agent can enable you to monitor your servers in a centralized and systematic manner.

Installation path

- **Linux:** `/usr/local/cloudmonit` or
- **Windows:** `C:\Program Files\Alibaba\cloudmonit` or

Process information

After the agent is installed, the following two processes run on your server:

- **Linux 32-bit:** `CmsGoAgent.linux-386`
- **Linux 64-bit:** `CmsGoAgent.linux-amd64`
- **Windows 32-bit:** `CmsGoAgent.windows-386.exe`
- **Windows 64-bit:** `CmsGoAgent.windows-amd64.exe`

Port description

- TCP port 3128, 8080, or 443 of remote servers is accessed for heartbeat monitoring and the reporting of monitoring data. Port 3128 or 8080 is used for Alibaba Cloud hosts, and port 443 is used for other hosts.
- HTTP port 80 of remote servers is accessed for CloudMonitor agent upgrades.

Agent logs

- Logs of monitoring data are stored in the log directory.

- You can adjust the level of a log by modifying the `cms . log . level` field in the `config / conf . properties` file. If the field does not exist, you can manually create it. The level of a log can be DEBUG, INFO, WARNING, ERROR, or FATAL.

Resource usage

- The agent process occupies a memory of 10 to 20 MB and 1% to 2% of a single core CPU.
- The size of the agent installation package is 10 to 15 MB.
- Logs use up to 40 MB and are cleared if they use more than 40 MB.
- Monitoring data is sent every 15 seconds, occupying about 10 KB of intranet bandwidth.
- Heartbeat data is sent every 3 minutes, occupying about 2 KB of intranet bandwidth.

Installation instructions

For details, see [Install CloudMonitor GoLang agent](#).

Install the agent on a host not provided by Alibaba Cloud

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Host Monitoring.
3. At the top of the displayed page, click Not Aliyun ecs install. In the Monitor Install Guide dialog box that is displayed, select the agent type and host type to view the corresponding installation method.

2.9 Install CloudMonitor GoLang agent

Requirements on systems

Operating system	Hardware architecture	Note
Windows 7, Windows Server 2008 R2, or later versions	amd64, 386	None
Linux 2.6.23 or later with glibc	amd64, 386	CentOS/RHEL 5.x are not supported.

Resource usage

- Installation package size: 10–15 MB

- **Memory:** 10–15 MB, or 20 MB if you include shared space. Actual numbers vary depending on the size of your system memory.
- **CPU:** 1–2%
- **Network:** intranet. No Internet bandwidth is used.

Install a CloudMonitor GoLang agent on Linux



Note:

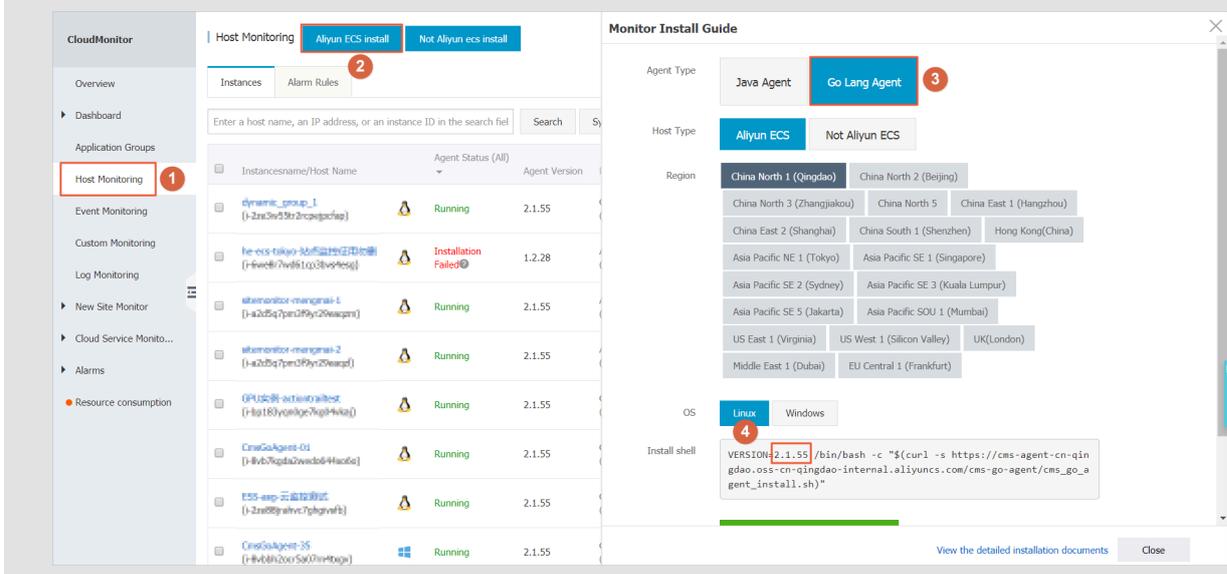
1. The binary file name of the agent

```
CmsGoAgent . linux -${ ARCH }
```

The value of "ARCH" can be "amd64" or "386" depending on the architecture of your Linux system.

2. Version

In this topic, the version 2.1.55 is used. We recommend that you use the latest version. You can find the number of the latest version on the host monitoring page in the CloudMonitor console.



Frequently used commands

```
# Register the agent as a system service .
/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH }
install
# Remove the agent from system services .
/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH }
uninstall
# Start the agent .
/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH } start
# Stop the agent .
/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH } stop
# Restart the agent .
```

```

/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH }
restart
# Uninstall the agent .
/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH } stop
&& \
/ usr / local / cloudmonit or / CmsGoAgent . linux -${ ARCH }
uninstall && \
rm - rf / usr / local / cloudmonit or

```

Installation command

Copy the installation command of the region you require and then run the command on your server with root permissions.



Note:

You can also find the command on the [Monitor Install Guide](#) page in the CloudMonitor console.

China North 1 (Qingdao) cn-qingdao

```

REGION_ID = cn - qingdao  VERSION = 2 . 1 . 55 \
bash - c "$( curl https :// cms - agent - cn - qingdao . oss
- cn - qingdao - internal . aliyuncs . com / cms - go - agent /
cms_go_age nt_install . sh )"

```

China North 2 (Beijing) cn-beijing

```

REGION_ID = cn - beijing  VERSION = 2 . 1 . 55 \
bash - c "$( curl https :// cms - agent - cn - beijing . oss
- cn - beijing - internal . aliyuncs . com / cms - go - agent /
cms_go_age nt_install . sh )"

```

China North 3 (Zhangjiakou) cn-zhangjiakou

```

REGION_ID = cn - zhangjiako u  VERSION = 2 . 1 . 55 \
bash - c "$( curl https :// cms - agent - cn - zhangjiako u .
oss - cn - zhangjiako u - internal . aliyuncs . com / cms - go -
agent / cms_go_age nt_install . sh )"

```

China North 5 (Hohhot) cn-huhehaote

```

REGION_ID = cn - huhehaote  VERSION = 2 . 1 . 55 \
bash - c "$( curl https :// cms - agent - cn - huhehaote . oss
- cn - huhehaote - internal . aliyuncs . com / cms - go - agent /
cms_go_age nt_install . sh )"

```

China East 1 (Hangzhou) cn-hangzhou

```

REGION_ID = cn - hangzhou  VERSION = 2 . 1 . 55 \

```

```
bash -c "$( curl https://cms-agent-cn-hangzhou.oss-cn-hangzhou-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

China East 2 (Shanghai) cn-shanghai

```
REGION_ID = cn-shanghai VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-cn-shanghai.oss-cn-shanghai-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

China South 1 (Shenzhen) cn-shenzhen

```
REGION_ID = cn-shenzhen VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-cn-shenzhen.oss-cn-shenzhen-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Hong Kong (China) cn-hongkong

```
REGION_ID = cn-hongkong VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-cn-hongkong.oss-cn-hongkong-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

US West 1 (Silicon Valley) us-west-1

```
REGION_ID = us-west-1 VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-us-west-1.oss-us-west-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

US East 1 (Virginia) us-east-1

```
REGION_ID = us-east-1 VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-us-east-1.oss-us-east-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Asia Pacific SE 1 (Singapore) ap-southeast-1

```
REGION_ID = ap-southeast-1 VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-ap-southeast-1.oss-ap-southeast-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Asia Pacific SE 2 (Sydney) ap-southeast-2

```
REGION_ID = ap-southeast-2 VERSION = 2.1.55 \  
bash -c "$( curl https://cms-agent-ap-southeast-2.oss-ap-southeast-2-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Asia Pacific SE 3 (Kuala Lumpur) ap-southeast-3

```
REGION_ID = ap-southeast-3 VERSION = 2.1.55 \  

```

```
bash - c "$( curl https://cms-agent-ap-southeast-3.oss-ap-southeast-3-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Asia Pacific SE 5 (Jakarta) ap-southeast-5

```
REGION_ID = ap-southeast-5 VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-ap-southeast-5.oss-ap-southeast-5-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Asia Pacific NE 1 (Tokyo) ap-northeast-1

```
REGION_ID = ap-northeast-1 VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-ap-northeast-1.oss-ap-northeast-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Asia Pacific SOU 1 (Mumbai) ap-south-1

```
REGION_ID = ap-south-1 VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-ap-south-1.oss-ap-south-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

EU Central 1 (Frankfurt) eu-central-1

```
REGION_ID = eu-central-1 VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-eu-central-1.oss-eu-central-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

UK (London) eu-west-1

```
REGION_ID = eu-west-1 VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-eu-west-1.oss-eu-west-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

Middle East 1 (Dubai) me-east-1

```
REGION_ID = me-east-1 VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-me-east-1.oss-me-east-1-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

China East 1 Finance Cloud (Hangzhou) cn-hangzhou

```
REGION_ID = cn-hangzhou VERSION = 2.1.55 \
bash - c "$( curl https://cms-agent-cn-hangzhou.oss-cn-hangzhou-internal.aliyuncs.com/cms-go-agent/cms_go_age_nt_install.sh )"
```

China East 2 Finance Cloud (Shanghai) cn-shanghai-finance-1

```
REGION_ID = cn-shanghai-finance-1 VERSION = 2.1.55 \
```

```
bash -c "$( curl https://cms-agent-cn-shanghai-finance-1.oss-cn-shanghai-finance-1-pub-internal.aliyuncs.com/cms-go-agent/cms_go_agent_install.sh )"
```

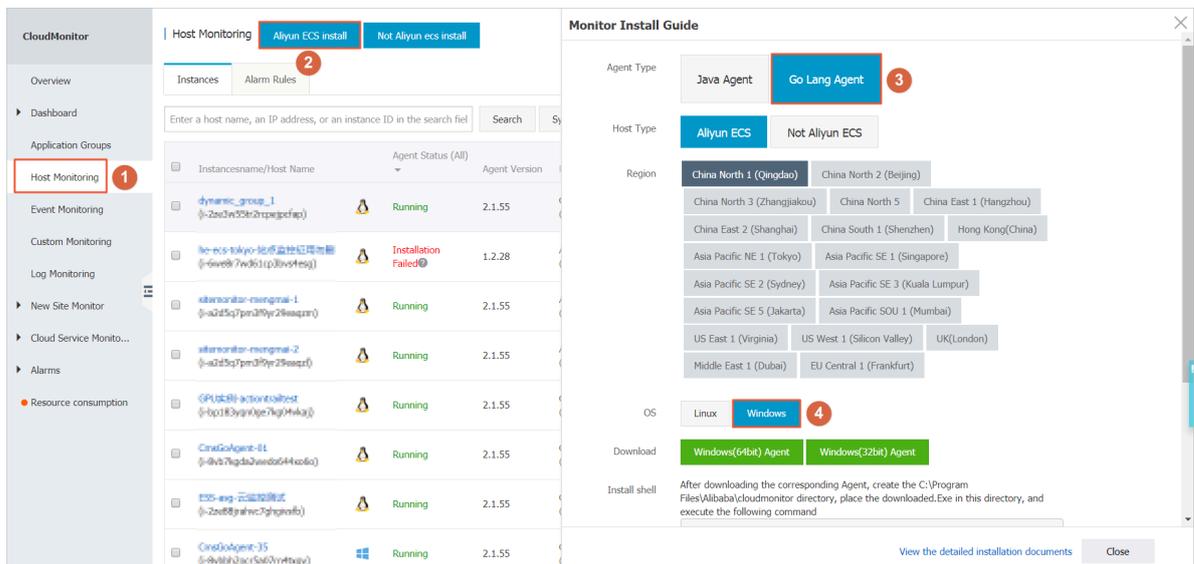
China South 1 Finance Cloud (Shenzhen) cn-shenzhen-finance-1

```
REGION_ID = cn-shenzhen-finance-1 VERSION = 2.1.55 \
bash -c "$( curl http://cms-agent-cn-shenzhen-finance-1.oss-cn-shenzhen-finance-1-internal.aliyuncs.com/cms-go-agent/cms_go_agent_install.sh )"
```

Install a CloudMonitor GoLang agent on Windows

Installation procedure

1. Select your region and host type. Then, depending on your operating system version, download a [64-bit agent version](#) or [32-bit agent version](#) and save it in C:\Program Files\Alibaba\cloudmonitor.



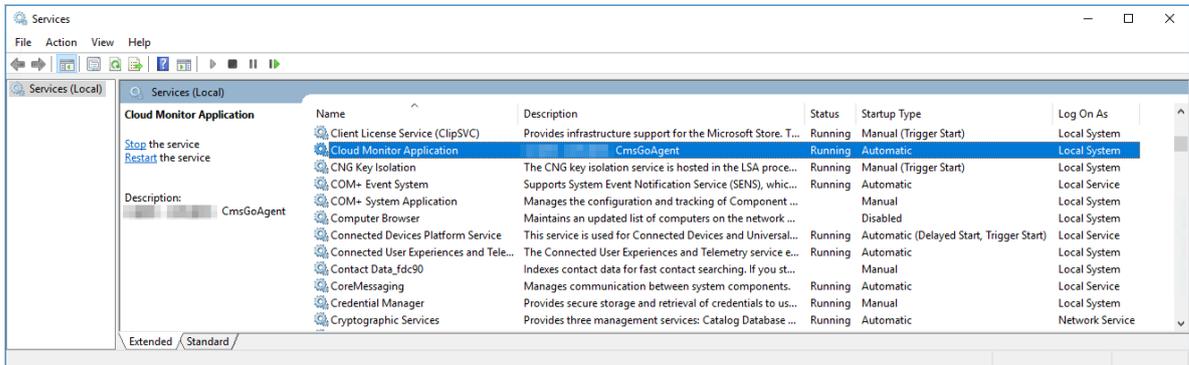
2. Open the Command Prompt as an administrator.

3. Run the following command:

```
cd " C : \ Program Files \ Alibaba \ cloudmonit or "
CmsGoAgent . windows - amd64 . exe install
```

```
CmsGoAgent . windows - amd64 . exe start
```

- After the installation is complete, you can use Windows Services to view, start, and stop the agent.



Uninstall procedure

- Open the Command Prompt as an administrator.
- Run the following command:

```
cd " C :\ Program Files \ Alibaba \ cloudmonit or "
CmsGoAgent . windows - amd64 . exe stop
CmsGoAgent . windows - amd64 . exe uninstall
```

- Close the Command Prompt, and delete the directory `C :\ Program Files \ Alibaba \ cloudmonit or .`

Download the agent with no Internet connection

If you do not have an Internet connection, you can download the installation package from the intranet. For example, if the region of your host is Qingdao and the host uses a 64-bit system, then the intranet download address is as follows: <http://cms-agent-cn-qingdao.oss-cn-qingdao.aliyuncs.com/cms-go-agent/2.1.55/CmsGoAgent.windows-amd64.exe>.

- For a host in another region, change "cn-qingdao" to the corresponding region ID.
- For a host that uses a 32-bit system, change "amd64" to "386".
- For another version, change "2.1.55" to the corresponding version number.

Security configuration instructions

The following table lists the ports that the CloudMonitor agent uses to interact with your server. Note that monitoring data may not be collected if these ports are disabled by security software. Therefore, in the case that your ECS server requires a higher level of security, we recommend that you add one of the following IP addresses to your whitelist.

**Note:**

1. Future maintenance and version updates of CloudMonitor may cause changes to the following IP addresses. Therefore, to simplify the configuration of your firewall rules, we recommend that you directly allow the 100.0.0.0/8 CIDR block in the egress direction. This CIDR block is reserved for the intranet of Alibaba Cloud and is free of security issues.
2. The IP addresses in square brackets ([]) are optional. They can be used as backup addresses in the situation that your network connection is poor.

Region	IP	Direction	Description
China East 1 (Hangzhou) cn-hangzhou	100.100.19.43:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.45.73:80	Egress	Used to collect monitoring data to CloudMonitor
China North 2 (Beijing) cn-beijing	100.100.18.22:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.18.50:80	Egress	Used to collect monitoring data to CloudMonitor
China North 1 (Qingdao) cn-qingdao	100.100.36.102:3128	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.15.23:80	Egress	Used to collect monitoring data to CloudMonitor
China South 1 (Shenzhen) cn-shenzhen	100.100.0.13:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.0.31:80	Egress	Used to collect monitoring data to CloudMonitor
Hong Kong (China) cn-hongkong	100.103.0.47:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.103.0.45:80	Egress	Used to collect monitoring data to CloudMonitor
China North 5 (Hohhot) cn-huhehaote	100.100.80.135:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.12:80	Egress	Used to collect monitoring data to CloudMonitor
China North 3 (Zhangjiakou) cn-zhangjiakou	100.100.80.92:8080	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.0.19:80	Egress	Used to collect monitoring data to CloudMonitor
China East 2 (Shanghai) cn-shanghai	100.100.36.11:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.36.6:80	Egress	Used to collect monitoring data to CloudMonitor
China SW 1 (Chengdu) cn-chengdu	100.100.80.229:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.14:80	Egress	Used to collect monitoring data to CloudMonitor
US East 1 (Virginia) us-east-1	100.103.0.95:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.103.0.94:80	Egress	Used to collect monitoring data to CloudMonitor
US West 1 (Silicon Valley) us-west-1	100.103.0.95:3128	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.29.7:80	Egress	Used to collect monitoring data to CloudMonitor
EU Central 1 (Frankfurt) eu-central-1	100.100.80.241:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.72:80	Egress	Used to collect monitoring data to CloudMonitor
UK (London) eu-west-1	100.100.0.3:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.0.2:80	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 1 (Singapore) ap-southeast-1	100.100.30.20:3128	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.103.7:80	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 2 (Sydney) ap-southeast-2	100.100.80.92:8080	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.80.13:80 [47.91.39.6:443]	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 3 (Kuala Lumpur) ap-southeast-3	100.100.80.153:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.140:80	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SE 5 (Jakarta) ap-southeast-5	100.100.80.160:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.180:80	Egress	Used to collect monitoring data to CloudMonitor
Middle East 1 (Dubai) me-east-1	100.100.80.142:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.151:80 [47.91.99.5:443]	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific NE 1 (Tokyo) ap-northeast-1	100.100.80.184:8080	Egress	Used for monitoring configuration management, and other management and control operations

Region	IP	Direction	Description
	100.100.80.137:80 [47.91.8.7:443]	Egress	Used to collect monitoring data to CloudMonitor
Asia Pacific SOU 1 (Mumbai) ap-south-1	100.100.80.152:8080	Egress	Used for monitoring configuration management, and other management and control operations
	100.100.80.66:80	Egress	Used to collect monitoring data to CloudMonitor

FAQ

- Where are CloudMonitor logs saved?
 - Linux: /usr/local/cloudmonitor/logs
 - Windows: C:\Program Files\Alibaba\cloudmonitor\logs

2.10 Agent release notes

This topic describes the different versions of the host monitoring agent.

2.1.55

Release date: January 24, 2019

Feature optimization and bug fixes:

Fixed the bug that prevented the agent from collecting monitoring data after an ECS instance restarts.

Upgrade recommendations:

If your host runs a Go language agent of a version earlier than 2.1.55, we recommend that you upgrade the agent to this version.

2.1.54

Release date: January 3, 2019

Feature optimization and bug fixes:

Fixed the bug that prevented the agent from collecting monitoring data from graphics processing unit (GPU) servers running a Windows operating system.

Upgrade recommendations:

If your host runs a Go language agent of a version earlier than 2.1.54 on a Windows operating system, we recommend that you upgrade the agent to this version.

2.1.53

Release date: December 25, 2018

Feature optimization and bug fixes:

Fixed the bug that prevented the agent from collecting ECS monitoring data from classic networks.

Upgrade recommendations:

If your host runs a Go language agent of a version earlier than 2.1.53 in a classic network, we recommend that you upgrade the agent to this version.

2.1.51

Release date: December 4, 2018

Feature optimization and bug fixes:

- Fixed the bug that displayed the disk monitoring mount point as a hexadecimal string.
- Pre-check: Check the operating system version, system memory, remaining disk capacity, and connectivity to the CloudMonitor server before installing the agent, to determine whether the agent can be successfully installed.

Upgrade recommendations:

If your host runs a Go language agent of a version earlier than 2.1.50, we recommend that you upgrade the agent to this version.

2.1.50

Release date: November 29, 2018

New features:

The Go programming language version is officially released. Compared with the Java version, the Go programming language version significantly reduces host

performance consumption and provides more stable monitoring services. For more information, see [Introduction to the CloudMonitor GoLang agent](#).

Upgrade recommendations:

If your host runs a Java agent of 1.X.X version, we recommend that you upgrade the agent to this version. On the Host Monitoring page, select a host from the instance list, and click Install Plugins.

1.2.11

New features:

Protocol-dependent local and remote detection through Telnet and HTTP

Feature optimization and bug fixes:

- Fixed the bug that may cause the privilege escalation loophole to occur when the tmp directory is used as the temporary download directory of the installation script.
- Fixed the bug that submitted identical device data when the same disk is attached more than once.
- Fixed the bug that prevented certain processes from obtaining the path and name.
- Optimized the file download method to prevent the download process from blocking the monitoring process.

Upgrade recommendations

When using the local health check function, upgrade the agent to this version.

1.1.64

Feature optimization and bug fixes:

The memory usage collection logic is adjusted. For versions later than CentOS 7.2, the `/proc/meminfo MemAvailable` field is used for available memory estimation to improve the accuracy of memory usage calculation.

Upgrade recommendations:

If your host runs CentOS 7.2 or later, we recommend that you upgrade the agent to this version.

1.1.63

Feature optimization and bug fixes:

- Changed the default wrapper log to the info level.
- Added log information of the error level for easy failure location.
- Fixed the bug that may cause memory leakage for logs at the debug level.

1.1.62

Feature optimization and bug fixes:

- Optimized the HTTP Proxy selection logic to improve the agent installation success rate.
- Added key logs for easy failure location.

1.1.61

Feature optimization and bug fixes:

Fixed the bug that may cause exceptions to occur when certain systems collect process user names, thus causing incorrect topN process collection.

1.1.59

Feature optimization and bug fixes:

- Optimized the process count collection method to improve performance.
- Adjusted process monitoring so that two CloudMonitor agent processes are excluded from process count collection.

3 Site Monitoring

4 Alarm service

4.1 Alarm service overview

You can set alarm rules for metrics in host monitoring, instances in cloud service monitoring, and metrics in custom monitoring. Alarm rules can be applied to all resources, to application groups, or to a single instance.

The alarm service supports alarm notifications through various channels such as emails, TradeManager, and DingTalk chatbots. TradeManager only supports alarm notifications through PC clients. You can also install the Alibaba Cloud app to receive alarm notifications in this method.

Host monitoring alarm rules

Alarm rules can be set for all metrics in host monitoring. Alarm detection frequency can be set to a minimum of once per minute.

Cloud service alarm rule

CloudMonitor allows you to set threshold alarms to monitor the consumption of your cloud resources, and set event alarms to monitor the status of instances and services.

Custom monitoring alarm rules

After reporting monitoring data through the custom monitoring API, you can set alarm rules for corresponding metrics. Then, when the value of a metric exceeds the specified threshold, an alarm is triggered and an alarm notification is sent through the specified notification method.

Custom event alarm rules

After reporting event exceptions through custom event API, you can set alarm rules for the events. Then, when an alarm rule is met, an alarm is triggered and an alarm notification is sent with the specified notification method.

4.2 Use alarm templates

This topic describes how to simplify the creation and management of alarm rules by using alarm templates.

Scenarios

If you have multiple cloud resources (such as ECS instances, RDS services, SLB instances, and OSS buckets), we recommend that you use alarm templates to save alarm rules for these various resources. With having created alarm templates, you can directly apply the templates when creating alarm rules. This process can help you to simplify the creation and management of alarm rules, improving your overall O&M efficiency.

By default, CloudMonitor provides an initialized alarm template that contains common metrics for products such as ECS, RDS, SLB, and OSS, so that you can quickly and easily start to use alarm templates.

Before you begin

Alarm templates are used in combination with application groups. Therefore, we recommend that you create application groups for your resources before you use alarm templates in the creation of related alarm rules. For more information about how to create application groups, see [Create application groups](#).

Create an alarm template



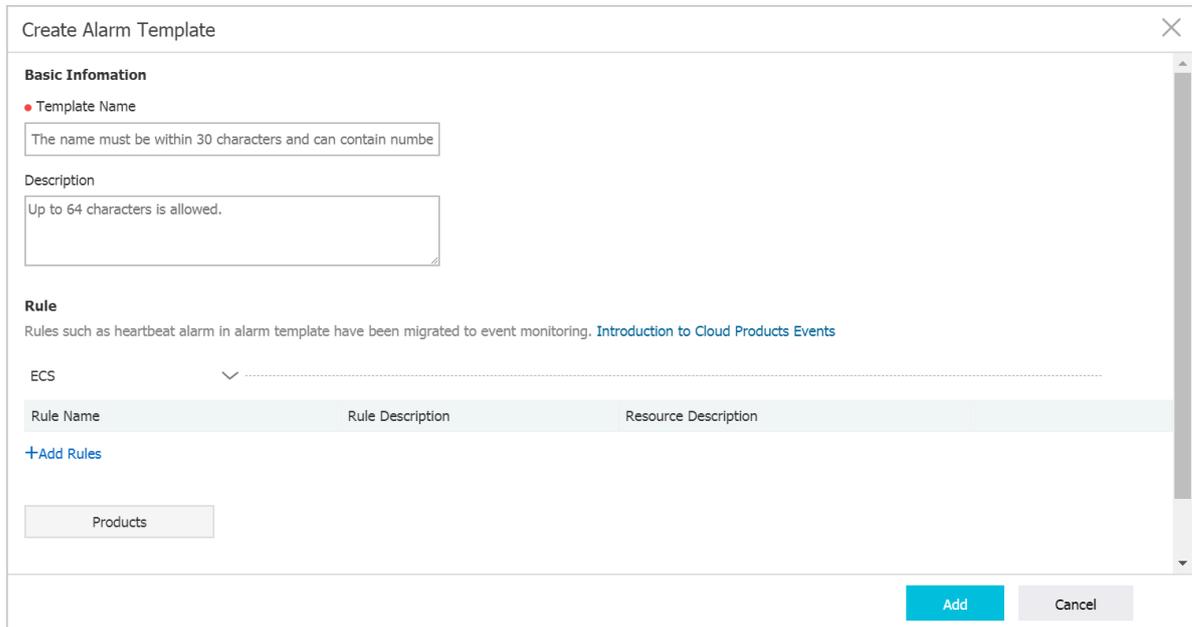
Note:

- Alarm templates can be applied only to application groups.
- Each Alibaba Cloud account can contain up to 100 alarm templates.
- Each alarm template can contain up to 30 metrics.
- The alarm template function is only a shortcut to create multiple alarm rules. Alarm rules are not bound to alarm templates. After an alarm template is modified, alarm rules generated by using this template will remain unchanged. To modify the alarm rules for different application groups in batches, you must apply the modified template to each application group.

Procedure

1. Log on to the [CloudMonitor console](#).

2. In the left-side navigation pane, choose Alarms > Alarm Templates.
3. Click Create Alarm Template to go to the Create Alarm Template page.



4. Enter a Template Name and Description in the Basic Information area.
5. Set an alarm rule. To add more alarm rules, click Add Rules.
6. Click Add.

Use an alarm template

- Use an alarm template when you create an application group

When you create an application group for your resources, you can select an existing alarm template in the MonitorAlarm area. After you have successfully created the application group, CloudMonitor generates alarm rules for this group based on the selected alarm template.

- Apply an alarm template directly to an existing application group

If you have created an application group but have not created alarm rules for the group, you can create an alarm template and then quickly apply the template to the group.

4.3 Alarm rules

4.3.1 Create a threshold alert rule

This topic describes how to create a threshold alert rule so that you can receive alert notification when a metric value reaches the threshold, and handle the exception in a timely manner.

Background information

You can create threshold alert rules to manage and monitor the usage and operation of cloud service resources. When a metric value reaches the threshold, you will receive alert notification. This way, you can trace exceptions as they occur and automate handling of the exceptions in a timely manner.

Prerequisites

We recommend that you create an alert contact and alert contact group before creating a threshold alert rule. When you create an alert rule, you can select the alert contact group as the alert notification receiver. For more information about how to create an alert contact and an alert contact group, see [Create an alert contact and an alert contact group](#).

If you want to use alert callback in alert rules, you must prepare a callback URL that is accessible from the Internet. In addition, you must enable URL callback as a notification method in the existing O&M or message notification system.

Procedure

Precautions

Alert notification can be sent through phone calls, text messages, emails, TradeManager messages, or DingTalk Chatbot. If you want to receive alert notification through multiple methods, be sure to enter correct information when configuring alert contacts.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Alarms > Alarm Rules. On the Alarm Rules page that appears, the Threshold Value Alarm tab is displayed by default.

3. Click Create Alarm Rule. The Create Alarm Rule page is displayed.

The screenshot shows the 'Create Alarm Rule' page with the following configuration:

- 1 Related Resource:**
 - Products: ECS
 - Resource Range: All Resources
- 2 Set Alarm Rules:**
 - Alarm Rule: (empty)
 - Rule Describe: (Agent) Host.cpu.total(Recommend)
 - 1Minute cycle: 1Minute cycle
 - 1 periods: 1 periods
 - Average: Average
 - >=: >=
 - Threshold: Threshold %
- 3 Notification Method:**
 - Notification Contact: Contact Group
 - Selected Groups 0 count

4. Select a resource range, set alert rule parameters, and select a notification method. Click OK.

Parameter description

- **Product** : Select a cloud service, such as ECS, RDS, or OSS.
- **Resource Range** : Select the resources to which the alert rule applies. You can choose from All Resources and Instance .
 - **All Resources** : indicates that the alert rule applies to all your instances of the specified service. For example, you set Resource Range to All Resources and the alert threshold for MongoDB CPU utilization to 80%. The alert rule is hit when the CPU utilization of any of your MongoDB instances is greater than 80%. If you set Resource Range to All Resources , the alert rule is applied to up to 1,000 instances. If there are more than 1,000 instances, you may not receive alert notification when the specified metric reaches the threshold. We recommend that you add resources to service-specific application groups before creating alert rules.
 - **Instance** : indicates that the alert rule only applies to a specific instance. For example, if you set Resource Range to Instance and the alert threshold for ECS CPU utilization to 80%. The alert rule is hit when the ECS CPU utilization of the specified instance is greater than 80%.
- **Rule Name** : the name of the alert rule.

- **Rule Description** : the content of the alert rule. It defines the metric data conditions in which an alert is triggered. For example, the rule description is the average CPU utilization in 5 minutes is greater than or equal to 90%. The rule is hit if the average CPU utilization in 5 minutes is greater than or equal to 90%.

Alert rule example: In ECS monitoring, a data point on the metric of a single server is reported every 15 seconds. 20 data points are reported in 5 minutes.

- Average CPU utilization in 5 minutes greater than 90% means the average value of the 20 data points on CPU utilization that are reported in 5 minutes is greater than 90%.
 - CPU utilization in 5 minutes always greater than 90% means the values of all the 20 data points on CPU utilization that are reported in 5 minutes are greater than 90%.
 - CPU utilization in 5 minutes greater than 90% for once means the value of at least one of the 20 data points on CPU utilization that are reported in 5 minutes is greater than 90%.
 - Total Internet outbound traffic in 5 minutes greater than 50 Mbit/s indicates the sum of the values of the 20 data points on Internet outbound traffic in 5 minutes is greater than 50 Mbit/s.
- **Muted For** : the period when your alert rule is muted so that no new alerts for the rule are sent even if the conditions specified in the rule are met.
 - **Triggered when threshold is exceeded for** : the number of times a rule must be hit before an alert is sent. For example, the rule description is "average CPU utilization in 1 minute greater than 80%, three consecutive times." The rule is hit if the average CPU utilization in 1 minute is greater than 80% for three consecutive times.
 - **Effective Time** : the period when the alert rule is effective. The metric data is checked against the alert rule only when the alert rule is effective.
 - **Notification Object** : the alert contact group to which alert notification is sent.
 - **Alarm Level** :
 - **Email + DingTalk Chatbot**
 - **Email Subject** : the service name + metric name + instance ID by default.

- **Email Remarks** : custom supplementary information of an alert email. The remarks are sent together with the alert notification email.
- **Alarm Callback** : a URL that is accessible from the Internet. CloudMonitor pushes the alert notification to this address by using the POST request. Currently, only HTTP is supported.

More information

- [Create an alert callback](#)
- [Create and use alert templates](#)
- [Use one-click alert](#)

4.3.2 Create an event alert rule

This topic describes how to create an event alert rule so that you can receive alert notification when system exceptions occur to an Alibaba Cloud service and handle the exceptions in a timely manner.

Background information

When an exception occurs to an Alibaba Cloud service, users need to receive alert notification and handle the exception in a timely manner. The CloudMonitor alert service provides the following types of event alert notification so that you can trace exceptions as they occur and automate handling of the exceptions in a timely manner :

- Event alerts can be sent to you through phone calls, text messages, emails, or DingTalk Chatbot.
- Events are distributed to your MNS queue, Function Compute, and URL callback so that you can automate handling of exceptions based on your business scenario.

Prerequisites

We recommend that you create an alert contact and alert contact group before creating an event alert rule. When you create an alert rule, you can select the alert contact group to receive alert notification. For more information about how to create an alert contact and an alert contact group, see [Create an alert contact and an alert contact group](#).

If you want to use alert callback as an alert notification method for system events, you must prepare a callback URL that is accessible from the Internet. In addition, you

must enable URL callback as a notification method in the existing O&M or message notification system.

If you want to use MNS queue or Function Compute as the notification method of a system event, create a message queue or function.

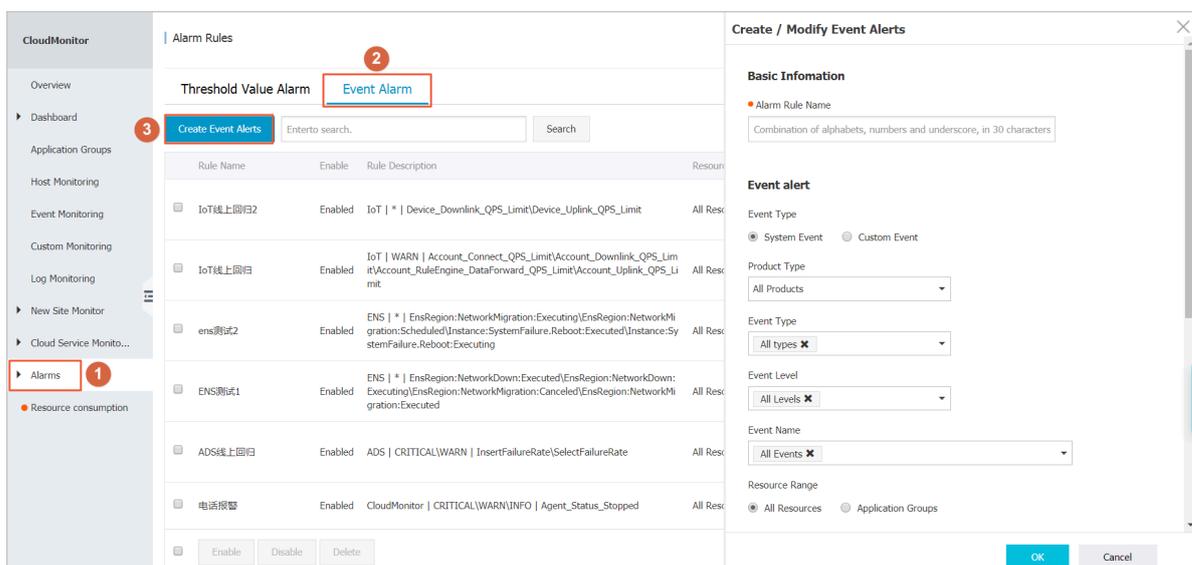
Procedure

Precautions

Events are classified into system events and custom events. The alert rule and notification method vary with event type.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Alarms > Alarm Rules. On the Alarm Rules page that appears, the Threshold Value Alarm tab is displayed by default.
3. Click the Event Alarm tab. On the Event Alarm tab that appears, click Create Event Alarms in the upper-right corner. The Create/Modify Event Alarms dialog box is displayed.



4. In the Basic Information section, enter the alert rule name.

5. Set Event Alarm Rule:
 - a. If you set event type to System Event:
 - Product Type, Event Level, and Event Name: Set these parameters as needed.
 - Resource Range: If you select All Resources, notification is sent based on the configuration for any resource-related events. If you select Application

Group, notification is sent only based on events related to the resources in the specified group.

- b. If you set event type to Custom Event, set Application Group, Event Name, and Rule Description as needed.
6. Set Alarm Type. System events can be distributed to alert notification, MNS queue, Function Compute, and URL callback. Custom events can be distributed to alert notification and alert callback.
7. Click OK.

Subsequent operations

After creating an event alert rule, you can use system event testing to simulate the occurrence of system events. In this way, you can verify whether the MNS queue configured in the alert rule can receive events, and whether the function of Function Compute can be triggered.

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Alarms > Alarm Rules. On the Alarm Rules page that appears, the Threshold Value Alarm tab is displayed by default.
3. Click the Event Alarm tab. The Event Alarm tab that appears shows an alert rule list.
4. Click Test in the Actions column corresponding to an alert rule. The Create Event Test page is displayed.

The screenshot shows the 'Create event test' dialog box in the CloudMonitor console. The dialog box is titled 'Create event test' and has a close button (X) in the top right corner. It contains the following information:

- Product Type: CloudMonitor
- Event Level: CRITICAL
- Event Name: 系统报警消息 (selected from a dropdown menu)
- Content(JSON):


```
{
    "product": "CloudMonitor",
    "resourceId": "acs:ecs:cn-hongkong:1270670679467014:instance(instanceId)",
    "level": "CRITICAL",
    "instanceName": "instanceName",
    "regionId": "cn-hangzhou",
    "name": "Agent_Status_Stopped",
    "content": {
      "ipGroup": "0.0.0.0,0.0.0.1",
      "banjimonVersion": "1.2.11"
    },
    "status": "stopped"
  }
```

At the bottom of the dialog box, there are 'OK' and 'Cancel' buttons.

5. Select an event that you want to test. The event content is displayed. You can modify the fields such as instance ID in the content as needed.

6. Click OK. The system will send an event based on the content, triggering alert notification, MNS queue, Function Compute, or URL callback that you configure in the alert rule.

4.3.3 Manage alarm rules

The alarm service provides powerful capabilities to monitor alarms so that you can easily detect metric exceptions and quickly troubleshoot faults.

Parameter description

- **Products:** ECS, RDS, OSS, among others
- **Resource Range:** The range for which an alarm rule takes effect. There are three alarm rule ranges available: All Resources, Application Group, and Instances. When you set Resource Range to All Resources, you can report an alarm for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold set in your alarm rules. Therefore, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - **All Resources:** Indicates that the specified alarm rule applies to all instances under a user name. For example, if you set the resource range to all resources, and set the alarm threshold for MongoDB CPU usage to 80%, then an alarm is triggered when the CPU usage of any MongoDB instance exceeds 80%.
 - **Application Group:** Indicates that the specified rule applies to all instances under an application group. For example, if you set the resource range to application group and set the alarm threshold for host CPU usage to 80%, then an alarm is triggered when the CPU usage of a host instance exceeds 80%.
 - **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to instances and set the alarm threshold for host CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
- **Alarm Rule:** The alarm rule name.
- **Rule Describe:** The main content of the alarm rule where you define the alarm-triggering condition, or value threshold, for related metrics. For example, if you describe the rule as 1-minute average CPU usage $\geq 90\%$, the alarm service will

check every minute whether the average value of metrics within one minute meets or exceeds 90%.

Consider the following example. For the alarm service in host monitoring, a single server metric item reports one data point in 15 seconds, and 20 data points in five minutes. This relates to the following alarm rules.

- 5-minute average CPU usage > 90%: Indicates that the average CPU usage value of the 20 data points for five minutes exceeds 90%.
- 5-minute CPU usage always > 90%: Indicates that the CPU usage values of the 20 data points for five minutes all exceed 90%.
- 5-minute CPU usage once > 90%: Indicates that the CPU usage value of at least one of the 20 data points for five minutes exceeds 90%.
- Total 5-minute Internet outbound traffic > 50 MB: Indicates that the sum of the outbound traffic values of the 20 data points for five minutes exceeds 50 MB.
- Triggered when threshold is exceeded for: An alarm notification is sent if the detected values reach the alarm rule threshold multiple times in a row.
- Effective Period: the period of time for which an alarm rule is valid. The alarm service checks metrics and determines whether to generate an alarm only during this period of time.
- Alarm Contact: a group of contacts who receive alarm notifications.
- Notification Methods: Different notification methods are available based on different alarm levels. Three alarm levels are available: Critical, Warning, and Info.
 - Critical: voice calls, SMS messages, emails, and DingTalk chatbot
 - Warning: SMS messages, emails, and DingTalk chatbot
 - Info: emails and DingTalk chatbot
- Email Remark: supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email.

Manage alarm rules

CloudMonitor provides three alarm rule management portals: the application group page, metric list page, and alarm rule list page of the alarm service.

- For alarm rule management in application groups, see [Manage alarm rules](#).
- For alarm rule management in host monitoring, see [Manage alarm rules](#).
- For setting alarm rules in custom monitoring, see [Set alarm rules](#).

- You can also set alarm rules in cloud service monitoring.

4.3.4 Create an alert callback

This topic describes how to create an alert callback to integrate CloudMonitor alerts to your existing O&M or message system.

Background information

CloudMonitor provides the alert callback feature for alert notification in addition to the methods such as emails, and DingTalk Chatbot. Alert callback allows O&M engineers and developers to handle alert events flexibly.

CloudMonitor pushes alerts to a specified Internet URL through HTTP POST requests . You can take actions based on received notification.

Prerequisites

- You have a callback URL that is accessible through the Internet.
- URL callback is enabled as an alert notification method in your existing O&M or message system.

Procedure

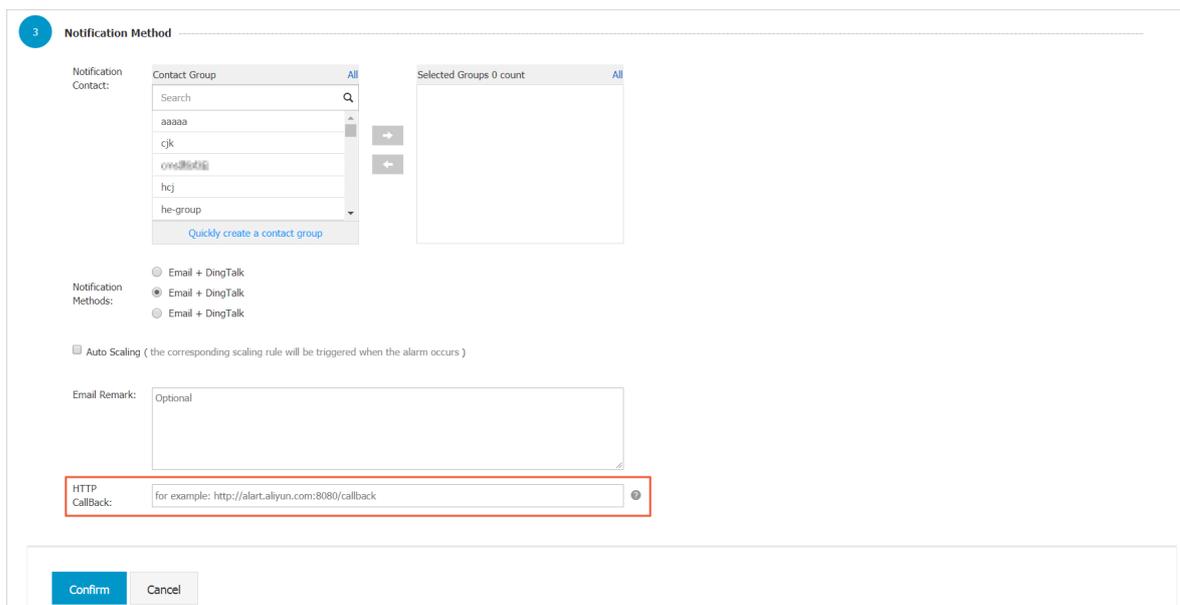
Precautions

- According to the retry policy of alert callback, the number of retries is 3 and the timeout period is 5 seconds.
- Currently, only HTTP is supported.

Procedure

1. Log on to the [CloudMonitor console](#).

2. Modify an existing alert rule by creating a callback or create an alert rule.



3. In the notification method section, enter the URL address for alert callback and click OK. When an alert rule is triggered, CloudMonitor sends an alert to your specified URL.

Callback parameters

The following table lists the content of a POST request that is pushed when an alert rule calls back a URL.

Parameter	Data type	Description
userId	String	The user ID.
alertName	String	The alert name.
timestamp	String	The time stamp when the alert is generated.
alertState	String	The alert state. One of the following states is returned: OK, ALERT, and INSUFFICIENT_DATA.
dimensions	String	The object that has triggered the alert. For example: [{"userId":"12345","instanceId":"i-12345"}]

Parameter	Data type	Description
expression	String	The alert conditions. For example, [{"expression": "\$value > 12", "level": 4, "times": 2}] indicates that an alert is triggered when the threshold value is greater than 12 for two consecutive times. If the value of level is 4, an alert is sent to you through an email. If the value of level is 3, an alert is sent to you through a text message and an email. The times field indicates the number of consecutive times of reaching the alert threshold that you selected when configuring the alert rule.
curValue	String	The current value of the metric when an alert is triggered or cleared.
metricName	String	The metric name.
metricProject	String	The service name. For more information about the metric and service names, see Preset metrics reference .

An example of a POST request is as follows:

```
{
  "userId ":" 12345 ",
  "alertName ":" putNewAlarm_group_a3_7cd898 - ea6b - 4b7b -
a8a8 - de017a8327_f6 ",
  "timestamp ":" 1508136760 ",
  "alertState ":" ALARM ",
  "dimensions ":[
    {
      "userId ":" 12345 ",
      "instanceId ":" i - 12345 "
    }
  ],
  "expression ":[{"expression \":" $ Average > 90 \", \" level \":
4, \" times \": 2 }]",
  "curValue ":" 95 ",
```

```
" metricName ":" CPUUtiliza tion ",
" metricProj ect ":" acs_ecs_da shboard "
}
```

4.4 Alarm contacts

4.4.1 Create an alert contact and an alert contact group

This topic describes how to create an alert contact and an alert contact group, and add the contact to the group. With this configuration, contacts in the alert contact group can receive alert notification.

Background information

CloudMonitor sends alert notification based on alert contacts and contact groups. To receive alert notification, you must first create an alert contact and an alert contact group, add the contact to the contact group, and then select the contact group when creating an alarm rule.

An alert contact group is a group of one or more alert contacts. An alert contact can be added to multiple alert contact groups. In alert rule configurations, alert notification receivers are alert contact groups, not alert contacts.

Prerequisites

An alert contact is created and the contact information is correct.

Procedure

Precautions



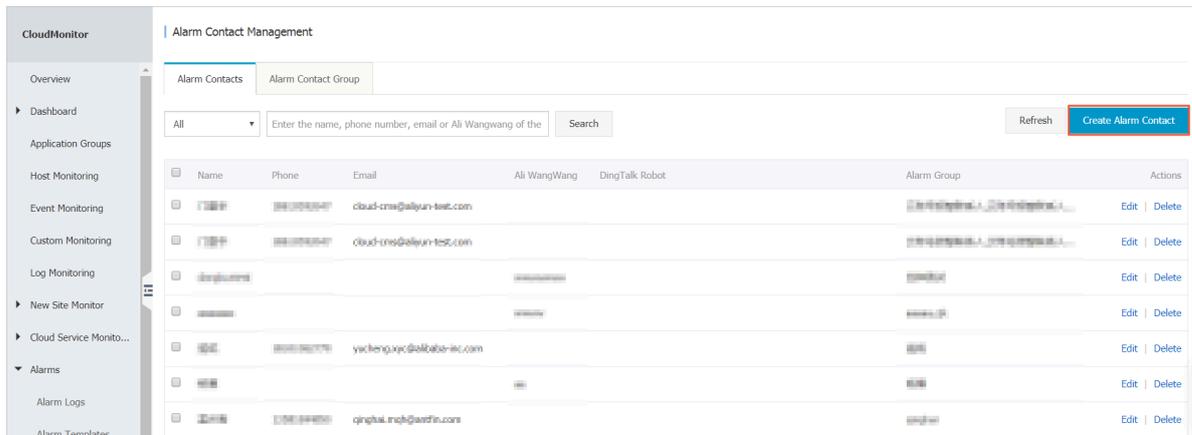
Note:

You must verify the email address to ensure that it can receive alert notification.

Create an alert contact

1. Log on to the [CloudMonitor console](#).

2. In the left-side navigation pane, choose Alarms > Alarm Contact. The Alarm Contact Management page is displayed.



3. Click Create Alarm Contact in the upper-right corner. In the dialog box that appears, enter your name and email address.

The 'Set Alarm Contact' dialog box contains the following fields and controls:

- Name:** Input field with a note: "The name must be 2-40 characters, can include English letters, numbers, . , and underscores, and should start with a Chinese or English character."
- Phone:** Input field with a "Send verification code." button.
- Verification code:** Input field with a note: "Fill in the phone verification code."
- Email ID:** Input field with a "Send verification code." button.
- Verification code:** Input field with a note: "Fill in the E-mail verification code."
- Ali WangWang:** Input field.
- DingTalk Robot:** Input field with a link: "How to get the DingTalk robot address".
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

4. After the information is verified, click Save.

Create an alert contact group

1. Log on to the [CloudMonitor console](#).

2. In the left-side navigation pane, choose Alarms > Alarm Contact. The Alarm Contact Management page is displayed.
3. Click the Alarm Contact Group tab.
4. On the tab that appears, click Create Alarm Contact Group in the upper-right corner. The Create Alarm Contact Group dialog box is displayed.

Create Alarm Contact Group

Group Name:

Description:

Select contacts:

Existing Contacts (Create Alarm Contact) All

Selected Contacts All

Enter the contact name

You have selected 0 contacts.

OK Cancel

5. Enter the group name, select the contacts to be added to the group, and click OK.

Add multiple contacts to a contact group

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Alarms > Alarm Contact. The Alarm Contact Management page is displayed.
3. Select the contacts that you want to add from the alert contact list.
4. Click Add to a Contact Group at the bottom of the list.
5. In the dialog box that appears, select a contact group and click OK.

4.4.2 Manage alarm contacts and alarm contact groups

Alarm notifications are sent to alarm contacts and alarm contact groups. When creating an alarm rule, you will need to create an alarm contact and an alarm

contact group so that you can select the contact and contact group to receive alarm notifications.

Manage an alarm contact

You can create, edit, or delete contact information, such as an email address.

- Create an alarm contact.

1. Log on to the [CloudMonitor Console](#).
2. In the left navigation pane, click Alarm contacts under Alarms. The Alarm Contact Management page is displayed.
3. Click Create Alarm Contact in the upper-right corner of the page. In the displayed dialog box, enter the contact email address and other information.

The specified email address needs to be verified so that you can avoid entering incorrect information that may cause you to not receive alarm notifications.

- Edit an alarm contact.

1. Log on to the [CloudMonitor Console](#).
2. In the left navigation pane, click Alarm contacts under Alarms. The Alarm Contact Management page is displayed.
3. Click Edit in the Actions column to edit the contact information.

- Delete an alarm contact.

1. Log on to the [CloudMonitor Console](#).
2. In the left navigation pane, click Alarm contacts under Alarms. Alarm Contact Management page is displayed.
3. Click Delete in the Actions column.



Note:

Once you delete an alarm contact, CloudMonitor alarm notifications are not longer sent to that contact.

Manage an alarm contact group

An alarm contact group may contain one or more alarm contacts. The same alarm contact can be added to multiple alarm contact groups. , When setting alarm rules, all alarm notifications need to be sent through an alarm contact group.

- Create an alarm contact group.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left navigation pane, click Alarm contacts under Alarms. The Alarm Contact Management page is displayed.
 3. Click the Alarm Contact Group tab at the top of the page to switch to the alarm contact group list.
 4. Click Create Alarm Contact Group in the upper-right corner to open the Create Alarm Contact dialog box.
 5. Enter a group name and select the contacts you want to add to the group.
- Edit an alarm contact group.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left navigation pane, click Alarm contacts under Alarms. The Alarm Contact Management page is displayed.
 3. Click the Alarm Contact Group tab at the top of the page to switch to the alarm contact group list.
 4. Click Edit in the Actions column to edit the contact group information.
- Delete an alarm contact group.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left navigation pane, click Alarm contacts under Alarms. The Alarm Contact Management page is displayed.
 3. Click the Alarm Contact Group tab at the top of the page to switch to the alarm contact group list.
 4. Click Delete in the Actions column to delete the contact group.
- Add contacts to a contact group in batches.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left navigation pane, click Alarm contacts under Alarms. The Alarm Contact Management page is displayed.
 3. Select the contacts that you want to add from the alarm contact list.
 4. Click Add to a contact group at the bottom of the page.
 5. In the displayed dialog box, select the target contact group and click OK.

4.5 Use one-click alert

This topic describes how to use the one-click alert function to enable key metric alerts with a single click.

Background information

One-click alert allows you to enable key metric alerts with a single click. One-click alert is designed for inexperienced cloud service developers and O&M engineers . It helps them quickly establish a basic monitoring and alert system on the cloud without the need for a wide range of knowledge on cloud services and metrics. With this system, the engineers can receive alert notification on exceptions for key metrics.

Prerequisites

Before using one-click alert, you must understand the services that support this function and related alert rules.

Service name	Metric name	Rule description
ECS	CPUUtilization	Maximum value in 1 minute greater than 90%, five consecutive times, 1-hour mute duration, email notification
	vm.DiskUtilization	Maximum value in 1 minute greater than 90%, five consecutive times, 1-hour mute duration, text message and email notification
	vm.MemoryUtilization	Maximum value in 1 minute greater than 90%, five consecutive times, 1-hour mute duration, email notification
	InternetOutRate_Percent	Maximum value in 1 minute greater than 90%, five consecutive times, 1-hour mute duration, email notification

Service name	Metric name	Rule description
RDS	CpuUsage	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	DiskUsage	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, text message and email notification
	IOPSUsage	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ConnectionUsage	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	DataDelay	Maximum value in 5 minutes greater than 5, five consecutive times, 1-hour mute duration, email notification
SLB	DropConnection	Maximum value in 1 minute greater than 0, five consecutive times, 1-hour mute duration, email notification
	DropTrafficRX	Maximum value in 1 minute greater than 0, five consecutive times, 1-hour mute duration, email notification

Service name	Metric name	Rule description
	DropTrafficTX	Maximum value in 1 minute greater than 0, five consecutive times, 1-hour mute duration, email notification
ApsaraDB RDS for Redis	CpuUsage	Maximum value in 1 minute greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ConnectionUsage	Maximum value in 1 minute greater than 80%, five consecutive times, 1-hour mute duration, email notification
	MemoryUsage	Maximum value in 1 minute greater than 80%, five consecutive times, 1-hour mute duration, email notification
	IntranetInRatio	Maximum value in 1 minute greater than 80%, five consecutive times, 1-hour mute duration, email notification
	IntranetOutRatio	Maximum value in 1 minute greater than 80%, five consecutive times, 1-hour mute duration, email notification
ApsaraDB RDS for MongoDB (replica set)	CPUUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	MemoryUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification

Service name	Metric name	Rule description
	DiskUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	IOPSUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ConnectionUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
ApsaraDB RDS for MongoDB (sharded cluster)	ShardingCPUUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ShardingMemoryUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ShardingDiskUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ShardingIOPSUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification
	ShardingConnectionUtilization	Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification

Service name	Metric name	Rule description
ApsaraDB RDS for HBase	LoadPerCpu	Maximum value in 5 minutes greater than 3, three consecutive times, 1-hour mute duration, email notification
	cpu_idle	Maximum value in 5 minutes smaller than 10, three consecutive times, 1-hour mute duration, email notification
	compactionQueueSize	Maximum value in 5 minutes greater than 2,000, three consecutive times, 1-hour mute duration, email notification
	rs_handlerQueueSize	Maximum value in 5 minutes greater than 1,000, three consecutive times, 1-hour mute duration, email notification
	CapacityUsedPercent	Maximum value in 5 minutes greater than 80%, three consecutive times, 1-hour mute duration, email notification
	zookeeper_tcp_count	Maximum value in 5 minutes greater than 2,000, three consecutive times, 1-hour mute duration, email notification
Elasticsearch	ClusterStatus	Maximum value in 1 minute greater than 2, ten consecutive times, 1-hour mute duration, email notification
	NodeDiskUtilization	Maximum value in 1 minute greater than 75%, ten consecutive times, 1-hour mute duration, email notification

Service name	Metric name	Rule description
	NodeHeapMemoryUtilization	Maximum value in 1 minute greater than 85%, ten consecutive times, 1-hour mute duration, email notification
Open Search	DocSizeRatioByApp	Maximum value in 10 minutes greater than 85%, one time, 1-hour mute duration, email notification
	ComputeResourceRatioByApp	Maximum value in 10 minutes greater than 85%, one time, 1-hour mute duration, email notification

Procedure

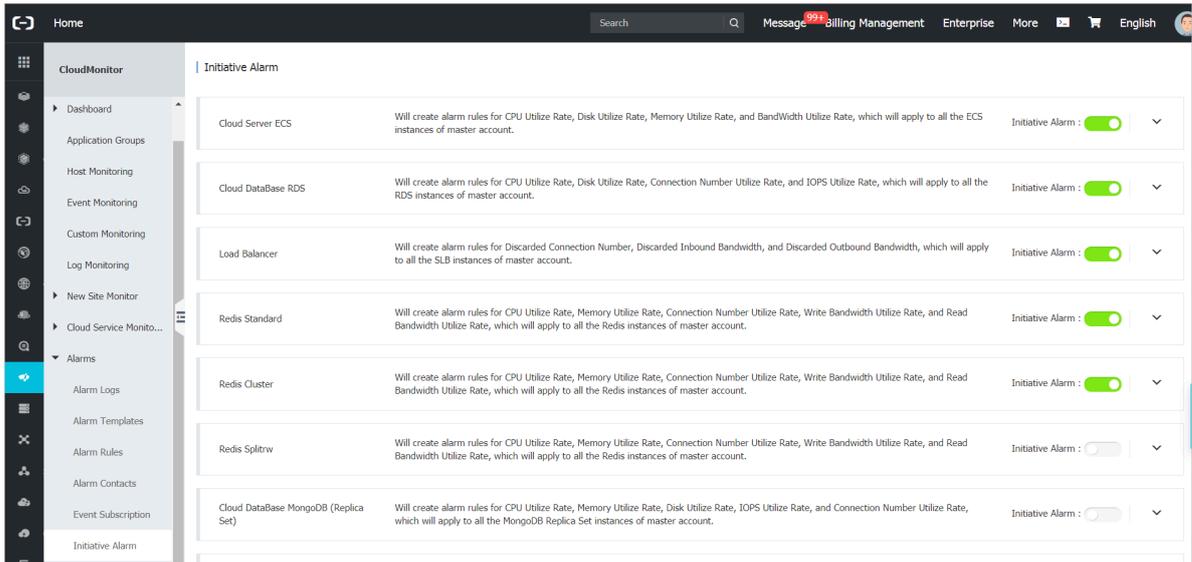
Precautions

- When one-click alert is enabled, the built-in alert rules of CloudMonitor are enabled by default. An alert system is quickly established to monitor key metrics, not all metrics.
- When one-click alert is enabled, the corresponding alert rules apply to the existing and to-be-created instances of the selected services.
- One-click alert allows you to modify, disable, and delete built-in alert rules.

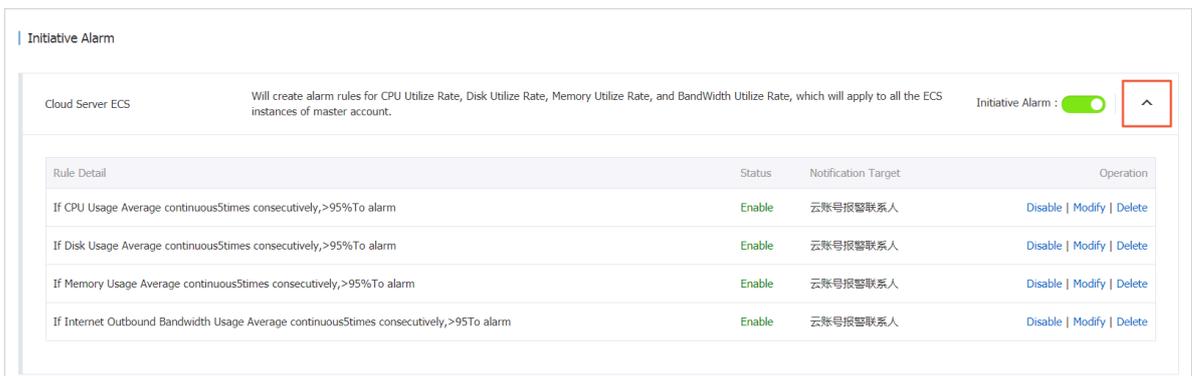
Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Alarms > One-click Alarm. The One-click Alarm page is displayed.

3. Turn on One-click Alarm corresponding to the cloud service for which you want to enable alert notification.



4. Click the drop-down arrow to the right of the One-click Alarm switch to view the alert rules that are automatically generated by CloudMonitor.



5. (Optional) You can click Disable, Modify, or Delete in the Actions column corresponding to an alert rule to disable, modify, or delete the rule.

5 Availability monitoring

5.1 Create an availability monitoring task

This topic describes how to create availability monitoring tasks to quickly identify situations where local or dependent remote services are unresponsive.

Background information

CloudMonitor availability monitoring helps you quickly identify situations where local or remote hosts are unresponsive. Alerts are sent if the service does not respond within the specified timeout period or when an error status code is returned.

Prerequisites

- A group is created for the resources for availability monitoring. For more information, see [Create application groups](#).
- The CloudMonitor agent is installed for the monitored host. For more information, see [Introduction to the CloudMonitor GoLang agent](#).

Procedure

Precautions



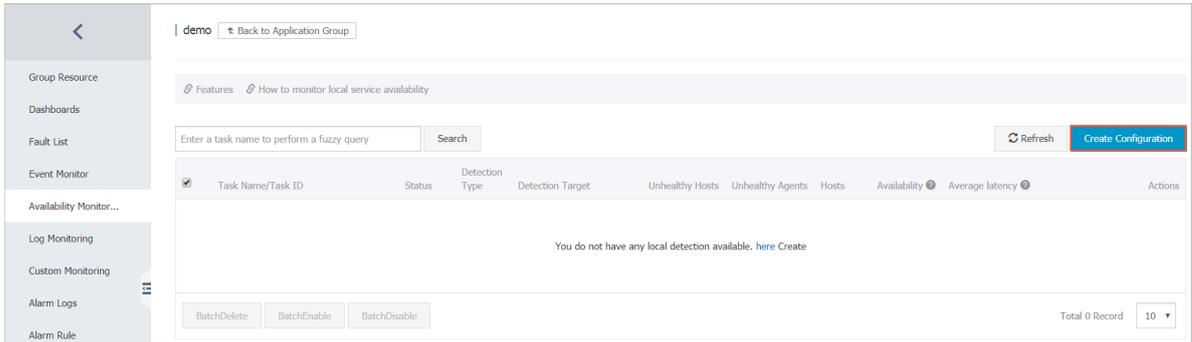
Note:

- Availability monitoring depends on the CloudMonitor agent. Ensure that the CloudMonitor agent has been installed on the monitored host.
- Monitoring is performed once a minute.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups. The Application Groups page is displayed.
3. Click the name of the application group for which an availability monitoring task is to be created. The application group details page is displayed.

4. In the left-side navigation pane, click Availability Monitoring. The Availability Monitoring page is displayed.



- Click **Create Configuration** in the upper-right corner. The **Create Availability Monitoring** page is displayed.

CreateAvailability Monitoring
✕

1 Monitoring Configurations

* Task Name :
Enter 4 to 50 characters. Only English letters, numbers, underlines, and Chinese characters are allowed.

* Target Server : All

91a2a7b442ac
cmastomule011175066070.usd3
cmastomule000189016219.usd2
cmastomule011175064120.usd3
cmastomule0147282800446c.usd3

* Detection Type :

* Detection Target :

* Request Method : HEAD GET POST

[Advanced Configuration](#) ▾

2 Alarm Configuration

Status Code : [Status Code Description](#)

Response Time :

Notification Method : Email + DingTalk [?](#)
 Email + DingTalk
 Email + DingTalk

[Advanced Configuration](#) ▾

The detection period is 1 minute. When the above alarm configurations are met, any server will send an alarm notification to the contact group associated with the application group.

- Set **Task Name** and **Target Server**. You can configure the same detection rule to be applied to all hosts in the group or just a portion of hosts in the group.
- Set **Detection Type** to **URL or IP Address**, **ApsaraDB for RDS**, or **ApsaraDB RDS for Redis**. Set **Detection Target**.
 - If you set **Detection Type** to **URL or IP Address**, you can set **Detection Target** to **HTTP(S)**, **TELNET**, or **PING**. If you set **Detection Target** to **HTTP(S)**, you can set

Request Method to HEAD, GET, or POST. You can also configure the returned value.

- If you set Detection Target to ApsaraDB for RDS or ApsaraDB RDS for Redis, the related instances in your group and their access addresses are displayed.

8. In the Alert Configuration section, set Status Code and Response Time. An alert is triggered if the conditions of either parameter are met. Alerts are sent to the contact group of the corresponding application group.

- **Status Code:** An alert is triggered if the returned status code meets the conditions set in the alert rule.
- **Response Time:** An alert is triggered if the response time meets the conditions set in the alert rule.
- **Notification Method:** the method by which alerts are sent.
- **Advanced Configuration:** You can configure Muted For and Effective From. Muted For is a period when your alert rules are muted so that no alerts are sent even if the conditions specified in your alert rules are met. Effective From is a period when the alert rules are effective. Alerts are sent if the conditions specified in your alert rules are met during this period.

9. Click OK.

5.2 Manage availability monitoring

Availability monitoring conducts periodical detection tasks to check whether specified local or remote paths or ports respond properly and sends alarm notifications if response timeouts occur or status codes indicate errors based on the conditions specified in your alarm rules. This function can help you to quickly learn if local or remote services are unresponsive or abnormal, improving overall O&M and management efficiency.

Viewing availability monitoring tasks

1. Log on to the [CloudMonitor Console](#).
2. Click Application Groups in the left-hand navigation bar to go to the application groups page.
3. Select the application groups for which you want to view availability monitoring, then click the application group name to enter the application group details page.

4. Select Availability Monitoring from the left-side navigation pane to go to the Availability Monitoring page. A list displaying the tasks that apply all availability monitoring in the group is displayed.

View monitoring results

1. Log on to the [CloudMonitor Console](#).
2. Click Application Groups in the left-hand navigation bar to go to the Application Groups page.
3. Select the Application Groups for which you want to view availability monitoring, then click the application group name to enter the application groups details page.
4. Select Availability Monitoring from the left-side navigation pane to go to the Availability Monitoring page.
5. You can view monitoring results in the list.
 - When the task probe does not trigger an alarm, the number of faulty instances in the list is 0.
 - When an alarm is triggered for a probe exception, the number of instances that triggered an alarm is displayed in the list, click exception numbers to view the faulty instance details.
 - Exception details.

Modify availability monitoring tasks

1. Log on to the [CloudMonitor Console](#).
2. Click Application Groups in the left-hand navigation bar to go to the Application Groups page.
3. Select the Application Groups that needs to modify the availability monitoring, click the application group name to go to the app grouping details page.
4. Select availability monitoring on the left-hand menu of the page to enter the management page for availability monitoring.
5. Select the task that needs to be modified, click Modify in the action to go to the modify application groups page.
6. Edit content on the modify application groups page and save the configuration.

View alarm logs

1. Log on to the [CloudMonitor Console](#).

2. Click Application Groups in the left-hand navigation bar to go to the Application Groups page.
3. Select the application groups that needs to view the alarm logs, click the application group name to go to the application group details page.
4. Select Alarm Logs on the left-hand menu of the page, and go to the alarm logs page to view the alarm log details.

Enable or disable monitoring tasks

Enabling or disabling monitoring tasks is supported for local health checks. When a task is disabled, health checks are no longer performed and alarms are no longer triggered for the task. However, when a task is enabled, probing is re-started and alarms will be triggered when the conditions specified in alarm rule settings are met.

1. Log on to the [CloudMonitor Console](#).
2. Click Application Groups in the left-hand navigation bar to go to the Application Groups page.
3. Select the application groups that needs to be enabled or disabled for availability monitoring, and click the application group name, enter the application group details page.
4. Select availability monitoring on the left-hand menu of the page to enter the task management page for availability monitoring.
5. Select the task that you want to enable or disable, and click enable or disable in the action to modify the task status.

5.3 Local service availability monitoring

This topic describes how to configure local service availability monitoring so that you can receive alert notification if the service does not respond within a specified timeout period or when an error status code is returned.

Background information

Local service availability monitoring sends alert notification to identify situations where local services are unresponsive or when an error status code is returned.

Prerequisites

- Local service availability monitoring depends on the CloudMonitor agent. Ensure that the CloudMonitor agent has been installed on the monitored host. For more information, see [Introduction to the CloudMonitor GoLang agent](#).
- Before you use local service availability monitoring, you must [Create application groups](#).

Procedure

Precautions



Note:

- Local service availability monitoring depends on the CloudMonitor agent to run properly. Ensure that the CloudMonitor agent has been installed on the monitored host.
- An availability test is performed once a minute.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click **Application Groups**. The **Application Groups** page is displayed.
3. Click the name of the application group for which you want to create a local service availability monitoring task. The application group details page is displayed.
4. In the left-side navigation pane, click **Availability Monitoring**. The **Availability Monitoring** page is displayed.

5. Click Create Configuration in the upper-right corner. The Create Availability Monitoring page is displayed.

CreateAvailability Monitoring
✕

1 Monitoring Configurations

* Task Name :
Enter 4 to 50 characters. Only English letters, numbers, underlines, and Chinese characters are allowed.

* Target Server : All

91a2e7b442ac
cmastomule011175066070.usd3
cmastomule000189016219.usd2
cmastomule011175064120.usd3
cmastomule014722200446c.usd3

* Detection Type :

* Detection Target :

* Request Method : HEAD GET POST

[Advanced Configuration](#) ▾

2 Alarm Configuration

Status Code : [Status Code Description](#)

Response Time :

Notification Method : Email + DingTalk [?](#)
 Email + DingTalk
 Email + DingTalk

[Advanced Configuration](#) ▾

The detection period is 1 minute. When the above alarm configurations are met, any server will send an alarm notification to the contact group associated with the application group.

6. Set Task Name and Target Server. You can configure the same detection rule to be applied to all hosts in the group or just a portion of hosts in the group.
7. Set Detection Type to URL or IP Address, ApsaraDB for RDS, or ApsaraDB RDS for Redis. Set Detection Target.
8. In the Alert Configuration section, set Status Code and Response Time. An alert is triggered if the conditions of either parameter are met. Alerts are sent to the contact group of the corresponding application group.

9. Click OK. If your service does not respond within the timeout period, you will receive alert notification through text messages, emails, or other channels.
- 10.(Optional) The availability monitoring list displays the number of unhealthy hosts. Click Unhealthy Hosts to view the details of the abnormal hosts.

Parameter description

- **Monitoring Configuration section:**

- **Target Server:** the host that initiates the test. Target Server and Detection Target are the same host.
- **Detection Type:** Select `URL` or `IP Address`.
- **Detection Target:** If you select `HTTP (S)`, enter the target address in the format of `localhost : port / path`. If you select `TELNET`, enter the target address in the format of `127 . 0 . 0 . 1 : port`. For example, to test whether Tomcat responds normally, select `HTTP (S)` and enter `localhost : 8080 / monitor`. To test the connectivity of MySQL, select `TELNET` and enter `127 . 0 . 0 . 1 : 3306`.

- **Alarm Configuration section:**

Both Status Code and Response Time are used as the metrics of availability monitoring. An alert is triggered if either metric value reaches the specified threshold. Alerts are sent to the contact group of the corresponding application group. For local availability monitoring, set the status code greater than 400.

- **Status Code:** An alert is triggered if the returned status code meets the conditions set in the alert rule.
- **Notification Method:** the method by which alerts are sent.
- **Advanced Configuration:**
 - **Muted For:** a period when your alert rules are muted so that no alerts are sent even if the conditions specified in your alert rules are met.
 - **Effective From:** a period when the alert rules are effective. Alerts are sent if the conditions specified in your alert rules are met during this period. You can configure these parameters based on your actual needs.

5.4 Status codes

The following is a list of the custom status codes returned whenever an exception is detected after an availability check is completed.

Protocol type	Status code	Definition
HTTP	610	Timeout due to no response within 5 seconds after the HTTP request was issued.
HTTP	611	The detection failed.
Telnet	630	Timeout due to no response within 5 seconds.
Telnet	631	The detection failed.

6 Cloud service monitoring

6.1 ApsaraDB for RDS

By monitoring multiple metrics of ApsaraDB for Relational Database Service (RDS), such as disk usage, IOPS usage, connection usage, and CPU usage, CloudMonitor helps you to monitor the running status of RDS. CloudMonitor automatically collects data for RDS metrics from the time after you purchase the RDS service.



Note:

- RDS provides monitoring and alarm services only for master and read-only instances.
- After you buy the RDS service, CloudMonitor automatically creates the following four alarm rules for each master instance and read-only instance: CPU usage > 80%, connection usage > 80%, IOPS usage > 80%, and disk usage > 80%. Alarm notifications are sent to alarm contacts through SMS messages and emails when the thresholds of the alarm rules are exceeded.

Monitoring service

- Metrics

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Disk Usage	The percentage of disk space in use in an RDS instance	Instance	%	5 minutes
IOPS Usage	The IOPS usage of an RDS instance	Instance	%	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Connections Usage	The percentage of active connections out of all possible connections that programs establish with an RDS instance.	Instance	%	5 minutes
CPU Usage	The percentage of CPU capacity consumed by an RDS instance (CPU performance is determined by the database memory size.)	Instance	%	5 minutes
Memory Usage	The percentage of the memory in use in an RDS instance. Currently, the memory usage metric is only supported by MySQL databases.	Instance	%	5 minutes
Read-only Instance Delay	MySQL read-only instance latency	Instance	second	5 minutes
Network Inbound Traffic	Inbound traffic to an instance per second	Instance	bit/s	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Network Outbound Traffic	Outbound traffic from an instance per second	Instance	bit/s	5 minutes
RDS Fault	An event-type metric for which alarm rules can be set	N/A	N/A	N/A
RDS Master/Slave Instance Switch	An event-type metric for which alarm rules can be set.	N/A	N/A	N/A

The inbound and outbound traffic metrics only support MySQL and SQLServer databases.

- View monitoring data

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for RDS.
3. Select the region to which the target RDS instance belongs and find the target instance.
4. Click the instance name or click Monitoring Charts from the Actions column to access the Monitoring Charts page.
5. To switch to another chart view, click the chart view button in the upper-left corner of the page.

Alarm service

- Set an alarm rule
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for RDS.
 3. Select the region to which the target RDS instance belongs and find the target instance.
 4. Click the instance name or click Alarm Rules from the Actions column to access the Alarm Rules page.
 5. In the upper-right corner of the displayed page, click Create Alarm Rule.
 6. Set parameters by referring to the following descriptions to create an alarm rule, and then click Confirm.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all RDS instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any RDS instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if

you describe the rule as 5mins Average CPU Usage \geq 90%, the alarm service will check every five minutes whether the average value of CPU usage within five minutes meets or exceeds 90%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage $>$ 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
- **5mins CPU Usage Always $>$ 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
- **5mins CPU Usage Once $>$ 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
- **Total 5mins Internet Outbound Traffic $>$ 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
- **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
- **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
- **Notification Contact:** a group of contacts who receive alarm notifications.
- **Notification Methods:** Emails and DingTalk chatbot.
- **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
- **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
- **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.2 SLB

By monitoring multiple metrics from Server Load Balancer (SLB), such as inbound and outbound traffic and the number of data packets and connections, CloudMonitor helps you to monitor the running status of instances and configure alarm rules accordingly. CloudMonitor automatically collects data from SLB from the time after you create an SLB instance.

Monitoring service

- Metrics
 - Layer-4 metrics

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Port inbound traffic	The traffic consumption for accessing a specified SLB port from the Internet	Port	bit/s	1 minute
Port outbound traffic	The traffic consumption for accessing the Internet from a specified SLB port	Port	bit/s	1 minute
Number of inbound data packets by port	The number of the request packets (per second) that a specified SLB port receives	Port	count/s	1 minute
Number of outbound data packets by port	The number of the request packets (per second) that a specified SLB port sends	Port	count/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Number of new port connections	The average number of times (per second) the first SYN_SENT status occurs in a TCP three-way handshake during a statistical period	Port	count/s	1 minute
Number of active port connections	The number of connections in ESTABLISHED status during a statistical period	Port	count	1 minute
Number of inactive port connections	The number of all TCP connections except the connections in ESTABLISHED status during a statistical period	Port	count	1 minute
Number of concurrent port connections	The total number of connections	Port	count	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Number of backend healthy ECS instances by port	The number of healthy instances reported by a health check	Port	count	1 minute
Number of backend unhealthy ECS instances by port	The number of unhealthy instances reported by a health check	Port	count	1 minute
Number of discarded port connections	The average number of connections discarded per second	Port	count/s	1 minute
Number of discarded inbound data packets by port	The average number of inbound packets discarded per second	Port	count/s	1 minute
Number of discarded outbound data packets by port	The average number of outbound packets discarded per second	Port	count/s	1 minute
Number of discarded inbound bandwidth by port	The average inbound traffic discarded per second	Port	bit/s	1 minute
Number of discarded outbound bandwidth by port	The average outbound traffic discarded per second	Port	bit/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Number of active instance connections	The number of connections in ESTABLISHED status during a statistical period	Instance	count/s	1 minute
Number of inactive instance connections	The number of connections except those in ESTABLISHED status during a statistical period	Instance	count/s	1 minute
Number of discarded instance connections	The number of connections discarded per second	Instance	count/s	1 minute
Number of discarded inbound data packets by instance	The number of inbound packets discarded per second	Instance	count/s	1 minute
Number of discarded outbound data packets by instance	The number of outbound packets discarded per second	Instance	count/s	1 minute
Discarded inbound bandwidth by instance	The amount of inbound traffic discarded per second	Instance	bit/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Discarded outbound bandwidth by instance	The amount of outbound traffic discarded per second	Instance	bit/s	1 minute
Number of concurrent instance connections	The total number of connections of the instance (the sum of active and inactive connections)	Instance	count/s	1 minute
Number of new instance connections	The average number of times (per second) the first SYN_SENT status occurs in a TCP three-way handshake during a statistical period	Instance	count/s	1 minute
Number of inbound data packets by instance	The number of request packets received per second	Instance	count/s	1 minute
Number of outbound data packets by instance	The number of packets sent per second	Instance	count/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Number of inbound bandwidth by instance	The traffic consumption for accessing the SLB instance from the Internet	Instance	bit/s	1 minute
Number of outbound bandwidth by instance	The traffic consumption for accessing the Internet from the SLB instance	Instance	bit/s	1 minute

- Layer-7 metrics

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Port QPS	The QPS of a specified port	Port	count/s	1 minute
Port RT	The average request latency of a specified port	Port	ms	1 minute
Status codes of the format 2xx	The number of status codes of the format 2xx that SLB returns to the client	Port	count/s	1 minute
Status codes of the format 3xx	The number of status codes of the format 3xx that SLB returns to the client	Port	count/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Status codes of the format 4xx	The number of status codes of the format 4xx that SLB returns to the client	Port	count/s	1 minute
Status codes of the format 5xx	The number of status codes of the format 5xx that SLB returns to the client	Port	count/s	1 minute
Other status codes	The number of other status codes that SLB returns to the client	Port	count/s	1 minute
Upstream status codes of the format 4xx	The number of status codes of the format 4xx that RS returns to SLB	Port	count/s	1 minute
Upstream status codes of the format 5xx	The number of status codes of the format 5xx that RS returns to the client	Port	count/s	1 minute
Upstream RT	The average request delay from RS to proxy	Port	ms	1 minute
Instance QPS	The QPS of an instance	Instance	count/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Instance RT	The average request latency of the instance	Instance	count/s	1 minute
Status codes of the format 2xx	The number of status codes of the format 2xx that SLB returns to the client	Instance	count/s	1 minute
Status codes of the format 3xx	The number of status codes of the format 3xx that SLB returns to the client	Instance	count/s	1 minute
Status codes of the format 4xx	The number of status codes of the format 4xx that SLB returns to the client	Instance	count/s	1 minute
Status codes of the format 5xx	The number of status codes of the format 5xx that SLB returns to the client	Instance	count/s	1 minute
Other status codes	The number of status codes of other formats that SLB returns to the client	Instance	count/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring frequency
Upstream status codes of the format 4xx	The number of status codes of the format 4xx that RS returns to SLB	Instance	count/s	1 minute
Upstream status codes of the format 5xx	The number of status codes of the format 5xx that RS returns to SLB	Instance	count/s	1 minute
Upstream RT	The average request delay from RS to proxy	Instance	ms	1 minute

**Note:**

The numbers of new connections, active connections, and inactive connections are all based on TCP connection requests from the client to SLB.

- View monitoring data
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Server Load Balancer to go to the SLB instance list page.
 3. Select the region to which the target instance belongs.
 4. Find the target instance and click the instance name or click Monitoring Charts in the Actions column.
 5. On the Monitoring Charts tab page, you can view the monitoring data.
 6. To switch to another chart view, click the chart view button in the upper-left corner of the page.

Alarm service

- Set an alarm rule
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Server Load Balancer to go to the SLB instance list page.
 3. Select the region to which the target instance belongs.
 4. Find the target instance and click Alarm Rules in the Actions column.
 5. On the Alarm Rules tab page, click Create Alarm Rules in the upper-right corner.
 6. Set the parameters according to the following parameter descriptions and click Confirm.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all SLB instances under your account. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any SLB instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5-minute average CPU usage $\geq 90\%$, the alarm service

will check every five minutes whether the average value of CPU usage within five minutes meets or exceeds 90%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5-minute average CPU usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
- **5-minute CPU usage always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
- **5-minute CPU usage once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
- **Total 5-minute Internet outbound traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Muted For:** the period of time that an alarm has been muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted up to 24 hours (or 1 day).
- **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
- **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
- **Notification Contact:** a group of contacts who receive alarm notifications.
- **Notification Methods:**
 - **Email and DingTalk chatbot**
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.3 OSS

By monitoring the basic service, performance, and metering data of the Object Service Storage (OSS) service, CloudMonitor enables you to gain insights into the overall performance of the OSS service and set alarm rules accordingly. Specifically, this can help you better track requests, analyze usage, collect statistics on business trends, and quickly discover and diagnose system issues.

Monitoring service

- Metrics

The metrics used for monitoring OSS mainly include basic service, performance, and metering indicators. For more information, see [Monitoring indicators reference](#).



Note:

To maintain consistency with the billing policies, the collection and presentation of metering data have the following characteristics:

- Metering data is collected hourly, so that the metering data for your resources is aggregated to a single value each hour. This value represents the overall metering condition of the hour monitored.
- Metering data has an output delay of nearly 30 minutes.
- The metering data time refers to the start time of the relevant statistical period.
- The cutoff time of metering data is the end time of the last statistical period of the current month. If no metering data is produced in the current month, the metering data cutoff time is 00:00 on the first day of the current month.
- For presentation purposes, the maximum quantity of metering data is pushed. For more information about metering data, see [Usage Records](#).

Example

Assuming that you only use PutObject requests to upload data and perform this operation at an average of 10 times per minute. Then, in the hour between 08:00:00 and 09:00:00 on May 10, 2016, the metering result of your PUT requests is 600 times (10 × 60 minutes), the time of metering data is 08:00:00 on May 10, 2016, and the result will be generated at around 09:30:00 on May 10, 2016. If the result is the last data record since 00:00:00 on May 1, 2016, the metering data cutoff time for the current month is 09:00:00 on May 10, 2016. If in May 2016,

you have not produced any metering data, the metering data cutoff time will be 00:00:00 on May 1, 2016.

Alarm service



Note:

The names of OSS buckets are unique. Given this, after you delete a bucket, if you create another one with the same name as the deleted one, the monitoring rules and alarm rules that were previously set for the deleted bucket will also apply to the new bucket.

You can set alarm rules for several metrics in addition to the preceding metering and statistical indicators. You can also add these metrics to your monitoring list. Moreover, multiple alarm rules can be set for a single metric.

Instructions

- For more information about the alarm service, see [Alarm service overview](#).
- For more information about the alarm service for OSS monitoring, see [OSS alarm service user guide](#).

6.4 CDN

By monitoring multiple metrics from Alibaba Cloud Content Delivery Network (CDN), such as QPS, BPS, and byte hit rate, CloudMonitor helps you to gain insights into the usage of domain names. CloudMonitor automatically begins to collect data from CDN domains after you add a CDN domain name. You can view the monitoring details on the Alibaba Cloud CDN page of the CloudMonitor console. You can also configure alarm rules for metrics so that you can be notified of any alarm when data exceptions occur.

Monitoring service

· Metrics

Metric	Description	Dimension	Unit	Minimum monitoring granularity
QPS	The number of visits per second	Domain name	Count/s	1 minute
Peak Bandwidth	The maximum network bandwidth within a certain period of time	Domain name	bit/s	1 minute
Hit Rate	The ratio of bytes served by the cache over the total number of bytes requested by the clients (bytes = number of requests × traffic). The byte hit rate directly reflects the back-to-source traffic.	Domain name	%	1 minute
Public Network Outbound Traffic	The downstream Internet traffic of CDN	Domain name	Byte	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring granularity
4xx Code	The percentage of HTTP 4xx codes out of all returned codes	Domain name	%	1 minute
5xx Code	The percentage of HTTP 5xx codes out of all returned codes	Domain name	%	1 minute

- Viewing monitoring data

- Log on to the [CloudMonitor console](#).
- In the left-side navigation pane, choose Cloud Service Monitoring > Alibaba Cloud CDN.
- Click the Domain Name List tab.
- Find the target domain name and click the domain name, or click Monitoring Charts in the Actions column.
- To switch to a larger view, click the enlargement icon in the upper-right corner of the chart.

Alarm service

- Set an alarm rule

- Log on to the [CloudMonitor console](#).
- In the left-side navigation pane, choose Cloud Service Monitoring > Alibaba Cloud CDN.
- Click the Domain Name List tab.
- Find the target domain name and click Alarm Rules in the Actions column.
- In the upper-right corner of the displayed page, click Create Alarm Rule.
- Configure the alarm rule by referring to the following parameter descriptions. Then, click Confirm.

- Parameters
 - **Products:** ECS, RDS, OSS, among others
 - **Resource Range:** the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Domain Name.
 - **All Resources:** Indicates that the specified alarm rule applies to all CDN domain names under your account. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - **Instances:** Indicates that the specified rule only applies to a specific domain name.
 - **Alarm Rule:** the alarm rule name
 - **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Average QPS \geq 300, the alarm service will check

every minute whether the average value of QPS within five minutes meets or exceeds 90.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm has been muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively. For example, if you set this parameter to 3 and set Rule Describe to 5mins Average CPU Usage > 80%, only when the average CPU usage in five minutes is detected to be greater than 80% for three times in a row, will an alarm be triggered.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Email and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.5 EIP

By monitoring multiple metrics from Elastic IP Address (EIP), such as inbound and outbound traffic and inbound and outbound packet rates, CloudMonitor helps you to monitor the running status of your services and configure alarm rules accordingly. CloudMonitor automatically collects data for the metrics from the time after you purchase an EIP.

Monitoring service

- Metrics

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Inbound Bandwidth	The volume of traffic that flows into an ECS instance through the EIP per second	Instance	bit/s	1 minute
Outbound Bandwidth	The volume of traffic that passes through the EIP from an ECS instance per second	Instance	bit/s	1 minute
Inbound packet rate	The number of packets that flow into an ECS instance through the EIP per second	Instance	Count/s	1 minute
Outbound packet rate	The number of packets that pass through the EIP from an ECS instance per second	Instance	Count/s	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Out Ratelimit Drop Speed	The rate at which packets are dropped due to the occupied bandwidth exceeding the specified peak bandwidth	Instance	Count/s	1 minute

- View monitoring data

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > Elastic IP Address.
3. Select the region to which the target instance belongs.
4. Find the target instance and click the instance name, or click Monitoring Charts in the Actions column.
5. To switch to a larger view, click the enlargement icon in the upper-right corner of the chart.

Alarm service

- Set an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > Elastic IP Address.
3. Select the region to which the target instance belongs.
4. Find the target instance and click Alarm Rules in the Actions column.
5. In the upper-right corner of the displayed page, click Create Alarm Rules.
6. Configure the alarm rule by referring to the following parameter descriptions. Then, click Confirm.

- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all EIPs under your account. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for inbound traffic to 100 MB, an alarm is triggered when the inbound traffic of the specified instance exceeds 100 MB.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Average Inbound Traffic \geq 100 MB, the alarm

service will check every five minutes whether the average value of inbound traffic within five minutes meets or exceeds 100 MB.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm has been muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively. For example, if you set this parameter to 3 and set Rule Describe to 5mins Average CPU Usage > 80%, only when the average CPU usage in five minutes is detected to be greater than 80% for three times in a row, will an alarm be triggered.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Email and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.6 ApsaraDB for Memcache

By monitoring multiple metrics of ApsaraDB for Memcache, such as the cache used and read hit rate, CloudMonitor helps you to monitor the running status of ApsaraDB for Memcache. CloudMonitor automatically collects data for ApsaraDB for Memcache metrics from the time after you purchase the ApsaraDB for Memcache service.

Monitoring service

- Metrics

Monitoring metrics	Meaning	Dimension	Unit	Minimum monitoring granularity
Cache used	The amount of cache used	Instance	Bytes	1 minute
Read hit rate	The probability of reading key-values (KVs) successfully	Instance	Percentage	1 minute
QPS	Total times of reading KVs per second	Instance	Times	1 minute
Number of records	Total number of KVs in the current measurement period	Instance	KVs	1 minute
Cache inbound bandwidth	Traffic generated by accessing the cache	Instance	Bit/s	1 minute
Cache outbound bandwidth	Traffic generated by reading the cache	Instance	Bit/s	1 minute

Monitoring metrics	Meaning	Dimension	Unit	Minimum monitoring granularity
Eviction	Number of KVs evicted per second	Instance	KVs per second	1 minute

**Note:**

- Monitoring data is saved for up to 31 days.
- You can view metric data for up to 14 consecutive days.

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Memcache.
 3. Click an instance name or click Monitoring Charts in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click the Time Range toggle button at the top of the page or use the specific selection function.
 5. Click the Zoom In button in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

CloudMonitor provides alarm services for all Memcache monitoring metrics. After setting an alarm rule for an important monitoring metric, you can receive an alarm notification once the monitoring data exceeds the set threshold value, so that you can handle the problem rapidly to avoid malfunction.

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Memcache.
 3. Click an instance name or click Monitoring Charts in the Actions column to access the instance monitoring details page.
 4. Click the bell icon in the upper-right corner of the monitoring chart to set an alarm rule for corresponding monitoring metrics of this instance.
- Set batch alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Memcache.
 3. Select the appropriate instance on the instance list page. Click Set Alarm Rules at the bottom of the page to add alarm rules in batches.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all ApsaraDB for Memcache instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any ApsaraDB for Memcache instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm

threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Average CPU Usage \geq 90%, the alarm service

will check every five minutes whether the average value of CPU usage within five minutes meets or exceeds 90%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP Callback:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.7 ApsaraDB for Redis

By monitoring multiple metrics from ApsaraDB for Redis, such as CPU usage, memory usage, and connection usage, CloudMonitor helps you to monitor the running status and usage of Redis. CloudMonitor automatically collects data for the metrics from the time after you create a Redis instance. You can view the Redis monitoring data on the ApsaraDB for Redis page of the CloudMonitor console. You can also configure alarm rules for metrics so that you can be notified of any alarm when data exceptions occur.

Monitoring service

- Metrics

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Capacity in use	The Redis capacity that is occupied	Instance	Byte	1 minute
Number of connections used	Total number of client connections	Instance	Count	1 minute
Write speed	The network traffic written per second	Instance	bit/s	1 minute
Read Speed	The network traffic read per second	Instance	bit/s	1 minute
Failed operations	The number of failed operations on KVSTORE	Instance	Count	1 minute
Capacity usage	The percentage of currently occupied capacity to total capacity	Instance	%	1 minute

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Connections Usage	The percentage of established connections out of all possible connections	Instance	%	1 minute
Inbound Bandwidth Usage	The inbound bandwidth as a percentage of total bandwidth	Instance	%	1 minute
Outbound Bandwidth Usage	The outbound bandwidth as a percentage of total bandwidth	Instance	%	1 minute
Redis fault	Event-type metric, for which alarm rules can be set	N/A	N/A	N/A
Redis Master/ Slave Instance Switch	Event-type metric, for which alarm rules can be set.	N/A	N/A	N/A

- View monitoring data
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Redis.
 3. Select the region to which the target instance belongs and select the instance type.
 4. Find the target instance and click the instance name or click Monitoring Charts from the Actions column.
 5. To switch to a larger view, click the enlargement icon in the upper-right corner of the chart.

Alarm service

- Set an alarm rule
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Redis.
 3. Select the region to which the target instance belongs and select the instance type.
 4. Find the target instance and click Alarm Rules in the Actions column.
 5. In the upper-right corner of the displayed page, click Create Alarm Rules.
 6. Configure the alarm rule by referring to the following parameter descriptions. Then, click Confirm.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all Redis instances under your account. For example, if you set the resource range to All Resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Redis instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in

your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for inbound traffic to 100 MB, an alarm is triggered when the inbound traffic of the specified instance exceeds 100 MB.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Average Inbound Traffic \geq 100 MB, the alarm service will check every five minutes whether the average value of inbound traffic within five minutes meets or exceeds 100 MB.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

■ **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.

■ **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.

■ **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.

■ **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.

- **Mute For:** the period of time that an alarm has been muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted up to 24 hours (or 1 day).
- **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively. For example, if you set this parameter to 3 and set Rule Describe to 5min Average CPU Usage > 80%, only when the average CPU usage in five

minutes is detected to be greater than 80% for three times in a row, will an alarm be triggered.

- **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
- **Notification Contact:** a group of contacts who receive alarm notifications.
- **Notification Methods:** Email and DingTalk chatbot.
- **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
- **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email.
- **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.8 ApsaraDB for MongoDB

By monitoring multiple metrics from ApsaraDB for MongoDB, such as CPU usage and memory usage, CloudMonitor helps you to monitor the running status of instances. CloudMonitor automatically begins to collect data for the metrics after you purchase the MongoDB service.

Monitoring service

- **Metrics**

Metric	Description	Dimension	Unit	Minimum monitoring granularity
CPU Usage	The percentage of CPU in use of the instance	User, instance, and master/salve	%	5 minutes
Memory Usage	The memory in use of the instance	User, instance, and master/slave	%	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Disk Usage	The disk space in use of the instance	User, instance , and master/ slave	%	5 minutes
IOPS Usage	The percentage of the actual IOPS to the maximum IOPS of the instance	User, instance , and master/ slave	%	5 minutes
Connection Usage	The number of connections is the number of instances that an application can connect to a MongoDB instance. Connection usage is the percentage of the connections currently in use.	User, instance , and master/ slave	%	5 minutes
Number of Query Operations	The number of SQL queries per second of the instance	User, instance , and master/ slave	Count	5 minutes
Connection Usage	The number of connections in use	User, instance , and master/ slave	Count	5 minutes
Disk Size Occupied by Instances	Total amount of disk space in use of the instance	User, instance , and master/ slave	Byte	5 minutes
Disk Size Occupied by data	The amount of disk space used by data	User, instance , and master/ slave	Byte	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Disk Size Occupied by logs	The amount of disk space used by logs	User, instance , and master/ slave	Byte	5 minutes
Intranet Inbound Traffic	The network inbound traffic for an instance	User, instance , and master/ slave	Byte	5 minutes
Intranet Outbound Traffic	The network outbound traffic for an instance	User, instance , and master/ slave	Byte	5 minutes
Number of Requests	Total number of requests sent to the server	User, instance , and master/ slave	Count	5 minutes
Number of Insert Operations	The number of Insert commands received during the time from the last start of the instance to the current time	User, instance , and master/ slave	Count	5 minutes
Number of Query Operations	The number of Query commands received during the time from the last start of the instance to the current time	User, instance , and master/ slave	Byte	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Number of Update Operations	The number of Update commands received during the time from the last start of the instance to the current time	User, instance , and master/ slave	Count	5 minutes
Number of Delete Operations	The number of Delete commands executed during the time from the last start of the instance to the current time	User, instance , and master/ slave	Count	5 minutes
Number of Getmore Operations	The number of Getmore commands executed during the time from the last start of the instance to the current time	User, instance , master/slave	Count	5 minutes
Number of Command Operations	The total number of commands sent to the database during the time from the last start of the instance to the current time	User, instance , and master/ slave	Count	5 minutes

Metric	Description	Dimension	Unit	Minimum monitoring granularity
Instance Failure	Event-type metric, for which alarm rules can be set	N/A	N/A	N/A

**Note:**

- Monitoring data is saved for up to 31 days.
- You can view metric data for up to 14 consecutive days at a time.

- View monitoring data

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for MongoDB.
3. Select the region to which the target instance belongs and select the instance type.
4. Find the target instance and click the instance name or click Monitoring Charts from the Actions column.
5. Select the Time Range by clicking the time duration button or choose a custom time range. You can view the monitoring data in up to 14 days at a time.
6. To switch to a larger view, click the enlargement icon in the upper-right corner of the chart.

Alarm service

- Set a single alarm rule
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for MongoDB.
 3. Select the region to which the target instance belongs and select the instance type.
 4. Find the target instance and click the instance name or click Monitoring Charts from the Actions column.
 5. On the Monitoring Charts tab page, click the bell icon in the upper-right corner of a chart to set an alarm rule.
- Set alarm rules in batches
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for MongoDB.
 3. Select the target instances on the instance list page. Then, click Set Alarm Rules below the list to add alarm rules in batches.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all MongoDB instances under your account. For example, if you set the resource range to All Resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any MongoDB instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm

threshold for inbound traffic to 100 MB, an alarm is triggered when the inbound traffic of the specified instance exceeds 100 MB.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Average Inbound Traffic \geq 100 MB, the alarm service will check every five minutes whether the average value of inbound traffic within five minutes meets or exceeds 100 MB.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
- **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
- **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
- **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm has been muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted up to 24 hours (or 1 day).
- **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively. For example, if you set this parameter to 3 and set Rule Describe to 5mins Average CPU Usage > 80%, only when the average CPU usage in five

minutes is detected to be greater than 80% for three times in a row, will an alarm be triggered.

- **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
- **Notification Contact:** a group of contacts who receive alarm notifications.
- **Notification Methods:** Email and DingTalk chatbot.
- **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
- **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email.
- **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.9 Message Service monitoring

Cloud monitoring through the monitoring of Message Service delay messages, invalid messages, active messages, three monitoring items, help Users get message service Queue usage.

When users create message queues for message service, cloud monitoring automatically begins to monitor them, you are logged in to the cloud monitoring Message Service You can view monitoring details on the page. You can also set alarm rules on monitoring items so that you receive alarm information when the data is abnormal.

Monitoring Services

- Monitoring item description

Monitoring items	Meaning	Dimensions	Unit	Minimum monitor Granularity
Activemessages	The total number of messages that are active in this queue	Userid, region, bid, queue	Number	5 minutes

Monitoring items	Meaning	Dimensions	Unit	Minimum monitor Granularity
Inactivemessages	The total number of messages that are inactive in this queue	Userid, region, bid, queue	Number	5 minutes
Delaymessage	The total number of messages in the delayed state in this queue	Userid, region, bid, queue	Number	5 minutes
Sendmessagecount	Send message requests	Userid, region, queue	Items	3600
Batchsendmessagecount	Number of bulk send message requests	Userid, region, queue	Items	3600
Receivemessagecount	Number of receive message requests	Userid, region, queue	Items	3600
Batchreceivecount	Number of bulk receive message requests	Userid, region, queue	Items	3600
Batchdeletecount	Bulk Delete message request quantity	Userid, region, queue	Items	3600
Changemessagevisibilitycount	Change message visibility count	Userid, region, queue	Items	3600

- Viewing Monitoring Data

1. Log in to the cloud monitoring console.
2. Enter the list of message service instances that the cloud service monitors.
3. Click the Instance name or the monitor chart in the operation to go to the monitor details page.
4. Click the size chart toggle button to toggle the larger image display (optional).

Alarm service

- Parameter descriptions

- **Monitor:** the monitoring metrics provided by the message service.
- **Statistical Cycle:** the alarm system checks if your monitoring data exceeds the alarm threshold for this cycle. For example, to set a memory usage alarm rule, the statistics cycle is 1 minute, an interval of 1 minute checks if memory usage exceeds the threshold.
- **Statistical Methods:** Statistical Methods refer to settings that exceed the threshold range. You can set the average, maximum, minimum, count, and value in a statistical method.
 - **Average:** the average of the monitored data during the statistical cycle. The result is an average of all monitoring data collected within 15 minutes, when this average is greater than 80. Only when the threshold is exceeded.
 - **Maximum:** the maximum value of the monitor data during the statistics cycle. The maximum value of monitoring data collected during the statistical cycle exceeds the threshold of 80.
 - **Minimum:** the minimum value of the monitored data during the statistics cycle. The minimum value for monitoring data collected during the statistical cycle exceeds the threshold of 80.
 - **Value:** Sum of monitoring data during the statistics cycle. Sum up the monitoring data collected during the statistical period with more than 80

% results after the sum That is, the threshold is exceeded. Such statistical methods are required for traffic-class metrics.

- **Alarm after several consecutive exceeds threshold:** refers to a number of consecutive statistical cycle monitoring items whose values continue to exceed the threshold to trigger an alarm.

Example: Set CPU usage to more than 80% alarm with a statistical cycle of 5 In minutes, alarm after 3 consecutive times exceeds the threshold, when the first time the CPU usage is detected exceeds 80, no alarm notification will be issued. The second time in 5 minutes to probe the CPU Usage is more than 80%, and no alarm will be issued. The third time the probe is still over 80%, alarm notification will be issued. That is, from the first time the actual data exceeds the threshold to the final alarm rule, the minimum time required is the statistical cycle (number of consecutive probes-1) = 5 (3-1) = 10 minutes.

- Set alarm rules
 1. Log in to the cloud monitoring console.
 2. Enter the list of message service instances that the cloud service monitors.
 3. Click the alarm rule in the instance List action to enter the alert Rule Page for the instance.
 4. Click new alarm rule in the upper-right corner of the alarm Rule Page to create an alarm rule based on the parameter.

6.10 AnalyticDB monitoring

By monitoring multiple metrics of AnalyticDB, such as rated disk space, disk space in use, and disk usage, CloudMonitor helps you to monitor the usage of the AnalyticDB service. CloudMonitor automatically collects data for AnalyticDB from the time after you purchase the AnalyticDB service. You can view monitoring details on the AnalyticDB monitoring page in the CloudMonitor console. You can also configure alarm rules for metrics so that an alarm is generated when any data exception occurs.

Monitoring service

- Metrics

Metric	Description	Dimension	Unit	Minimum monitoring frequency
diskSize	The rated disk capacity	instanceId, tableSchema, workerId	MB	1 minute
diskUsed	The disk capacity in use	instanceId, tableSchema, workerId	MB	1 minute
diskUsedPercent	The disk usage	Instanceid, tableschema, workerid	%	1 minute

- View monitoring data

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > Analytic DB.
3. Find the target instance and click the instance name or click Monitoring Charts from the Actions column.
4. To switch to another chart view, click the chart view button in the upper-left corner of the page.

Alarm service

- Set an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > Analytic DB.
3. Find the target instance and click Alarm Rules in the Actions column.
4. In the upper-right corner of the displayed page, click Create Alarm Rule.
5. Set parameters by referring to the following descriptions to create an alarm rule, and then click Confirm.

- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all AnalyticDB instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Average CPU Usage \geq 90%, the alarm service

will check every five minutes whether the average value of CPU usage within five minutes meets or exceeds 90%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.11 Log Service

By monitoring multiple metrics of Log Service, such as the outbound traffic, inbound traffic, overall QPS, and log statistic method, CloudMonitor helps you to monitor the running status of Log Service. CloudMonitor automatically collects data for Log Service metrics from the time after you purchase the Log Service service. CloudMonitor also allows you to set alarm rules for these metrics so that you can receive alarm notifications once data exceptions occur.

Monitoring service

- Metrics

Metric	Description	Dimension	Unit	Minimum monitor granularity
Inflow	Logstore inbound and outbound traffic per minute	userId, Project, and Logstore	Bytes	1 minute
Outflow	Logstore traffic per minute	userId, Project, and Logstore	Bytes	1 minute
SumQPS	Total number of writes per minute in the Logstore	userId, Project, and Logstore	Count	1 minute
LogMethodQPS	Total number of writes per minute to the Logstore	userId, Project, and Logstore	Count	1 minute
LogCodeQPS	Number of writes per minute mapped to a specific status code in the Logstore	userId, Project, and Logstore	Count	1 minute

Metric	Description	Dimension	Unit	Minimum monitor granularity
SuccessdByte	Number of successfully resolved bytes in the Logstore	userId, Project , and Logstore	Bytes	10 minutes.
SuccessdLines	Number of lines in resolved logs in the Logstore	userId, Project , and Logstore	Count	10 minutes
Failedlines	Number of lines in logs failed to be resolved in the Logstore	userId, Project , and Logstore	Count	10 minutes
AlarmPV	Total number of ECS configuration errors in the LogStore	userId, Project , and Logstore	Count	5 minutes
AlarmUv	Total number of ECS instances with incorrect configurations in the Logstore	userId, Project , and Logstore	Count	5 minutes
AlarmIPCount	Number of errors incurred by a specific IP address in the Logstore	userId, Project , Logstore, alarm_type, and source_ip	Count	5 minutes

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Log Service.
 3. Click an instance name from the product instance list or click Monitoring Charts in the Actions column to access the instance monitoring details page.
 4. (Optional) Click the Chart Size button to switch to large chart display.

Alarm service

- Set an alarm rule
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Log Service.
 3. Click Alarm Rules in the Actions column to access the instance's Alarm Rules page.
 4. Enter all the required information and click Confirm.
- Parameters



Note:

- When setting alarm rules, you can specify the status for a specific metric. Status codes include 200, 400, 401, 403, 405, 500, and 502.
- You can specify the method for a specific metric. Valid values of the method fields are PostLogStoreLogs, GetLogtailConfig, PutData, GetCursorOrData, GetData, GetLogStoreHistogram, GetLogStoreLogs, ListLogStores, and ListLogStoreTopics.
- Products: ECS, RDS, OSS, among others
- Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all Log Service instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Log Service instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you

set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as `1mins Overall QPS >= 70`, the alarm service will check

every one minute whether the overall PQS within one minute meets or exceeds 70.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.12 Container Service monitoring

Cloud monitoring through monitoring CPU usage, memory usage, and so on for container services 7 Monitoring items to help users get use of Container Services.

After the user creates the container service, cloud monitoring automatically begins to monitor the container service, you can log in to the Container Services page for cloud monitoring to view monitoring details. You can also set alarm rules on the monitor so that you receive an alarm notification when the data is abnormal.

Monitoring Services

- Monitoring item description

Monitoring items	Meaning	Dimensions	Unit	Minimum monitor Granularity
containerCpuUtilization	Container CPU usage	User dimension , container dimension	Percentage	30 seconds
Containermemoryutilization	Container memory usage	User dimension , container dimension	Percentage	30 seconds
Containermemoryamount	Container memory usage	User dimension , container dimension	Bytes	30 seconds
Containerinternetin	Container traffic into the network	User dimension , container dimension	Bytes	30 seconds
containerinternetOut	Container traffic out of the Network	User dimension , container dimension	Bytes	30 seconds
containerIOWRead	Container Io read	User dimension , container dimension	Bytes	30 seconds

Monitoring items	Meaning	Dimensions	Unit	Minimum monitor Granularity
containerIOWrite	Container I/O write	User dimension, container dimension	Bytes	30 seconds



Note:

- Monitor Data is saved for up to 31 days.
- Users can view monitoring data for up to 14 days in a row.

· Viewing Monitoring Data

1. Log in to the cloud monitoring console.
2. Enter the list of Container service instances that the cloud service monitors.
3. Click the Instance name or the monitor chart in the operation to go To the instance monitoring details page to view the metrics.
4. Click the time range quick select button or the exact select function at the top of the page, maximum monitoring data support view continuous 14 Monitoring data for days.
5. Click the zoom in button in the upper-right corner of the monitor MAP to view the monitor larger image.

Alarm service

- Set single alarm rule: Click the bell button in the upper right corner of the monitor chart, alarm rules can be set for monitoring items corresponding to this instance.
- Sets the bulk alarm rule: the list of instances page selects the desired instance, you can add alarm rules in bulk by clicking set alarm rules below the page.

6.13 Shared Bandwidth monitoring

By monitoring multiple metrics from Shared Bandwidth, such as inbound and outbound bandwidth, CloudMonitor helps you to monitor the usage of the bandwidth . CloudMonitor automatically begins to collect data for the metrics after you create a cluster. You can also set alarm rules for clusters through CloudMonitor so that you can be notified in the case of any exceptions.

Monitoring Services

- Monitoring items

Monitoring items	Dimensions	Unit	Minimum monitor Granularity
Bandwidth packet network inflows bandwidth	User dimension , instance dimension	Bits/s	1 minute
Bandwidth packet network outgoing bandwidth	User dimension , instance dimension	Bits/s	1 minute
Bandwidth packets network inflows packets	User dimension , instance dimension	packages/s	1 minute
Bandwidth packet network flow Packet	User dimension , instance dimension	packages/s	1 minute
Bandwidth packet network outgoing bandwidth usage	User dimension , instance dimension	%	1 minute



Note:

- Monitor Data is saved for up to 31 days.
- Users can view monitoring data for up to 7 days in a row.

- Viewing Monitoring Data

1. Log on to the [CloudMonitor console](#).
2. Enter the list of instances of shared bandwidth that the cloud service monitors.
3. Click the Instance name or the monitor chart in the operation to go To the instance monitoring details page to view the metrics.
4. Click the time range on the top of the page to quickly select a button or select an exact function, monitoring data supports viewing monitoring data for seven consecutive days.
5. Click the zoom button in the upper-right corner of the monitor MAP to view the monitor larger image.

Alarm service

- Parameter descriptions
 - **Monitor:** that is, the monitoring metrics provided by services that share bandwidth.
 - **Statistical Cycle:** the alarm system checks if your monitoring data exceeds the alarm threshold for this cycle. For example, to set a memory usage alarm rule, the statistics cycle is 1 minute, an interval of 1 minute checks if memory usage exceeds the threshold.
 - **Continuous number of times:** refers to the continuous number of statistical cycle monitoring items that continue to exceed the threshold value to trigger the alarm.
- Set single alarm rule
 1. Log on to the [CloudMonitor console](#).
 2. Enter the list of instances of shared bandwidth that the cloud service monitors.
 3. Click the Instance name or the monitor chart in the operation to go To the instance monitoring details page.
 4. Click the bell button in the upper right corner of the monitor chart or the new alarm rule in the upper right corner of the page, alarm rules can be set for monitoring items corresponding to this instance.
- Set up bulk alarm rules
 1. Log on to the [CloudMonitor console](#).
 2. Enter the list of shared bandwidth instances that the cloud service monitors.
 3. When the instance list page selects the desired instance, click set alarm rule below the page, you can add alarm rules in bulk.

6.14 Global Acceleration

By monitoring multiple metrics of Global Acceleration, such as the such as inbound and outbound network bandwidth, CloudMonitor helps you to monitor the running status of Global Acceleration. CloudMonitor automatically collects data for Global Acceleration metrics from the time after you purchase the Global Acceleration service

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitor granularity
Inbound bandwidth	User and instance	Bits/s	1 minute
Outbound bandwidth	User and instance	Bits/s	1 minute
Inbound package	User and instance	PPS	1 minute
Outbound package	User and instance	PPS	1 minute



Note:

- Monitoring data is saved for up to 31 days.
- You can view the monitoring data for up to 7 consecutive days.

- View monitoring data.

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > Global Acceleration.
3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
4. Click the Time Range quick selection button from the upper menu of the page or use the specific selection function. You can view the monitoring data for up to 7 consecutive days.
5. Click Zoom In in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Global Acceleration.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page.
 4. Click bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Global Acceleration.
 3. Select the appropriate instances on the instance list page. Click Set Alarm Rules to add multiple alarm rules.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all Global Acceleration instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Global Acceleration instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm

threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 1mins Inbound Bandwidth Value \geq 10 Mbit/s, the

alarm service will check every one minutes whether the inbound bandwidth within one minute meets or exceeds 10 Mbit/s.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.15 HiTSDB

By monitoring multiple metrics of HiTSDB, such as the disk usage, number of timelines, and incremental number of time points, CloudMonitor helps you to monitor the running status of HiTSDB. CloudMonitor automatically collects data for HiTSDB metrics from the time after you purchase the HiTSDB service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitor granularity
Disk usage	User and instance	%	20 seconds
Number of timelines	User and instance	Count	20 seconds
Incremental number of time points	User and instance	Count/second	20 seconds



Note:

- Monitoring data is saved for up to 31 days.
- You can view the monitoring data for up to 14 consecutive days.

- View monitoring data.

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > HiTSDB.
3. Click an instance name or click Monitoring Chart in the Action column to access the instance monitoring details page and view various metrics.
4. Click a Time Range shortcut on the top of the page or use the specific selection function. Up to 14 consecutive days of metric data can be viewed.
5. Click the Zoom In button in the upper-right corner of the monitor map to view the monitor larger image.

Alarm service

- Description
 - **Monitor:** the monitoring indicator provided by hitsdb's service.
 - **Statistical Cycle:** the alarm system checks if your monitoring data exceeds the alarm threshold for this cycle. For example, if the statistical cycle of the alarm rule for memory usage is set to one minute, the system checks whether the memory usage has exceeded the threshold value every other minute.
 - **Consecutive times:** an alarm is triggered when the value of the monitoring metrics continuously exceeds the threshold value for the set consecutive cycles.
- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > HiTSDB.
 3. Click the instance name or the monitor chart in the operation to go to the instance monitoring details page.
 4. Click the Bell button in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > HiTSDB.
 3. When the instance list page selects the desired instance, click set alarm rule below the page, you can add alarm rules in bulk.
- Parameters
 - **Products:** ECS, RDS, OSS, among others
 - **Resource Range:** the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all HiTSDB instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any HiTSDB instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported

for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if

you describe the rule as 1mins Disk Usage \geq 30%, the alarm service will check every minute whether the disk usage within one minute meets or exceeds 30%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.16 VPN

By monitoring multiple metrics of VPN, such as the inbound bandwidth and outbound bandwidth, CloudMonitor helps you to monitor the running status of VPN. CloudMonitor automatically collects data for VPN metrics from the time after you purchase the VPN service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitoring granularity
Inbound network bandwidth of a bandwidth package	User and instance	Bit/s	1 minute
Outbound network bandwidth of a bandwidth package	User and instance	Bit/s	1 minute
Incoming packet of a bandwidth package	User and instance	PPS	1 minute
Outgoing packet of a bandwidth package	User and instance	PPS	1 minute



Note:

- Monitoring data is saved for up to 31 days.
- You can view the monitoring data for up to 7 consecutive days.

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > VPN.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click the Time Range quick selection button at the top of the page or use the specific selection function. You can view the monitoring data for up to 7 consecutive days.
 5. Click the Zoom In button in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > VPN.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page.
 4. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > VPN.
 3. Select the appropriate instance on the instance list page. Click Set Alarm Rules at the bottom of the page to add alarm rules in batches.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all VPN instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any VPN instance exceeds 80%. When

you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 1mins Inbound Bandwidth \geq 1 Mbit/s, the alarm

service will check every minute whether the inbound bandwidth within one minute meets or exceeds 1 Mbit/s.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.17 API Gateway

By monitoring multiple metrics of API Gateway, such as the cache used and read hit rate, CloudMonitor helps you to monitor the running status of API Gateway. CloudMonitor automatically collects data for API Gateway metrics from the time after you purchase the API Gateway service. CloudMonitor also allows you to set alarm rules for these metrics so that you can receive alarm notifications once data exceptions occur.

Monitoring service

- Metrics

Metric	Description	Dimension	Unit	Minimum monitor granularity
Error Distribution	Number of times of 2XX, 4XX, and 5XX status codes returned for an API in one monitoring period.	User and API	Count	1 minute
Inbound traffic	The sum of traffic of requests from an API in one monitoring period	User and API	Bytes	1 minute
Outbound traffic	The sum of traffic of requests from an API in one monitoring period.	User and API	Bytes	1 minute

Metric	Description	Dimension	Unit	Minimum monitor granularity
Response time	The difference between the time when the gateway calls a backend service through an API and the time when the backend service receives the return result in a monitoring period.	User and API	Second	1 minute
The sum of requests	Total number of requests received by an API in a monitoring period	User and API	Count	1 minute

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > API Gateway.
 3. Click an instance name in the product instance list or click Metric Chart in the Actions column to access the instance monitoring details page.
 4. (Optional) Click the Chart Size button to switch to large chart display.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > API Gateway.
 3. Click an instance name in the product instance list or click Alarm Rules in the Actions column to access the instance monitoring details page.
 4. Click Create Alarm Rules at the upper right of the alert policies page to create an alert policy based on the entered parameters.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all API Gateway instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any API Gateway instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 1mins Inbound Bnadwidth >= 1024 KByte/s, the alarm

service will check every minute whether the inbound bandwidth within one minute meets or exceeds 1024 KByte/s.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.18 DirectMail

By monitoring multiple metrics of DirectMail, such as the WEB/API messaging, SMTP messaging, and metrics related to account exceptions, CloudMonitor helps you to monitor the running status of DirectMail. CloudMonitor automatically collects data for DirectMail metrics from the time after you purchase the DirectMail service.

Monitoring service

- Metrics

Metric	Unit	Minimum monitor granularity
Web/API error-QPS delayed	Count/Min	1 minute
Web/API error-over-quota QPS	Count/Min	1 minute
Web/API error-spam QPS	Count/Min	1 minute
Web/API message success QPS	Count/Min	1 minute
SMTP authentication failed QPS	Count/Min	1 minute
SMTP authentication is successful QPS	Count/Min	1 minute
SMTP error-length exceeded QPS	Count/Min	1 minute
SMTP error-over-quota QPS	Count/Min	1 minute
SMTP error-spam QPS	Count/Min	1 minute



Note:

- Monitor Data is saved for up to 31 days.
- Users can view monitoring data for up to 14 days in a row.

- Viewing monitoring data.

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, choose Cloud Service Monitoring > DirectMail.

Alarm service

CloudMonitor provides alarm functions for DirectMail metrics, so that you are notified immediately in case of any metric exceptions.

- Set alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > DirectMail.
 3. Click Alarm Rules to go to the Alarm Rules list page. Click Create Alarm Rules in the upper-right corner to create alarm rules.

Or click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.

- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all ApsaraDB for Memcache instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any ApsaraDB for Memcache instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 1mins Average CPU Usage >= 80%, the alarm service will

check every minute whether the average value of CPU usage within one minute meets or exceeds 80%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.19 Elasticsearch

By monitoring multiple metrics of Elasticsearch, such as the cluster status, cluster QPS, and cluster write QPS, CloudMonitor helps you to monitor the running status of Elasticsearch. CloudMonitor automatically collects data for Elasticsearch metrics from the time after you purchase the Elasticsearch service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitoring granularity
Cluster status	Cluster		1 minute
Cluster query QPS	Cluster	Count/Second	1 minute
Cluster writing QPS	Cluster	Count/Second	1 minute
Node CPU usage	Node	%	1 minute
Node disk usage	Node	%	1 minute
Node heapmemory usage	Node	%	1 minute
Node: load_1m	Node		1 minute
Node FullGc times	Node	Count	1 minute
Node Exception times	Node	Count	1 minute
Cluster snapshot status	Cluster	The value -1 indicates that there is no snapshot; 0 indicates success ; 1 indicates in progress; and 2 indicates failure.	1 minute



Note:

- Monitoring data is saved for up to 31 days.
- You can view the monitoring data for up to 14 consecutive days.

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Elasticsearch.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click a Time Range shortcut on the top of the page or use the specific selection function. Up to 14 consecutive days of metric data can be viewed.
 5. Click Zoom In in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Elasticsearch.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page.
 4. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - **All Resources:** Indicates that the specified alarm rule applies to all Elasticsearch instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Elasticsearch instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that

you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as `1mins Average CPU Usage >= 50%`, the alarm service will

check every minute whether the average value of CPU usage within one minute meets or exceeds 50%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.20 Auto Scaling

By monitoring multiple metrics of Auto Scaling, such as the maximum and minimum numbers of instances, CloudMonitor helps you to monitor the running status of Auto Scaling. CloudMonitor automatically collects data for Auto Scaling metrics from the time after you purchase the Auto Scaling service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitor granularity
Minimum number of instances	User dimension , elastic scaling group	Items	5 minutes
Maximum number of instances	User dimension , elastic scaling group	Items	5 minutes
Total number of instances	User dimension , elastic scaling group	Items	5 minutes
Number of running instances	User dimension , elastic scaling group	Items	5 minutes
Joining instance number	User dimension , elastic scaling group	Items	5 minutes
Removing number of instances	User dimension , elastic scaling group	Items	5 minutes



Note:

- Monitor Data is saved for up to 31 days.
- Users can view monitoring data for up to 14 days in a row.

- Viewing monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Auto Scaling.
 3. Click the instance name or the monitor chart in the operation to go to the instance monitoring details page to view the metrics.
 4. Click the Time Range toggle button or the exact select function at the top of the page, the maximum monitoring data allows you to view the monitored data for 14 consecutive days.
 5. Click Zoom In in the upper-right corner of the monitor map to view the monitor larger image.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Auto Scaling.
 3. Click the instance name or the monitor chart in the operation to enter the instance monitoring details page.
 4. Click the bell button in the upper right corner of the monitor chart or the new alarm rule in the upper right corner of the page, alarm rules can be set for monitoring items corresponding to this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Auto Scaling.
 3. Enter the list of elastic scale monitoring instances monitored by the cloud service.
 4. Select the appropriate instance on the instance list page. Then, click Set Alert Policies at the bottom of the page to add multiple alert policies.

- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all Auto Scaling instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Auto Scaling instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Maximum Number of Instance \geq 10, the

alarm service will check every five minutes whether the maximum number of instances within five minutes meets or exceeds 10.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.21 E-MapReduce

By monitoring multiple metrics of E-MapReduce, such as the CPU idle rate, memory capacity, and disk capacity, CloudMonitor helps you to monitor the running status of E-MapReduce. CloudMonitor automatically collects data for E-MapReduce metrics from the time after you purchase the E-MapReduce service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitoring granularity
Inbound traffic rate	User, cluster, and role	bits/s	30s
Outbound network rate Network drain Rate	User, cluster, and role	bits/s	30s
CPU idleness	User, cluster, and role	%	1 minute
User-mode CPU usage	User, cluster, and role	%	30s
System-mode CPU usage	User, cluster, and role	%	30s
Idle disk capacity	User, cluster, and role	Bytes	30s
Total disk capacity	User, cluster, and role	Bytes	30s
Average load within 15 minutes	User, cluster, and role	-	30s
Average load within 5 minutes	User, cluster, and role	-	30s
Average load within 1 minutes	User, cluster, and role	-	30s
Idle memory capacity	User, cluster, and role	Bytes	30s

Metric	Dimension	Unit	Minimum monitoring granularity
Total memory capacity	User, cluster, and role	Bytes	30s
Inbound data packet rate	User, cluster, and role	Packets/s	30s
Outbound data packet rate	User, cluster, and role	Packets/s	30s
Number of running processes	User, cluster, and role	Processes	30s
Total number of processes	User, cluster, and role	Processes	30s
Number of blocked processes	User, cluster, and role	Processes	30s
Number of created processes/threads	User, cluster, and role	Processes/threads	30s
MemNonHeapUsedM	User, cluster, and role	Bytes	30s
MemNonHeapCommittedM	User, cluster, and role	Bytes	30s
Memnonheapmaxm	User, cluster, and role	Bytes	30s
MemHeapUsedM	User, cluster, and role	Bytes	30s
MemHeapCommittedM	User, cluster, and role	Bytes	30s
MemHeapMaxM	User, cluster, and role	Bytes	30s
MemMaxM	User, cluster, and role	Bytes	30s
Threadsnew	User, cluster, and role	-	30s
ThreadsRunnable	User, cluster, and role	-	30s

Metric	Dimension	Unit	Minimum monitoring granularity
ThreadsBlocked	User, cluster, and role	-	30s
ThreadsWaiting	User, cluster, and role	-	30s
ThreadsTimedWaiting	User, cluster, and role	-	30s
ThreadsTerminated	User, cluster, and role	-	30s
GcCount	User, cluster, and role	-	30s
GcTimeMillis	User, cluster, and role	-	30s
CallQueueLength	User, cluster, and role	-	30s
NumOpenConnections	User, cluster, and role	-	30s
ReceivedBytes	User, cluster, and role	-	30s
SentBytes	User, cluster, and role	-	30s
BlockCapacity	User, cluster, and role	-	30s
BlocksTotal	User, cluster, and role	-	30s
CapacityRemaining	User, cluster, and role	-	30s
CapacityTotal	User, cluster, and role	-	30s
CapacityUsed	User, cluster, and role	-	30s
CapacityUsedNonDFS	User, cluster, and role	-	30s

Metric	Dimension	Unit	Minimum monitoring granularity
CorruptBlocks	User, cluster, and role	-	30s
ExcessBlocks	User, cluster, and role	-	30s
ExpiredHeartbeats	User, cluster, and role	-	30s
MissingBlocks	User, cluster, and role	-	30s
PendingDataNodeMessageCount	User, cluster, and role	-	30s
PendingDeletionBlocks	User, cluster, and role	-	30s
PendingReplicationBlocks	User, cluster, and role	-	30s
PostponedMisreplicatedBlocks	User, cluster, and role	-	30s
ScheduledReplicationBlocks	User, cluster, and role	-	30s
TotalFiles	User, cluster, and role	-	30s
TotalLoad	User, cluster, and role	-	30s
UnderReplicatedBlocks	User, cluster, and role	-	30s
BlocksRead	User, cluster, and role	-	30s
BlocksRemoved	User, cluster, and role	-	30s
BlocksReplicated	User, cluster, and role	-	30s
BlocksUncached	User, cluster, and role	-	30s

Metric	Dimension	Unit	Minimum monitoring granularity
BlocksVerified	User, cluster, and role	-	30s
BlockVerificationFailures	User, cluster, and role	-	30s
BlocksWritten	User, cluster, and role	-	30s
BytesRead	User, cluster, and role	-	30s
BytesWritten	User, cluster, and role	-	30s
FlushNanosAvgTime	User, cluster, and role	-	30s
FlushNanosNumOps	User, cluster, and role	-	30s
FsyncCount	User, cluster, and role	-	30s
VolumeFailures	User, cluster, and role	-	30s
ReadBlockOpNumOps	User, cluster, and role	-	30s
ReadBlockOpAvgTime	User, cluster, and role	ms	30s
WriteBlockOpNumOps	User, cluster, and role	-	30s
WriteBlockOpAvgTime	User, cluster, and role	ms	30s
BlockChecksumOpNumOps	User, cluster, and role	-	30s
BlockChecksumOpAvgTime	User, cluster, and role	ms	30s
CopyBlockOpNumOps	User, cluster, and role	-	30s

Metric	Dimension	Unit	Minimum monitoring granularity
CopyBlockOpAvgTime	User, cluster, and role	ms	30s
ReplaceBlockOpNumOps	User, cluster, and role	-	30s
ReplaceBlockOpAvgTime	User, cluster, and role	ms	30s
BlockReportsNumOps	User, cluster, and role	-	30s
BlockReportsAvgTime	User, cluster, and role	ms	30s
NodeManager_AllocatedContainers	User, cluster, and role	-	30s
Containers Completed	User, cluster, and role	-	30s
ContainersFailed	User, cluster, and role	-	30s
ContainersIniting	User, cluster, and role	-	30s
ContainersKilled	User, cluster, and role	-	30s
Containers Launched	User, cluster, and role	-	30s
Containers Running	User, cluster, and role	-	30s
ActiveApplications	User, cluster, and role	-	30s
ActiveUsers	User, cluster, and role	-	30s
AggregateContainersAllocated	User, cluster, and role	-	30s
AggregateContainersReleased	User, cluster, and role	-	30s

Metric	Dimension	Unit	Minimum monitoring granularity
AllocatedContainers	User, cluster, and role	-	30s
AppsCompleted	User, cluster, and role	-	30s
AppsFailed	User, cluster, and role	-	30s
AppsKilled	User, cluster, and role	-	30s
AppsPending	User, cluster, and role	-	30s
AppsRunning	User, cluster, and role	-	30s
AppsSubmitted	User, cluster, and role	-	30s
AvailableMB	User, cluster, and role	-	30s
AvailableVCores	User, cluster, and role	-	30s
PendingContainers	User, cluster, and role	-	30s
ReservedContainers	User, cluster, and role	-	30s

**Note:**

- Monitoring data is preserved for at most 31 days.
- You can view monitoring data for a maximum of 14 consecutive days.

- Viewing monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > E-MapReduce.
 3. Click an instance name or click Monitoring Chart in the Action column to access the instance monitoring details page and view various metrics.
 4. Click the Time Range toggle button at the top of the page or use the specific selection function. You can view the monitoring data for up to 14 consecutive days.
 5. Click Zoom In in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > E-MapReduce.
 3. Click an instance name or click Monitoring Chart in the Action column to access the instance monitoring details page.
 4. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all E-MapReduce instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any E-MapReduce instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that

you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins CPU Idle rate \geq 10%, the alarm service will

check every five minutes whether the CPU idle rate within five minutes meets or exceeds 10%.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.22 Express Connect

By monitoring multiple metrics of Express Connect, such as inbound and outbound traffic, CloudMonitor helps you to monitor the running status of Express Connect. CloudMonitor automatically collects data for Express Connect metrics from the time after you purchase the Express Connect service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitoring granularity
Inbound network traffic	User and instance	Bytes	1 minute
Outbound network traffic	User and instance	Bytes	1 minute
Inbound network bandwidth	User and instance	Bits/s	1 minute
Outbound network bandwidth	User and instance	Bits/s	1 minute
Latency	User and instance	ms	1 minute
Packet loss rate	User and instance	%	1 minute



Note:

- Monitoring data is saved for up to 31 days.
- You can view the monitoring data for up to 14 consecutive days.

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Express Connect.
 3. Click an instance name or click Monitoring Chart in the Action column to access the instance monitoring details page and view metrics.
 4. Click the Time Range toggle button on the top of the page or use the specific selection function. Up to 14 consecutive days of metric data can be viewed.
 5. Click Zoom In button in the upper-right corner of the monitoring chart to view a large image.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Express Connect.
 3. Click an instance name or click Monitoring Chart in the Action column to access the instance monitoring details page and view metrics.
 4. Click the bell button in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Express Connect.
 3. Select instances on the instance list page. Click Set Alarm Rules to add multiple alarm rules.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all Express Connect instances under your name. For example, if you set the resource

range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Express Connect instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 5mins Inbound Bandwidth \geq 100 Mbit/s, the alarm

service will check every five minutes whether the inbound bandwidth within five minutes meets or exceeds 100 Mbit/s.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP Callback:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.23 Function Compute

By monitoring multiple metrics of Function Compute, such as the service-level and function-level TotalInvocations, average Duration, and distribution of request status, CloudMonitor helps you to monitor the running status of Function Compute. CloudMonitor automatically collects data for Function Compute metrics from the time after you purchase the Function Compute service.

Monitoring service

- Metrics

Metric	Dimension	Unit	Minimum monitor granularity
Billableinvocations	User, service, and function	Count	1 minute
BillableInvocationsRate	User, service, and function	Percent	1 minute
ClientErrors	User, service, and function	Count	1 minute
ClientErrorsRate	User, service, and function	Percent	1 minute
ServerErrors	User, service, and function	Count	1 minute
ServerErrorsRate	User, service, and function	Percent	1 minute
Throttles	User, service, and function	Count	1 minute
ThrottlesRate	User, service, and function	Percent	1 minute
Totalinvocations	User, service, and function	Count	1 minute
Average duration	User, service, and function	Millisecond	1 minute



Note:

- Monitoring data is saved for up to 31 days.
- You can view the monitoring data for up to 14 consecutive days.

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Function Compute.
 3. Click Service List to view the monitoring information on the Service or Function level.

Alarm service

CloudMonitor provides alarm functions for Function Compute metrics, so you are notified immediately in case of any metric exceptions.

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Function Compute.
 3. In the Service or Function list, click Monitoring Chart in the Action column to access the monitoring details page.
 4. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding monitoring metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. In the left-side navigation pane, choose Cloud Service Monitoring > Function Compute.
 3. Click Alarm Rules to go to the Alarm Rules list page. Click Create in the upper-right corner to create alarm rules.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all Function Compute instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any Function Compute instance

exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.

■ **Instances:** Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 1mins Maximum Memory Usage Value \geq 1024 MBytes,

the alarm service will check every minute whether the maximum memory usage value within one minute meets or exceeds 1024 MBytes.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.24 StreamCompute

By monitoring service latency, failover rate, and read and write RPS metrics, CloudMonitor helps you gain insights into the overall performance of the StreamCompute services you are using and set alarm rules accordingly. CloudMonitor automatically collects data from StreamCompute from the time when you begin to use this product.

Monitoring service

- Metrics

Metric	Dimensions	Unit	Description	Minimum monitoring frequency
Service latency	Project, job	s	The data processing latency of the current job	1 minute
Read RPS	Project, job	read/s	The average number of data lines read per second for tasks	1 minute
Write RPS	Project, job	write/s	The average number of data lines written per second for tasks	1 minute
Failover rate	Project, job	%	The sum of failover frequency of current job	1 minute



Note:

- Monitoring data is saved for up to 31 days. You can view up to 14 consecutive days of monitoring data.

- View monitoring data
 1. Log on to the [CloudMonitor console](#).
 2. Go to the StreamCompute instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click the Time Range toggle button from the upper menu of the page or use the precise selection function. You can view monitoring data from up to 14 consecutive days.
 5. Click Zoom In in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Parameters
 - **Metrics:** The monitoring metrics imported from StreamCompute.
 - **Statistical cycle:** The recurring period of time in which the alarm system checks whether monitoring data has exceeded the alarm threshold.
 - **Statistical method:** The calculation method and resulting value used to determine whether the data has exceeded the threshold specified in an alarm rule, which can be average, maximum, minimum, or sum.
 - **Consecutive times:** An alarm is triggered after a metric value continuously exceeds the threshold specified in an alarm rule for some set of consecutive cycles. For example, if the consecutive times is set to three, then the conditions specified for an alarm rule must be met for three consecutive statistical cycles before an alarm is triggered.
- Set an alarm rule
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the StreamCompute instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column.
 4. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding metrics of this instance.

- Set multiple alarm rules
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the StreamCompute instance list under Cloud Service Monitoring.
 3. Select the appropriate instances on the instance list page. Click Set Alarm Rules to add multiple alarm rules.

6.25 ApsaraDB for HybridDB

By monitoring the CPU, disk, and memory usage, along with other metrics from ApsaraDB for HybridDB, CloudMonitor helps you quickly learn about product-wide instance usage and set alarm rules based on your specific requirements. CloudMonitor automatically collects data from ApsaraDB for HybridDB from the time when you begin to use this product.

Monitoring service

- Metrics

Metric	Dimensions	Unit	Minimum monitoring frequency
Disk usage	User and instance	%	5 minutes
Connection usage	User and instance	%	5 minutes
CPU usage	User and instance	%	5 minutes
Memory usage	User and instance	%	5 minutes
I/O throughput	User and instance	%	5 minutes



Note:

- Monitoring data is saved for up to 31 days.
- You can view up to 14 consecutive days of monitoring data.

- View monitoring data
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the Apsara for HybridDB instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click the Time Range toggle button from the upper menu of the page or use the precise selection function. You can view monitoring data from up to 14 consecutive days.
 5. Click the Zoom In button in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Parameters
 - Metrics: The monitoring metrics taken from ApsaraDB for HybridDB.
 - Statistical cycle: The recurring period of time in which the alarm system checks whether monitoring data has exceeded the alarm threshold.
 - Statistical method: The calculation method and resulting value used to determine whether the data has exceeded the threshold specified in an alarm rule, which can be average, maximum, minimum, or sum.
 - Consecutive times: An alarm is triggered after a metric value continuously exceeds the threshold specified in an alarm rule for some set of consecutive cycles. For example, if the consecutive times is set to three, then the conditions specified for an alarm rule must be met for three consecutive statistical cycles before an alarm is triggered.
- Set an alarm rule
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the Apsara for HybridDB instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click the bell icon in the upper-right corner of the monitoring chart or Create Alarm Rules in the upper-right corner of the page to set an alarm rule for corresponding metrics of this instance.

- Set multiple alarm rules
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the ApsaraDB for HybridDB instance list under Cloud Service Monitoring.
 3. Select the appropriate instances on the instance list page. Click Set Alarm Rules to add multiple alarm rules.

6.26 NAT Gateway

By monitoring multiple metrics from NAT Gateway, including SNAT connections and bandwidth package data, CloudMonitor helps you understand the overall network usage and performance of the NAT Gateway services you are using and set alarm rules accordingly. CloudMonitor automatically collects data from NAT Gateway from the time you begin to use this product.

Monitoring Service

- Metrics

Metric	Dimensions	Units	Minimum monitoring frequency
SNAT connections	User and instance	count/minute	1 minute
Bandwidth packets (inbound bandwidth)	User and instance	bit/s	1 minute
Bandwidth packets (outbound bandwidth)	User and instance	bits/s	1 minute
Bandwidth packets (inbound packets)	User and instance	packet/s	1 minute
Bandwidth packets (outbound packets)	User and instance	packet/s	1 minute
Bandwidth packets (outbound bandwidth usage)	User and instance	%	1 minute



Note:

- Monitoring data is saved for up to 31 days.
- You can view up to 14 consecutive days of monitoring data.

- View monitoring data.
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the NAT Gateway instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
 4. Click the Time Range toggle button from the upper menu or use the precise selection function. You can view monitoring data from up to 14 consecutive days.
 5. Click Zoom In in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the NAT Gateway instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page.
 4. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the NAT Gateway instance list under Cloud Service Monitoring.
 3. Select the appropriate instances on the instance list page. Click Set Alarm Rules to add multiple alarm rules.

- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.
 - Alarm Rule: the alarm rule name
 - Rule Describe: the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 1min Average Number of SNAT Connections >= 100

times, the alarm service will check every one minute whether the average number of SNAT connections within one minute meets or exceeds 100 times.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

6.27 Open Ad

By monitoring more than a dozen metrics taken from Open Ad, including RTB PV and QPS, and ad click PV, CloudMonitor helps you manage and interpret the real-time status of your Open Ad services and set alarm rules accordingly. CloudMonitor automatically collects data from Open Ad from the time when you begin to use this product.

Monitoring service

- Metrics

Metrics	Dimension	Unit	Minimum monitoring frequency
RTB PV	User	count	1 minute
RTB QPS	User	time/s	1 minute
Ad click PV	User	count	1 minute
Ad click QPS	User	time/s	1 minute
Ad click delay	User	ms	1 minute
Ad exposure PV	User	count	1 minute
Ad exposure QPS	User	time/s	1 minute
Ad exposure delay	User	ms	1 minute
DMP active crowd count	User	count/day	1 hour
DMP valid crowd requests	User	time/day	1 hour
Storage space utilized by DMP	User	byte/day	1 hour
League and dip effective crowd count	User	count/day	1 hour

Metrics	Dimension	Unit	Minimum monitoring frequency
Valid audience number in Umeng and DIP	User	time/day	1 hour

**Note:**

- Monitoring data is saved for up to 31 days.
- You can view up to 14 consecutive days of monitoring data.

- View monitoring data

1. Log on to the [CloudMonitor Console](#).
2. Go to the Open Ad monitoring page under Cloud Service Monitoring, and view the monitoring data of the Open Ad service.

Alarm service

CloudMonitor provides alarm functions for Open Ad monitoring metrics, so that you can be notified immediately in the case of any metric exceptions.

Set alarm rules

- Method 1

1. Log on to the [CloudMonitor Console](#).
2. Go to the Open Ad monitoring page under Cloud Service Monitoring.
3. Click the bell icon in the upper-right corner of the monitoring chart or New Alarm Rule in the upper-right corner of the page to set an alarm rule for corresponding metrics of this instance.

- Method 2

1. Log on to the [CloudMonitor Console](#).
2. Go to the Open Ad monitoring page under Cloud Service Monitoring.
3. Click Alarm Rules to go to the Alarm Rules list page. Click Create Alarm Rules in the upper-right corner to create alarm rules.

6.28 HybridDB for MySQL

By monitoring several metrics, specifically disk usage, inbound and outbound bandwidth and QPS, CloudMonitor helps you understand the status of your instances in HybridDB for MySQL scaling groups and set alarm rules accordingly. CloudMonitor automatically collects data from HybridDB for MySQL from the time when you begin to use this product.

Monitoring service

- Metrics

Metric	Dimensions	Unit	Minimum monitoring frequency
Disk usage	User and instance	byte	5 minutes
Inbound bandwidth	User and instance	byte/s	5 minutes
Outbound bandwidth	User and instance	byte/s	5 minutes
QPS	User and instance	count/s	5 minutes



Note:

- Monitoring data is saved for up to 31 days.
- You can view up to 14 consecutive days of monitoring data.

- View monitoring data.

1. Log on to the [CloudMonitor Console](#).
2. Go to the HybridDB for MySQL instance list under Cloud Service Monitoring.
3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page and view various metrics.
4. Click the Time Range toggle button from the upper menu of the page or use the precise selection function. You can view monitoring data from up to 14 consecutive days.
5. Click Zoom In in the upper-right corner of the monitoring chart to enlarge the chart.

Alarm service

- Set an alarm rule.
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the HybridDB for MySQL instance list under Cloud Service Monitoring.
 3. Click an instance name or click Monitoring Chart in the Actions column to access the instance monitoring details page.
 4. Click the bell icon or New Alarm Rule in the upper-right corner of the monitoring data page to set an alarm rule for corresponding metrics of this instance.
- Set multiple alarm rules.
 1. Log on to the [CloudMonitor Console](#).
 2. Go to the HybridDB for MySQL instance list under Cloud Service Monitoring.
 3. Select the appropriate instances on the instance list page. Click Set Alarm Rules to add multiple alarm rules.
- Parameters
 - Products: ECS, RDS, OSS, among others
 - Resource Range: the range for which an alarm rule takes effect. There are two alarm rule ranges available: All Resources and Instances.
 - All Resources: Indicates that the specified alarm rule applies to all HybridDB for MySQL instances under your name. For example, if you set the resource range to all resources, and set the alarm threshold for CPU usage to 80%, then an alarm is triggered when the CPU usage of any HybridDB for MySQL instance exceeds 80%. When you select All Resources, you can report alarms for up to 1,000 resources. If the number of your resources exceeds 1,000, alarms cannot be reported for some resources even if they exceed the threshold you set in your alarm rule. Therefore, for these scenarios, we recommend that you use application groups to divide resources by service before setting up alarm rules to avoid this issue.
 - Instances: Indicates that the specified rule only applies to a specific instance. For example, if you set the resource range to Instances and set the alarm

threshold for CPU usage to 80%, an alarm is triggered when the CPU usage of the specified instance exceeds 80%.

- **Alarm Rule:** the alarm rule name
- **Rule Describe:** the main content of the alarm rule where you define the alarm-triggering conditions, or value threshold, for related metrics. For example, if you describe the rule as 60mins Average CPU Usage >= 1GB, the alarm service

will check every five minutes whether the average value of CPU usage within 60 minutes meets or exceeds 1 GB.

Consider the following example. For the alarm service in host monitoring, one data point is reported in 15 seconds for a single server metric, and 20 data points in five minutes. This relates to the following alarm rules.

- **5mins Average CPU Usage > 90%:** The average CPU usage value of the 20 data points in five minutes exceeds 90%.
 - **5mins CPU Usage Always > 90%:** The CPU usage values of the 20 data points in five minutes all exceed 90%.
 - **5mins CPU Usage Once > 90%:** The CPU usage value of at least one of the 20 data points in five minutes exceeds 90%.
 - **Total 5mins Internet Outbound Traffic > 50 MB:** The sum of the outbound traffic values of the 20 data points in five minutes exceeds 50 MB.
- **Mute For:** the period of time that an alarm is muted so that alarm contacts do not receive any alarm notifications during this period. An alarm rule can be muted for up to 24 hours (or 1 day).
 - **Triggered when threshold is exceeded for:** An alarm notification is sent if the detected values reach the alarm rule threshold a certain number of times consecutively.
 - **Effective Period:** the period of time for which an alarm rule is effective. During this period of time, the alarm service checks metric data and determines whether to generate an alarm.
 - **Notification Contact:** a group of contacts who receive alarm notifications.
 - **Notification Methods:** Emails and DingTalk chatbot.
 - **Email Subject:** The email subject is set as the product name, metric, and instance ID involved in the alarm by default.
 - **Email Remark:** supplementary information customized for an alarm email. Remarks are sent as part of the alarm notification email body.
 - **HTTP CallBack:** Enter a URL accessible through the Internet and CloudMonitor will push the alarm information to the address through a POST request. Currently, only HTTP is supported.

7 RAM for CloudMonitor

RAM permissions are supported in CloudMonitor. Through the integration of the monitoring console with access control features, you can easily and quickly apply permissions for cloud service monitoring data, alarm rule management, alarm contact and alarm contact groups, and event subscription and related features.



Note:

RAM monitoring data queries are supported for the following cloud products:

- ECS
- RDS
- Server Load Balancer
- OSS
- CDN
- ApsaraDB for Memcache
- EIP
- ApsaraDB for Redis
- Message Service
- Log Service

Permissions

In RAM, if a user is authorized with read-only permissions for CloudMonitor , the user can only view relevant data, such as the monitoring data and alarm services, but cannot write data.

Authentication types

In addition to basic RAM account permission controls, time-based, multi-factor, and IP authentication are supported.

Resources

Fine-grained resource descriptions are not supported by RAM. The “*” wildcard is used for resource authorization.

Operation description

- Monitoring data

Data query actions are divided into two categories: Product instance lists and CloudMonitor metric data queries. When authorizing a RAM account to log on to the CloudMonitor portal and view metric data, you must also grant the account permissions for the corresponding product's instance list and metric data query.

The corresponding actions are listed in the following table.

Product	Action
CMS	QuerMetricList
CMS	QueryMetricLast
ECS	DescribeInstances
RDS	DescribeDBInstances
SLB	DescribeLoadBalancer*
OSS	ListBuckets
OCS	DescribeInstances
EIP	DescribeEipAddresses
Aliyun Cloud for Redis	DescribeInstances
Message Service	ListQueue
CDN	DescribeUserDomains

- Alarm service

The alarm service provides permission controls for alarm rule management, alarm contact and alarm contact group management, and event subscription and related features.

The query-related actions are listed in the following table.

Action	Description
QueryAlarm	Query an alarm rule
QueryAlarmHistory	Query an alarm history
QueryContactGroup	Query a contact group
QueryContact	Query a contact
QuerySms	Query the number of SMSs used

Action	Description
QueryMns	Querying an event subscription configuration

The management-related actions are listed in the following table.

Action	Description
UpdateAlarm	Modify an alarm rule
CreateAlarm	Create an alarm rule
DeleteAlarm	Delete an alarm rule
DisableAlarm	Disable an alarm rule
EnableAlarm	Enable an alarm rule
CreateContact	Create a contact
DeleteContact	Delete a contact
UpdateContact	Modify a contact
SendEmail	Send an email authentication code
SendSms	Send an SMS verification code
CheckEmail	Check an email verification code
CheckSms	Check an SMS verification code
CreateGroup	Create a contact group
DeleteGroup	Delete a contact group
UpdateGroup	Modify a contact group
CreateMns	Create an event subscription
DeleteMns	Delete an event subscription
UpdateMns	Modify an event subscription

8 Application groups

8.1 Application group overview

The application group feature of CloudMonitor allows you to group related resources and monitor these resources in a centralized manner. With application groups, you can easily monitor a group of target resources such as servers, databases, SLB instances, and storage, and apply alarm rules to the application group, thereby improving your overall O&M efficiency.



Note:

- A single account can create up to 100 application groups.
- Up to 1,000 resource instances can be added to one application group.

8.2 Create application groups

This topic describes how to group your cloud resources by creating application groups so that you can manage your resources and alarm rules on a grouped basis.

Scenarios

If you have purchased multiple products on Alibaba Cloud, you can group them together in a centralized manner by creating application groups. With application groups, you can manage resources of different regions and products, such as servers, databases, object storage, and cache, based on your business modules. In addition, you can easily manage alarm rules and view the monitoring data of these grouped resources.

Application group modes

Instances can be added to application groups using dynamic or static mode.

- **Dynamic mode:** When creating an application group, you can set name rules for instances so that instances which meet your name rules will be automatically added into the application group. If you want to add or remove instances to or from the group in the future, you only need to modify the instance names to complete these configurations. Currently, dynamic mode is supported only by ECS, ApsaraDB for RDS, and SLB instances.

- **Static mode:** With static mode, you need to manually add instances to an application group.

Create an application group



Note:

- Up to 1,000 resource instances can be added to each application group.
- Up to 100 application groups can be created under each account.

Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups.

- 3. In the upper-right corner of the displayed page, click Create Group.

Create Group

Basic Information

- Product Group Name
- Contact Group
 [?](#) [Quickly create](#)

MonitorAlarm

Select Template
 [Go to Create Alarm](#)

Initialize Agent Installation [?](#)

Event Monitor

- Subscribe Event notification
After subscription event notification, alarm notification will be sent when service is abnormal within the group. [Introduction to Cloud Products Events](#)

Add Instance dynamically

- Dynamic rules for ECS instances
- Dynamic rules
 - All rules Any rule
 - [!](#) instance created in future according with this rule would be added to group

[+Add Rules](#)

4. **Enter Basic Information:** Enter the group name and select one or more contact groups to receive alarm notifications.
5. **Set MonitorAlarm:** Select one or more templates to initialize alarm rules for the instances in the group (optional), and select the notification method. If you turn on the Initialize Agent Installation switch, the CloudMonitor agent will be installed on all servers in the group to collect monitoring data.
6. **Set Event Monitor:** If you select the Subscribe Event notification check box, alarm notifications will be sent when critical-level and warning-level events occur in related resources in the group.
7. **Set Add Instance dynamically.**
 - You can set name rules to automatically add ECS instances that match the name rules to the group. Specifically, instances, including future instances, whose names contain, start with, or end with the words you specify will be automatically added to the group. A maximum of three rules can be added, and the relationship among the rules can be AND or OR.
 - To add rules for ApsaraDB for RDS or SLB instances, click Add Product.
 - To add instances of other Alibaba Cloud products, you need to add them manually after creating the application group.
8. **Click Create Application Group.**

8.3 Check application group details

The group details page contains the fault list, alarm history, alarm rules, group resources, events, and group resource metric data. You can use this page to monitor the preceding details of your application groups.

Group list

All application groups on CloudMonitor, along with the resources and health status of each group, are displayed on the group details page.

Parameters

- **Group name (or ID):** The name or identification number of an application group.
- **Health status :** The alarm status of any group resource. An application group is healthy when no active alarms are triggered for any of the resources in the

group, but unhealthy whenever any metric threshold of a resource in the group is met and an alarm is triggered.

- **Instance count** : The total number of instances in an application group, both ECS and non-ECS instances.
- **Resource types** : The number of resource types in an application group. For example, if an application group contains ECS, ApsaraDB for RDS, and Server Load Balancer instances, then this number is three.
- **Unhealthy instances** : The total number of instances with active alarms in an application group. For example, if two ECS instances and one ApsaraDB for RDS instance have active alarms, the number of unhealthy instances is three.
- **Creation time** : The time when an application group is created.
- **Actions** : The actions that can be applied to an application group. Action types supported are manage, stop notifications, enable and disable all the alarm rules, and delete group.

Exception list

The resources with active alarms in your group are displayed in the fault list to help you to easily view unhealthy instances and quickly troubleshoot the causes.



Note:

- When multiple metrics of a resource have active alarms at the same time, the fault list displays the resource multiple times. Each row of the list shows a metric with an active alarm.
- Once you disable an alarm rule with an active alarm, the resources and metrics associated with the rule no longer appearing on the fault list.

Parameters

- **Faulty resource** : A resource with an active alarm.
- **Start time** : The time when the first alarm is generated for the resource.
- **Status** : Indicates whether a resource has an active alarm.
- **Duration** : The period of time when a faulty resource is in an alarm state.
- **Alarm rule name** : The name of the alarm rule applied to a faulty resource.

- **Actions** : The actions that can be applied to a faulty resource. You can click **Expand** to view the metric trends of a faulty resource with an active alarm over the past six hours, and compare the metric data with the alarm threshold value.

Alarm history

Alarm history provides the account of all the alarm rules applied to a group.



Note:

You can request the alarm history of the last three days. If the interval between the query start time and end time exceeds three days, the system prompts you to re-select the time range.

Parameters

- **Faulty resource** : A resource with an active alarm.
- **Duration** : The time during which a faulty resource is in an alarm state.
- **Occurrence time** : The time when the alarm is generated.
- **Alarm rule name** : The name of the alarm rule applied to a faulty resource.
- **Notification method** : The method by which alarm notifications are sent, which are SMS, email, and TradeManager.
- **Product type** : The product type to which a faulty resource belongs.
- **Status** : The status of the alarm rule, which are alarm status, cleared status, and muted states.
- **Notification target** : The group of contacts who receive alarm notifications.

Alarm rules

A list of all the alarm rules applied to a group is displayed in an alarm rules list. You can select the preferred alarm rule from the list and can enable, disable, or modify the rules based on your requirements.



Note:

The alarm rules list only shows the alarm rules applied to a specific application group. It does not show the alarm rules with **Resource Range** set to the **All Resources** or **Instance**.

Parameters

- **Alarm name** : Name of an alarm rule specified when the alarm rule was created.
- **Status** : Displays whether the resources associated with the alarm rules have active alarms.
 - **Normal state**: All resources associated with the alarm rules are normal.
 - **Alarm state**: At least one instance associated with the alarm rule has an active alarm.
 - **Insufficient data**: At least one instance associated with the alarm rule has insufficient data and no instance has an active alarm.
- **Enable** : Shows whether the alarm rule is enabled.
- **Product name** : The name of the product to which group resources belong.
- **Alarm description** : A brief description of alarm rules setting.
- **Actions** : The optional operations include Modify, Enable, Disable, Delete, and Alarm History.
 - **Modify**: Click to make changes in the alarm rule.
 - **Disable**: Click to disable the alarm rule. Once the alarm rule is disabled, the alarm service does not check whether metric data exceeds the threshold value.
 - **Enable**: Click to enable the alarm rule. Once you enable a previously disabled alarm rule, the alarm service checks the metric data and determines whether to trigger an alarm based on the alarm rule.
 - **Delete**: Click to delete the alarm rule.
 - **Alarm History**: Click to view the alarm history of the alarm rule.

Group resources

Display all the resources of a group and the health condition of these resource.

Parameters

- **Instance name (or ID)** : The instance name or ID of a resource.
- **Health status** : The alarm status of any group resource. An application group is healthy when no alarms are triggered for any of the resources in the group, but unhealthy whenever an alarm is triggered for any resource in the group.

Events

Alarm history and records for alarm rule operation events, such as add, modify, and delete actions, are supported, allowing you to trace any operation performed on a specific alarm rule.



Note:

You can query event information from the last 90 days.

Parameters

- `Occurrence time` : The time when an event occurred.
- `Event name` : The name of an event, which may be an alarm event such as alarm generated or alarm cleared, or an system event such as create alarm rule, modify alarm rule, or delete alarm rule.
- `Event type` : The type of event, which can be divided into system events and alarm events. Types of system events include create alarm rule, delete alarm rule, and modify alarm rule. Types of alarm events include alarm generated and alarm cleared.
- `Event details` : Detailed information associated with an event.

Charts

The lower area of the application group details page displays the monitoring details of group resources. By default, CloudMonitor initializes frequently used metric data. You can choose to customize the area, changing the chart type and metric data displayed.



Note:

To obtain the OS metrics of ECS, you must install the CloudMonitor agent.

Initialized metric data

By default, CloudMonitor initiates the following application group data, which are all displayed in line charts. If you want to view more metric data, click Add Metric Chart to add more metrics to the data.

Product	Metrics	Chart type	Description
ECS	CPU usage and outbound bandwidth (Internet)	Line chart	Displays the aggregate data of all servers in the group.
ApsaraDB for RDS	CPU usage, disk usage, IOPS usage, connection usage	Line chart	Displays the data of a single database instance.
Server Load Balancer	Outbound bandwidth and inbound bandwidth	Line chart	Displays the data of a single Server Load Balancer instance.
OSS	Storage size and GET/PUT request count	Line chart	Displays the data of a single bucket.
CDN	Downstream bandwidth and hit rate	Line chart	Displays the data of a single domain name.
EIP	outbound bandwidth (Internet)	Line chart	Displays the data of a single instance.
ApsaraDB for Redis	Memory usage, connection usage, and QPS usage	Line chart	Displays the data of a single instance.
ApsaraDB for MongoDB	CPU usage, memory usage, IOPS usage, and connection usage	Line chart	Displays the data of a single instance.

8.4 Modify an application group

Scenarios

When your applications use more cloud products to meet the requirements of service resizing or technical architecture improvement, you need to modify the resources in your application groups.

When the O&M and development personnel of your applications are changed, you need to modify the alarm contact groups of your application groups.



Note:

- After resources are removed from an application group, the alarm rule configured for the application group is no longer applicable to the removed instances.
- After an instance is added to a group, the instance automatically gets associated with the alarm rule configured for the application group. You do not need to create an alarm rule for the instance.

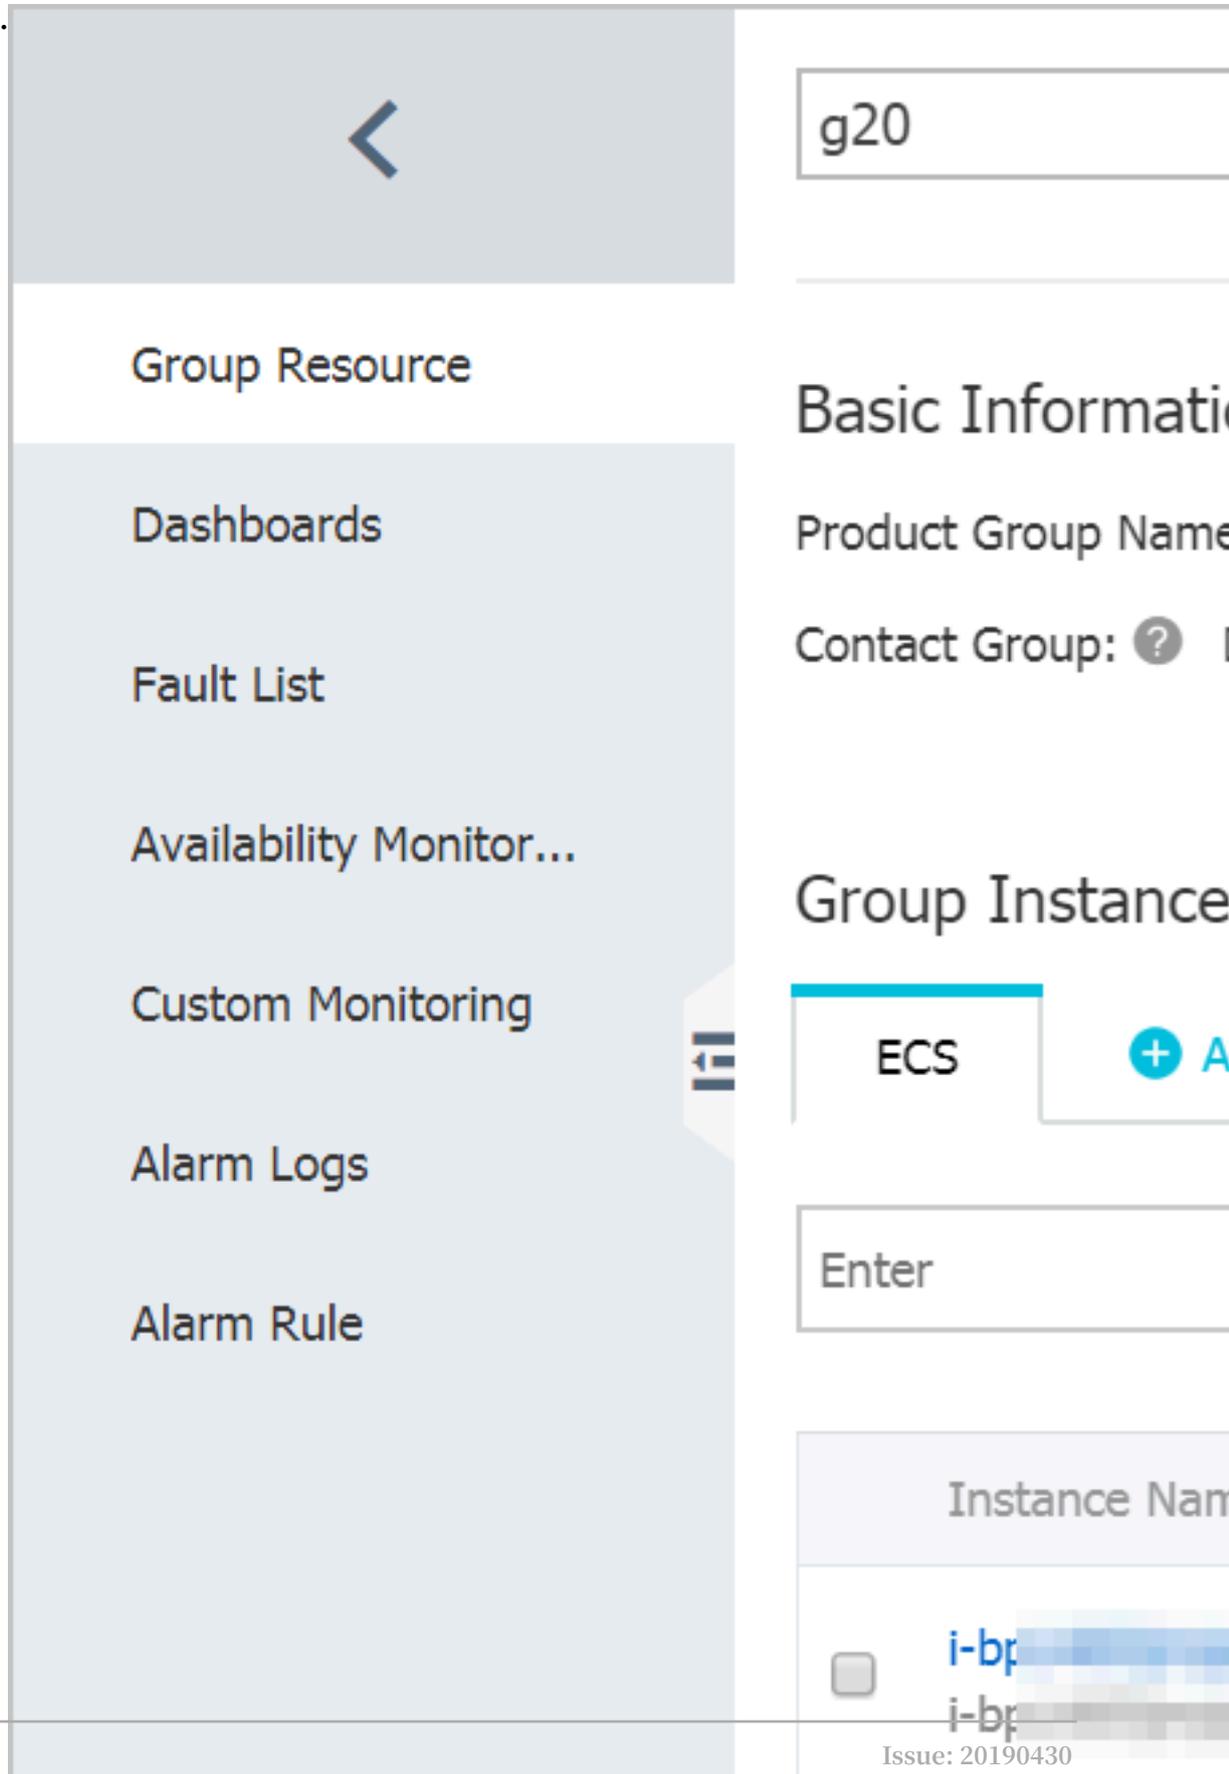
Modify basic information

To modify the application group name or the contact group, go to the details page of the target application group. In the Basic Information area, click the pencil icon next to the group name or contact group information. Modify the name or the contact group and click OK.

The screenshot shows the configuration page for an application group named 'g20'. On the left is a navigation sidebar with a back arrow at the top and the following menu items: Group Resource, Dashboards, Fault List, Availability Monitor..., Custom Monitoring, Alarm Logs, and Alarm Rule. The main content area is divided into three sections: 'Basic Information' with fields for 'Product Group Name: g20' and 'Contact Group: ? Default'; 'Group Instances' with a list containing 'ECS' and a '+ Add Prod' button; and an input field containing the text 'Enter'.

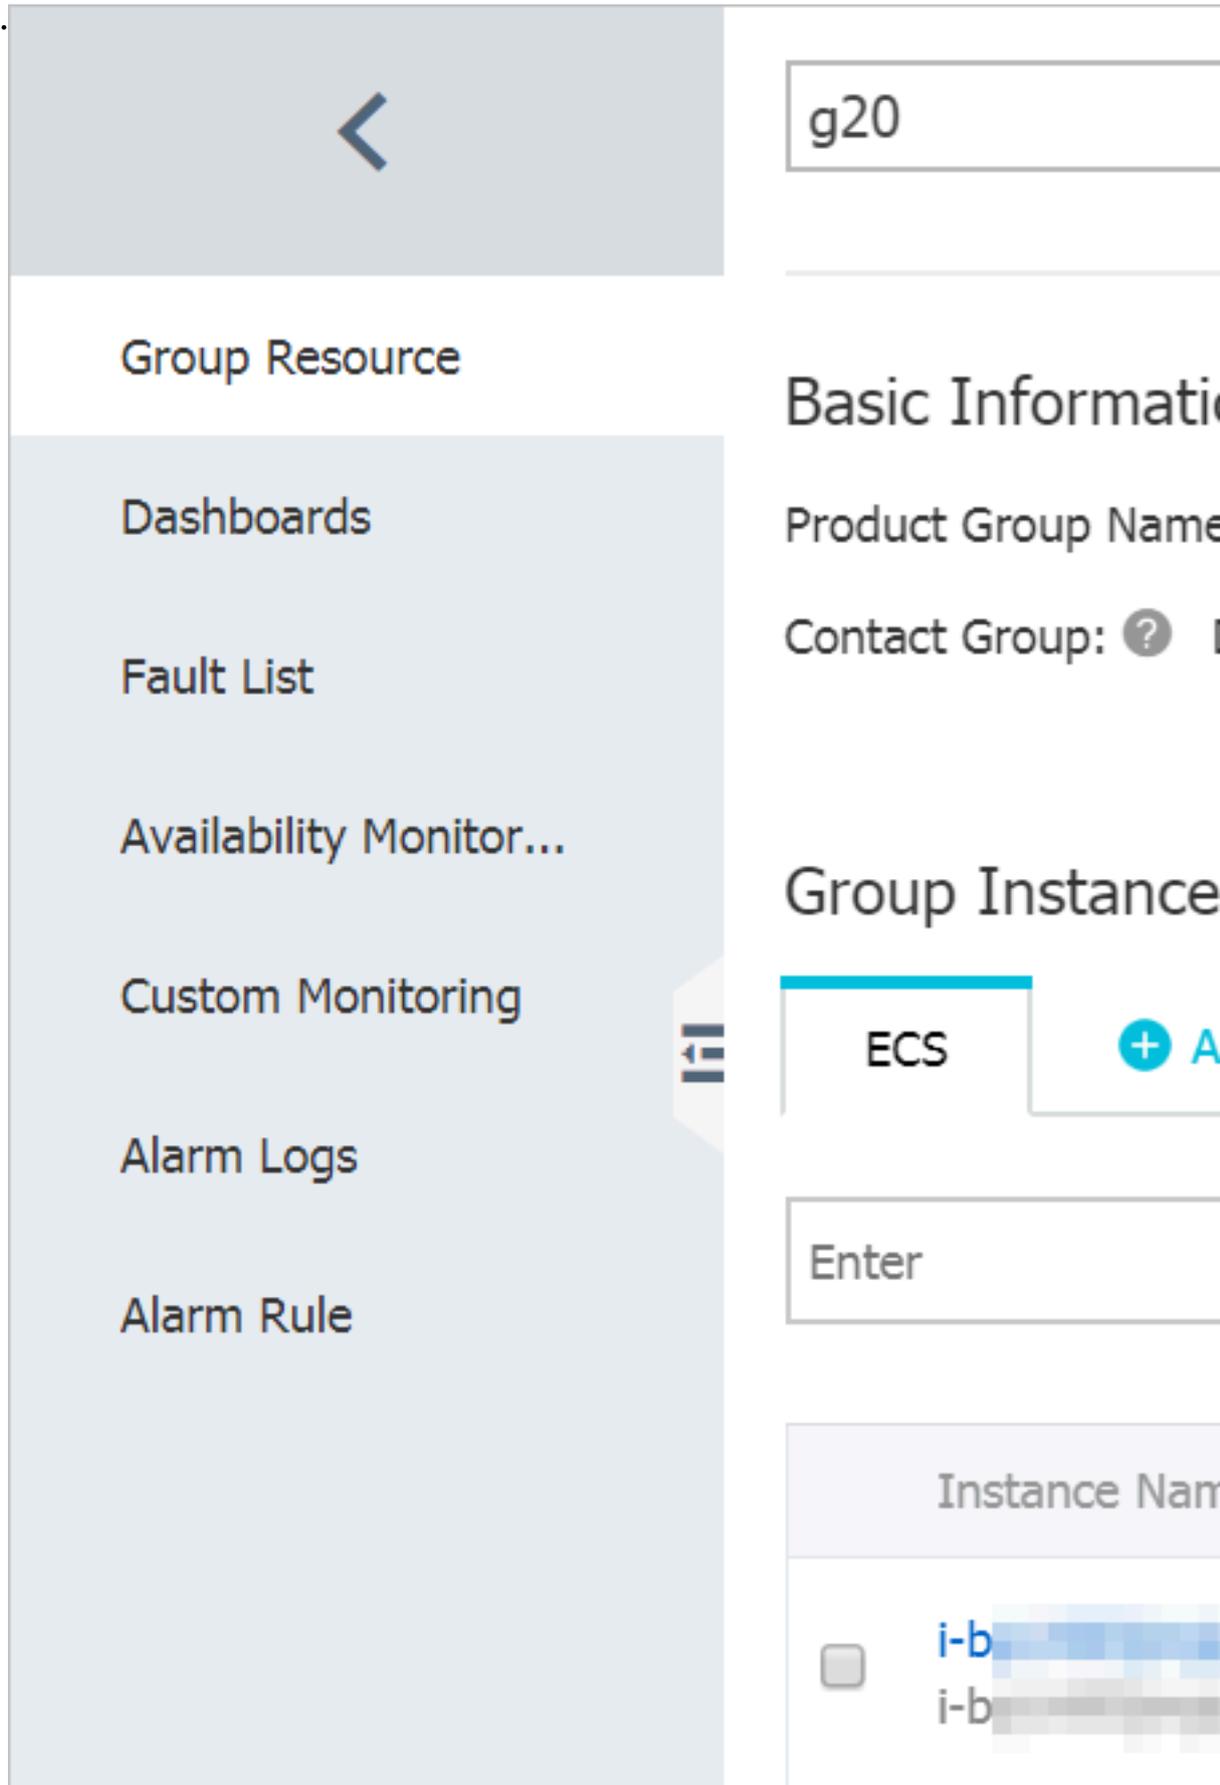
Add or remove an instance

1. To delete an instance, click the tab of the target product, find the target instance, and click Delete in the Actions column.



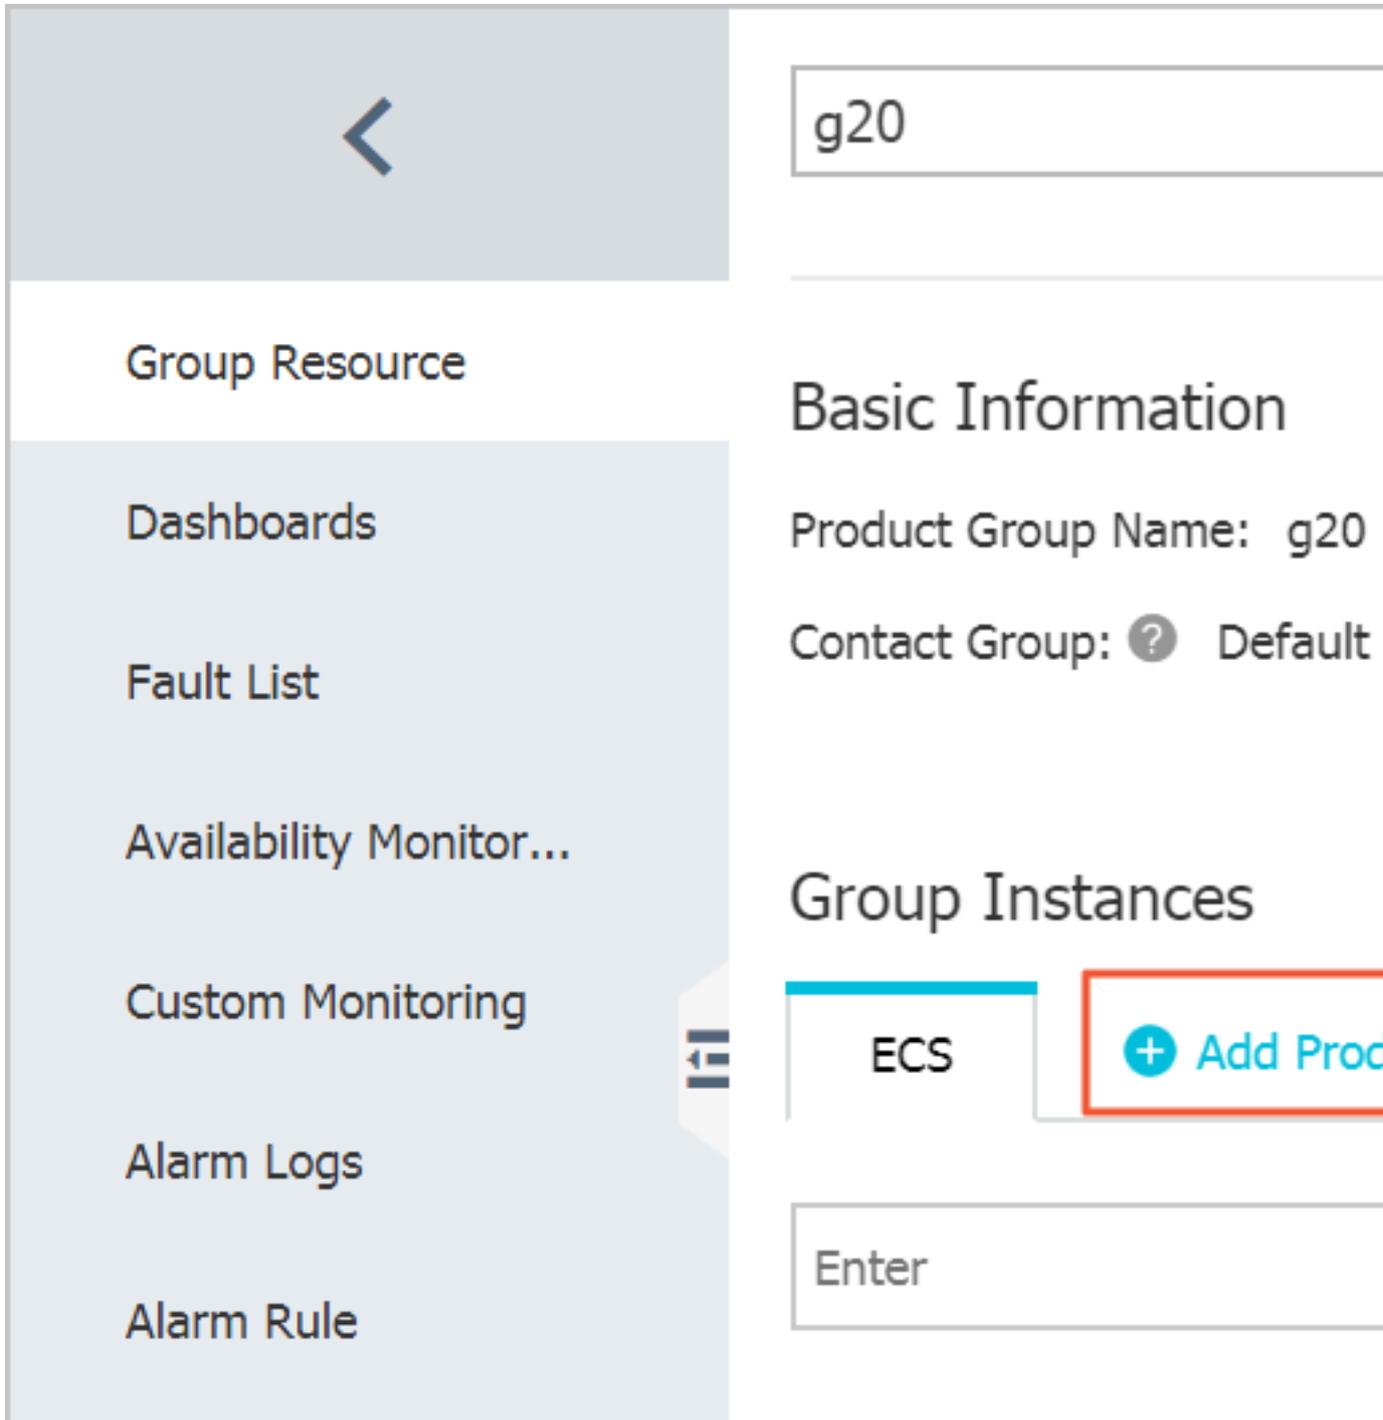
2. To add an instance, click the tab of the target product and in the upper-right corner of the tab page, click Add Instance. On the

displayed AddResource page, select the target instance and click Confirm.



Add a new product

Go to the details page of the target application group. Click Add Product. On the displayed AddResource page, select the target product and instance, and click Confirm.



8.5 Add resources to an application group

This topic describes how to add resources to an application group so that you can manage alert rules and view monitoring data by application group.

Background information

Only ECS, RDS, and SLB instances that meet the preconfigured dynamic matching rules can be automatically added to application groups. The other instances must be manually added to groups. This topic describes how to manually add an instance to an application group.

Prerequisites

- You have created the instances to be added to an application group.
- You have created an application group. If you have not created an application group, see [Create application groups](#).

Procedure

Precautions



Note:

Up to 1,000 instances can be added to an application group.

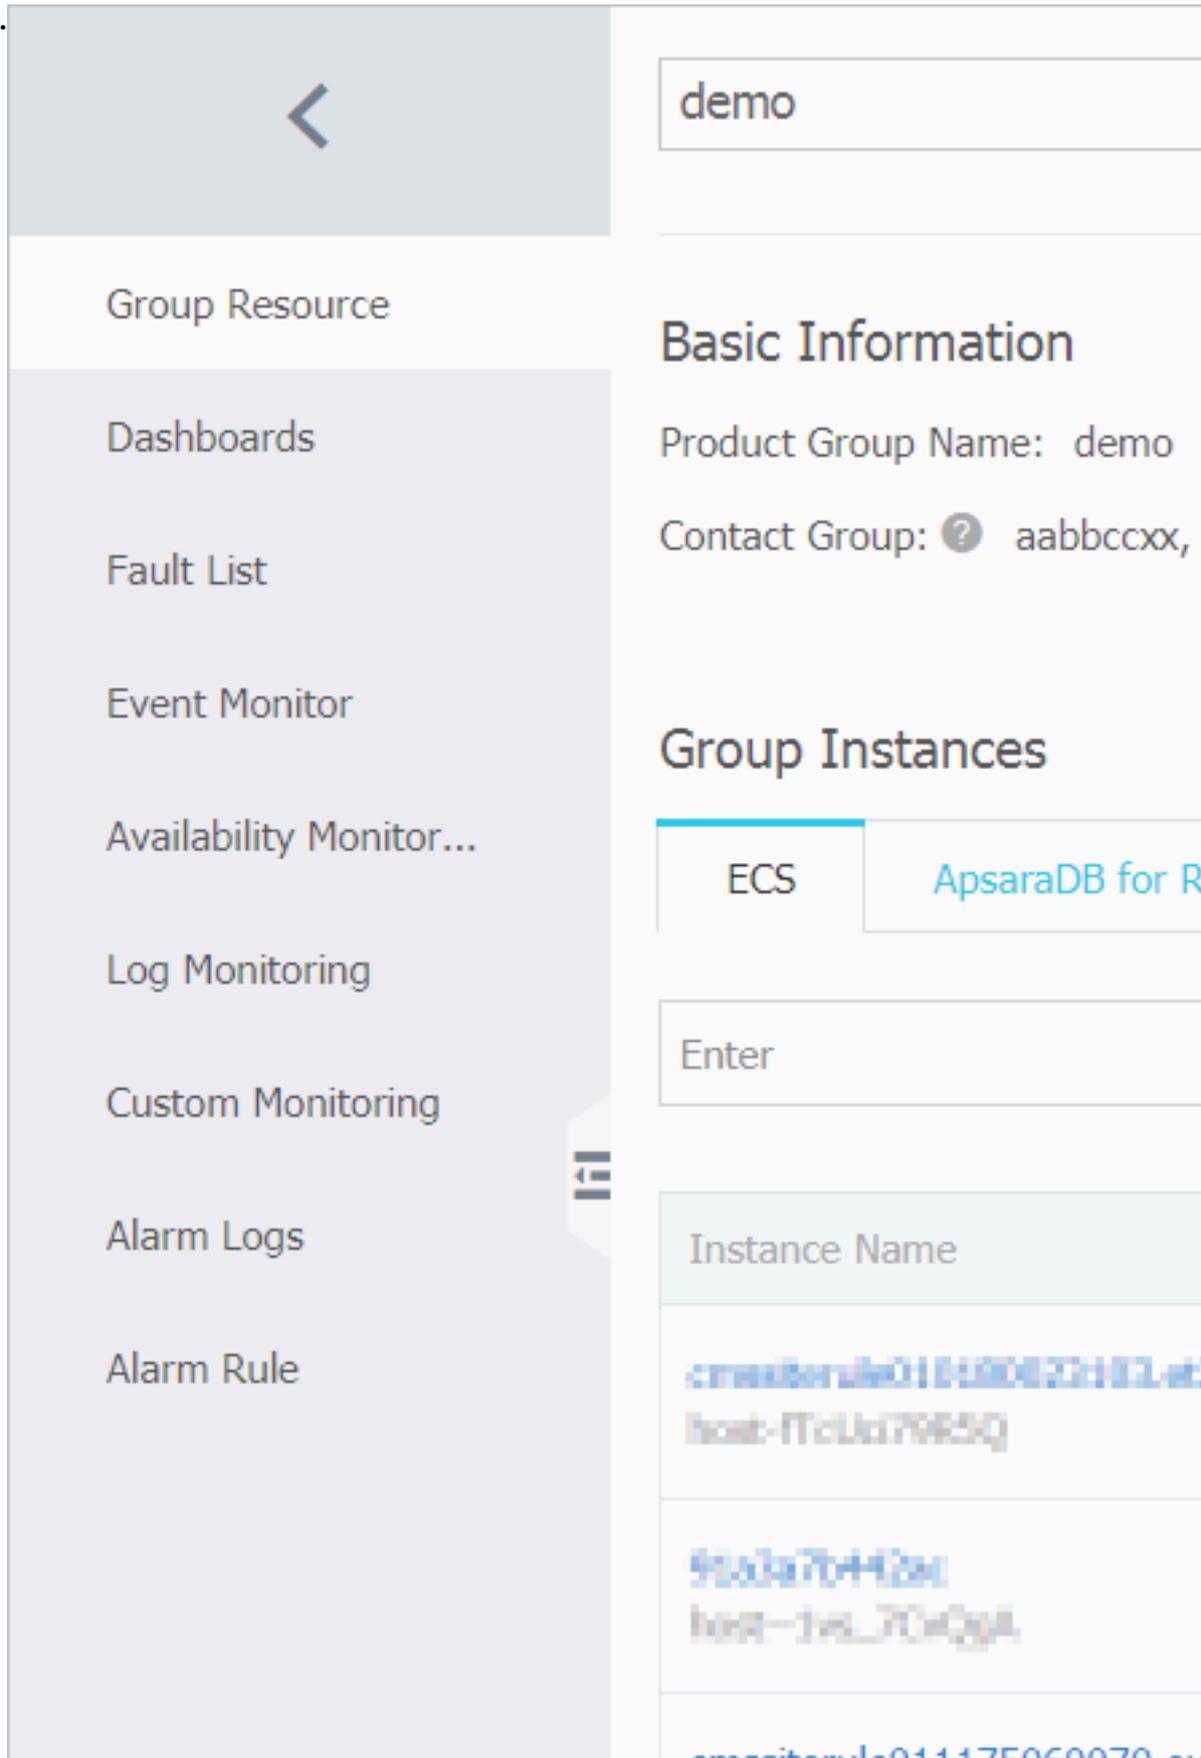
Add services

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups. The Application Groups page is displayed.
3. Click the name of the group to which you want to add resources. The Group Instances page is displayed.

The screenshot shows the 'Group Instances' page in the CloudMonitor console. The group name is 'demo'. The page includes a search bar, dynamic rules (Instance NameContainstest, Instance NameContains, Instance NameContaini), and a table of instances.

Instance Name	Health Status	Resource Description	CPU Usage(%)	Memory Usage(%)
cn-hangzhou-ecs-011890221003-ct2 host-ft10u07985Q	OK		2.15	38.72
cn-hangzhou-ecs-011890221003-ct2 host-ft10u07985Q	OK			

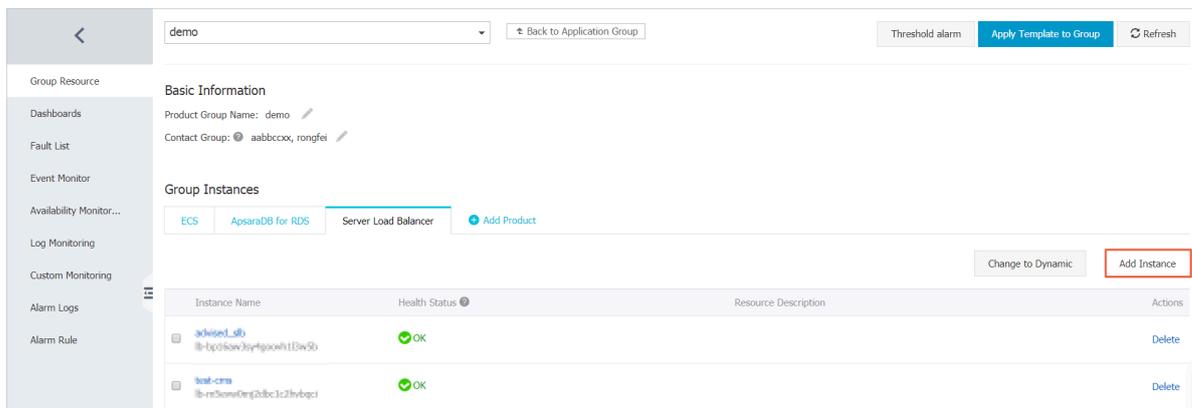
4. Click Add Product. The Add Resource page is displayed.



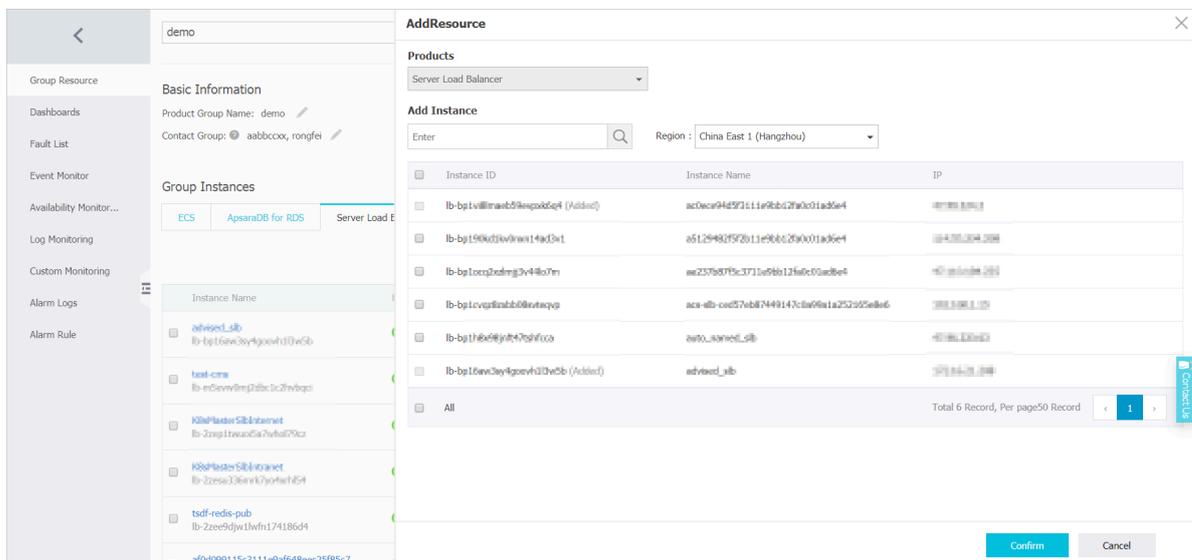
5. Select a service from the drop-down list, select the instances to be added from the instance list of the service, and click OK.

Add instances

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click **Application Groups**. The **Application Groups** page is displayed.
3. Click the name of the group to which you want to add resources. The **Group Instances** page is displayed.



4. Click the tab of a service to be added such as OSS.
5. On the tab that appears, click **Add Instance**. The **Add Resource** page is displayed.



6. Select the instances to be added and click **OK**.

8.6 Apply an alert template to an application group

This topic describes how to apply an alert template to an application group to quickly create alert rules for a business module.

Background information

If your account has a large amount of cloud resources such as ECS, RDS, SLB, and OSS instances, we recommend that you create service-related application groups and alert template and apply the alert templates to the application groups. This provides a simple way to create and maintain alert rules.

Alert templates must be used together with application groups. You can apply alert templates to application groups to quickly create alert rules for your business modules.

Prerequisites

Before you apply an alert template to an application group, you must create an alert template. For more information about how to create an alert template, see [Use alarm templates](#).

Procedure

Precautions

Alert templates must be used together with application groups. We recommend that you create application groups and alert templates for cloud resources based on your applications.



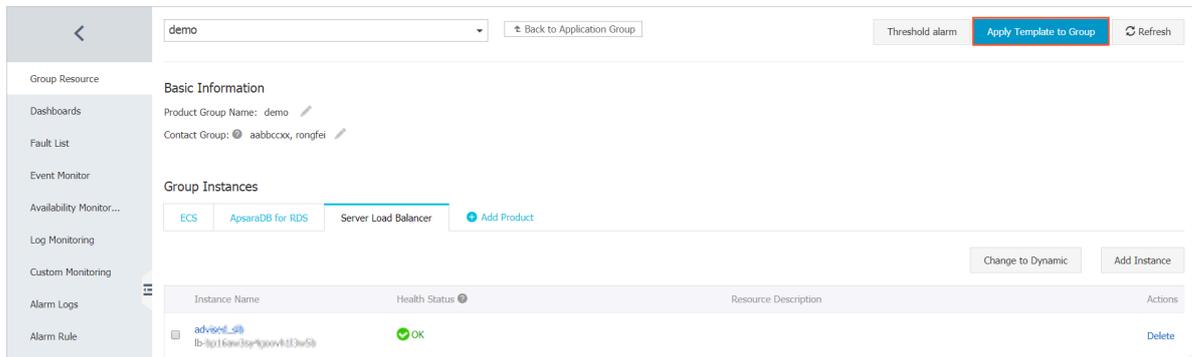
Note:

After you apply an alert template to a specified group, CloudMonitor deletes the original alert rule for this group and create a new one based on the template.

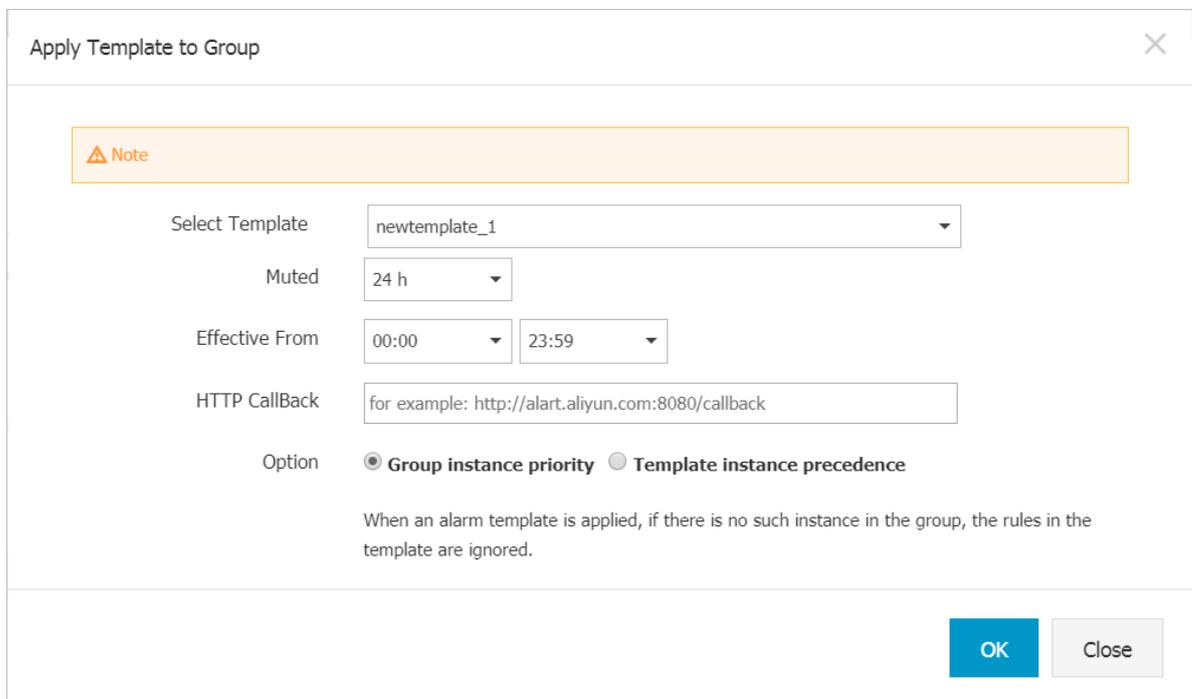
Procedure

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups. The Application Groups page is displayed.

- Click the name of the group to which you want to apply the alert template. The group details page is displayed.



- Click Apply Template to Group in the upper-right corner. The Apply Template to Group page is displayed.



- Select an alert template and click OK.

8.7 Manage alarm rules

You can create, view, modify, enable, disable, and delete threshold alarm rules in application groups.



Note:

When you view alarm rules of an application group, the system displays only the alarm rules applied to this application group. The alarm rules applied to the instances or resources in the group are not displayed.

Create an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups.
3. Find the target group and click the group name.
4. Click Threshold alarm in the upper-right corner.
5. Select the product type, add one or more alarm rules, set the alarm mechanism, select the contact group, and then click Add.

Create alarm rules by using an alarm template

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups.
3. Find the target group and click the group name.
4. In the upper-right corner of the displayed page, click Apply Template to Group.
5. Select the required alarm template and click OK.

Delete an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups.
3. Find the target application group and click the group name.
4. In the left-side navigation pane, click Alarm Rule.
5. Find the target alarm rule, and click Delete in the Actions column to delete this rule. To delete multiple rules at a time, select the rules to be deleted and click Delete under the alarm rule list.

Modify an alarm rule

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups.
3. Find the target application group and click the group name.
4. In the left-side navigation pane, click Alarm Rule.
5. Find the target alarm rule, and click Modify in the Actions column to modify this rule.

Disable or enable alarm rules

If you want to stop a service for application maintenance or upgrades, you can disable all alarm rules of the application group to avoid unnecessary alarm notifications. After the maintenance or upgrades are complete, you can enable the alarm rules.

- **Disable all alarm rules of an application group**
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Application Groups.
 3. Find the target application group and click More in the Actions column.
 4. Select Disable All Alarm Rules.
- **Enable all alarm rules of an application group**
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Application Groups.
 3. Find the target application group and click More in the Actions column.
 4. Select Enable All Alarm Rules.
- **Disable some alarm rules of an application group**
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Application Groups.
 3. Find the target application group and click the group name.
 4. In the left-side navigation pane, click Alarm Rule.
 5. Find the target alarm rule, and click Disable in the Actions column to disable this rule. Repeat this step to disable other alarm rules, or select multiple rules and click Disable under the alarm list.
- **Enable some alarm rules of an application group**
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Application Groups.
 3. Find the target application group and click the group name.
 4. In the left-side navigation pane, click Alarm Rule.
 5. Find the target alarm rule, and click Enable in the Actions column to enable this rule. Repeat this step to enable other alarm rules, or select multiple rules and click Enable under the alarm list.

9 Event monitoring

9.1 Event monitoring overview

Event monitoring covers cloud service faults, O&M events, and user business exceptions. It provides event statistics by service, level, name, and application group, to facilitate associated businesses and fault review. You can customize the receivers and methods of event notification to prevent key events from being ignored. The event details help you locate faults.

Cloud service events

Event monitoring provides you with a centralized platform to summarize and query system events that are generated by different types of cloud services. It enables you to track the use of cloud services.

After you classify resources into application groups, service-related system events are automatically associated with the resources of those groups. This helps you integrate monitoring information, and quickly analyze and troubleshoot problems.

Event monitoring also provides the event alert function. You can configure alerts based on the event level, notification through text messages, emails, or DingTalk Chatbot, or alert callbacks. With these configurations, you can be notified of critical events immediately after they occur and handle the events in an automated online O&M process.

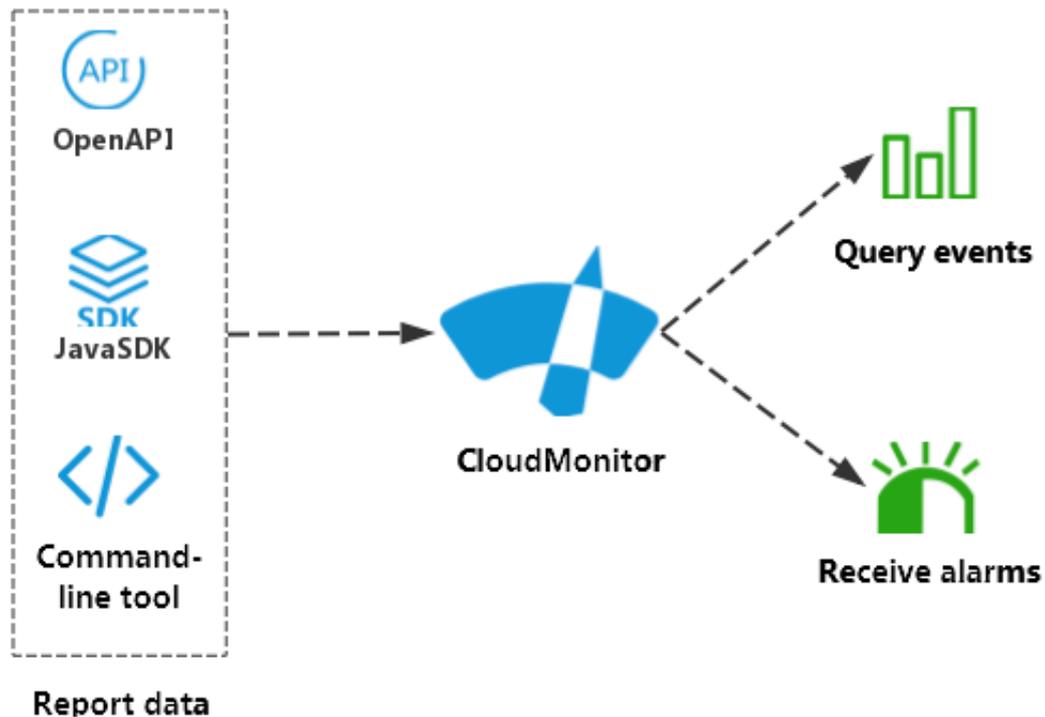
Event monitoring provides you with query and alert services for cloud service faults and O&M events.

- **ECS events:** critical ECS system events, such as unexpected restarts or disk performance degradation that are caused by system or instance errors
- **SLB events:** HTTPS certificate expiration events
- **OSS events:** The upstream or downstream bandwidth that is used by a bucket has exceeded the throttling threshold or report threshold.
- **Auto Scaling events:** successful or failed scale-in and scale-out of Auto Scaling
- **E-MapReduce events:** cluster creation failure, timeout, and service component status

For more cloud service events, see [Cloud service events](#).

Custom events

Event monitoring is able to report, query, and send alerts about events. It allows you to report exceptions or important changes in your business to CloudMonitor and receive alerts when exceptions occur.



The difference between custom event monitoring and custom monitoring is as follows:

- Custom event monitoring focuses on the data of non-continuous events.
- Custom monitoring focuses on periodically collected time series data.

Alert service and automated O&M

Event monitoring provides multiple alert methods for automated O&M.

- **Alert notification:** Alert notification can be sent through emails, DingTalk Chatbot, or other channels.
- **MNS queue:** Events can be written to the MNS queue, which can then be connected to your own O&M system.

- **Function Compute:** Events can trigger Function Compute to process subsequent O&M logic.
- **Alert callback:** Alert notification is pushed to the public URL of your existing O&M system or message notification system through HTTP POST requests. You can then handle the received alert notification based on its contents.

9.2 Cloud product events

9.2.1 Cloud service events

This topic describes the cloud service system events that are supported by event monitoring.

Elastic Compute Service system events

Name	Meaning	State	Level	Remarks
Instance: InstanceFailure. .Reboot	Beginning of instance restart due to an instance error.	Executing	CRITICAL	
Instance: InstanceFailure. .Reboot	End of instance restart due to an instance error.	Executed	CRITICAL	
Instance: SystemFailure. Reboot	Beginning of instance restart due to a system error.	Executing	CRITICAL	
Instance: SystemFailure. Reboot	End of instance restart due to a system error.	Executed	CRITICAL	
Instance: SystemMain tenance.Reboot	An instance restart is scheduled due to system maintenance.	Scheduled	CRITICAL	

Name	Meaning	State	Level	Remarks
Instance: SystemMain tenance.Reboot	The instance restart due to system maintenance is prevented.	Avoided	CRITICAL	
Instance: SystemMain tenance.Reboot	The instance restart due to system maintenance is being executed.	Executing	CRITICAL	
Instance: SystemMain tenance.Reboot	The instance restart due to system maintenance is completed.	Executed	CRITICAL	
Instance: SystemMain tenance.Reboot	The instance restart due to system maintenance is canceled.	Canceled	CRITICAL	
Instance: SystemMain tenance.Reboot	The instance restart due to system maintenance has failed.	Failed	CRITICAL	
Disk:Stalled	Beginning of disk performance impact.	Executing	CRITICAL	
Disk:Stalled	End of disk performance impact.	Executed	CRITICAL	

Name	Meaning	State	Level	Remarks
Instance: StateChange	Notification of instance state change.	Normal	INFO	The instance state is displayed in event details, including Running, Stopped (Stopped, Expired, About to Expire, Locked, Releasing, and Pending Release), and Deleted (the instance has been released). For more information about the instance lifecycle, see Instance lifecycle.

Name	Meaning	State	Level	Remarks
Instance: PreemptibleInstanceInterruption	Notification of preemptible instance interruption.	Normal	WARN	Preemptible instances enter the Pending Release state for certain reasons. The reasons include the market price is higher than your bid and the relationship between resource supply and demand changes. For more information, see Preemptible instance .

Server Load Balancer system events

Name	Meaning	Level
CertKeyExpired_1	The certificate will expire in 1 day.	WARN
CertKeyExpired_3	The certificate will expire in 3 days.	WARN
CertKeyExpired_7	The certificate will expire in 7 days.	WARN
CertKeyExpired_15	The certificate will expire in 15 days.	WARN
CertKeyExpired_30	The certificate will expire in 30 days.	WARN
CertKeyExpired_60	The certificate will expire in 60 days.	WARN

Object Storage Service system events

Name	Meaning	Level	Remarks
BucketEgressBandwidth	The downstream bandwidth that is used by buckets has exceeded the report threshold.	INFO	This event is triggered if the total downstream bandwidth that is used by all the buckets owned by a user exceeds the report threshold of 128 Mbit/s.
BucketEgressBandwidthThresholdExceeded	The downstream bandwidth that is used by a bucket has exceeded the throttling threshold.	WARN	This bucket is subject to regional throttling. Throttling is triggered if the total downstream bandwidth that is used by all the buckets in a region exceeds the throttling threshold. Alibaba Cloud throttling threshold is 10 Gbit/s for each region in mainland China, and 5 Gbit/s for Hong Kong and overseas regions by default. No bucket-level throttling threshold is configured by default.

Name	Meaning	Level	Remarks
BucketIngressBandwidth	The upstream bandwidth that is used by buckets has exceeded the report threshold.	INFO	This event is triggered if the total upstream bandwidth that is used by all the buckets owned by a user exceeds the report threshold of 128 Mbit/s.
BucketIngressBandwidthThresholdExceeded	The upstream bandwidth that is used by a bucket has exceeded the throttling threshold.	WARN	This bucket is subject to regional throttling. Throttling is triggered if the total upstream bandwidth that is used by of all the buckets in a region exceeds the throttling threshold. Alibaba Cloud throttling threshold is 10 Gbit/s for each region in Mainland China, and 5 Gbit/s for Hong Kong and overseas regions by default. No bucket-level throttling threshold is configured by default.
UserEgressBandwidth	The downstream bandwidth of a user has exceeded the report threshold.	INFO	This event is triggered if the total downstream bandwidth that is used by all the buckets owned by a user exceeds the report threshold of 128 Mbit/s.

Name	Meaning	Level	Remarks
UserEgressBandwidthThresholdExceeded	The downstream bandwidth of a user has exceeded the throttling threshold .	WARN	Throttling is triggered if the total downstream bandwidth that is used by all the buckets in a region exceeds the throttling threshold . Alibaba Cloud throttling threshold is 10 Gbit/s for each region in mainland China, and 5 Gbit/s for Hong Kong and overseas regions by default . No bucket-level throttling threshold is configured by default.
UserIngressBandwidth	The upstream bandwidth of a user has exceeded the report threshold.	INFO	This event is triggered if the total upstream bandwidth that is used by all the buckets owned by a user exceeds the report threshold of 128 Mbit/s.

Name	Meaning	Level	Remarks
UserIngressBandwidthThresholdExceeded	The upstream bandwidth of a user has exceeded the throttling threshold .	WARN	Throttling is triggered if the total upstream bandwidth that is used by all the buckets in a region exceeds the throttling threshold . Alibaba Cloud throttling threshold is 10 Gbit/s for each region in mainland China, and 5 Gbit/s for Hong Kong and overseas regions by default . No bucket-level throttling threshold is configured by default.

Auto Scaling system events

Name	Meaning	State	Level
AUTOSCALING:SCALE_IN_ERROR	The scale-in activity of the scaling group failed to be completed.	Unnormal	CRITICAL
AUTOSCALING:SCALE_IN_SUCCESS	The scale-in activity of the scaling group was successful.	Normal	INFO
AUTOSCALING:SCALE_OUT_ERROR	The scale-out activity of the scaling group failed to be completed.	Unnormal	CRITICAL
AUTOSCALING:SCALE_OUT_SUCCESS	The scale-out activity of the scaling group was successful.	Normal	INFO

Name	Meaning	State	Level
AUTOSCALING:SCALE_REJECT	The scaling activity of the scaling group was rejected.	Warn	WARN
AUTOSCALING:SCHEDULE_TASK_EXPIRING	Scheduled task expiration reminder.	Warn	WARN
AUTOSCALING:SCALE_OUT_START	The scale-out activity of the scaling group has started.	normal	INFO
AUTOSCALING:SCALE_IN_START	The scale-in activity of the scaling group has started.	normal	INFO

Alibaba Cloud IoT system events

Name	Meaning	State	Level
RuleEngineProcessFail	Rule Engine processing failed to be completed.	Failed	WARN

Smart Access Gateway system events

Name	Meaning	State	Level
AccessGatewayFailover	The access point fails over.	Agwfailover	INFO
ConnectionDisconnect	The network is disconnected.	Disconnect	CRITICAL
DeviceHacked	The device is under attack.	Hacked	CRITICAL
DeviceOffline	The device is offline.	Offline	CRITICAL
DeviceOnline	The device is online.	Online	INFO

CloudMonitor system events

Name	Meaning	State	Level
Group_AddResourcesFailed_QuotaReached	Machines failed to be dynamically added to a group because the resource usage limits have been reached.	Failed	CRITICAL
Agent_Status_Stopped	Heartbeat check failed.	Stopped	CRITICAL
Agent_Status_Running	Heartbeat check is resumed.	Running	CRITICAL

Database Backup system events

Name	Meaning	State	Level
CloseContBackup	Incremental backup is disabled.	Failed	INFO
ContBackupFail	An exception has occurred during incremental backup.	Failed	WARN
DataRestoreFail	An exception has occurred during data recovery.	Failed	WARN
DataRestoreSuccess	Data recovery is successful.	Running	WARN
FullBackupFail	An error has occurred during full backup.	Failed	WARN
InstancePause	The scheduled backup plan is suspended.	Failed	INFO
InstanceStart	The scheduled backup plan has started.	Running	INFO
OpenContBackup	Incremental log backup is enabled.	Running	INFO

Relational Database Service system events

Name	Meaning	State	Level
Instance_Failover	Instance failover	Executed	WARN
Instance_Failure_Start	Beginning of an instance failure	Executing	CRITICAL
Instance_Failure_End	End of an instance failure	Executed	CRITICAL

ApsaraDB RDS for Redis system events

Name	Meaning	State	Level
Instance_Failover	Instance failover	Executed	WARN
Instance_Failure_Start	Beginning of an instance failure	Executing	CRITICAL
Instance_Failure_End	End of an instance failure	Executed	CRITICAL

ApsaraDB RDS for MongoDB system events

Name	Meaning	State	Level
Instance_Failure_Start	Beginning of an instance failure	Executing	CRITICAL
Instance_Failure_End	End of an instance failure	Executed	CRITICAL

9.2.2 View cloud service events

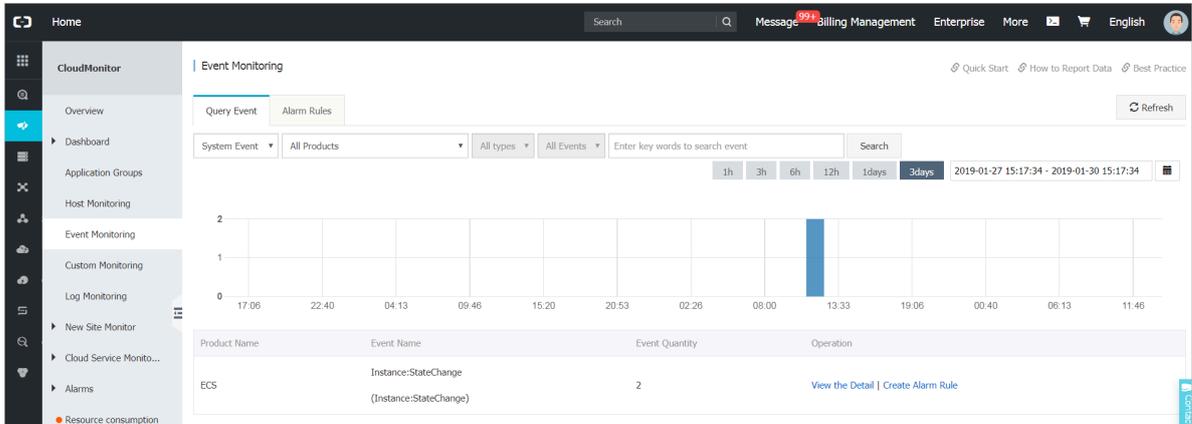
Event monitoring allows you to query and view statistics of all system events that are generated by various cloud services. You can obtain an overview of how those services are running.

After you use application groups to classify resources, service-related system events are automatically associated with the resources of different groups. This helps you integrate all kinds of monitoring information, and quickly analyze and locate problems.

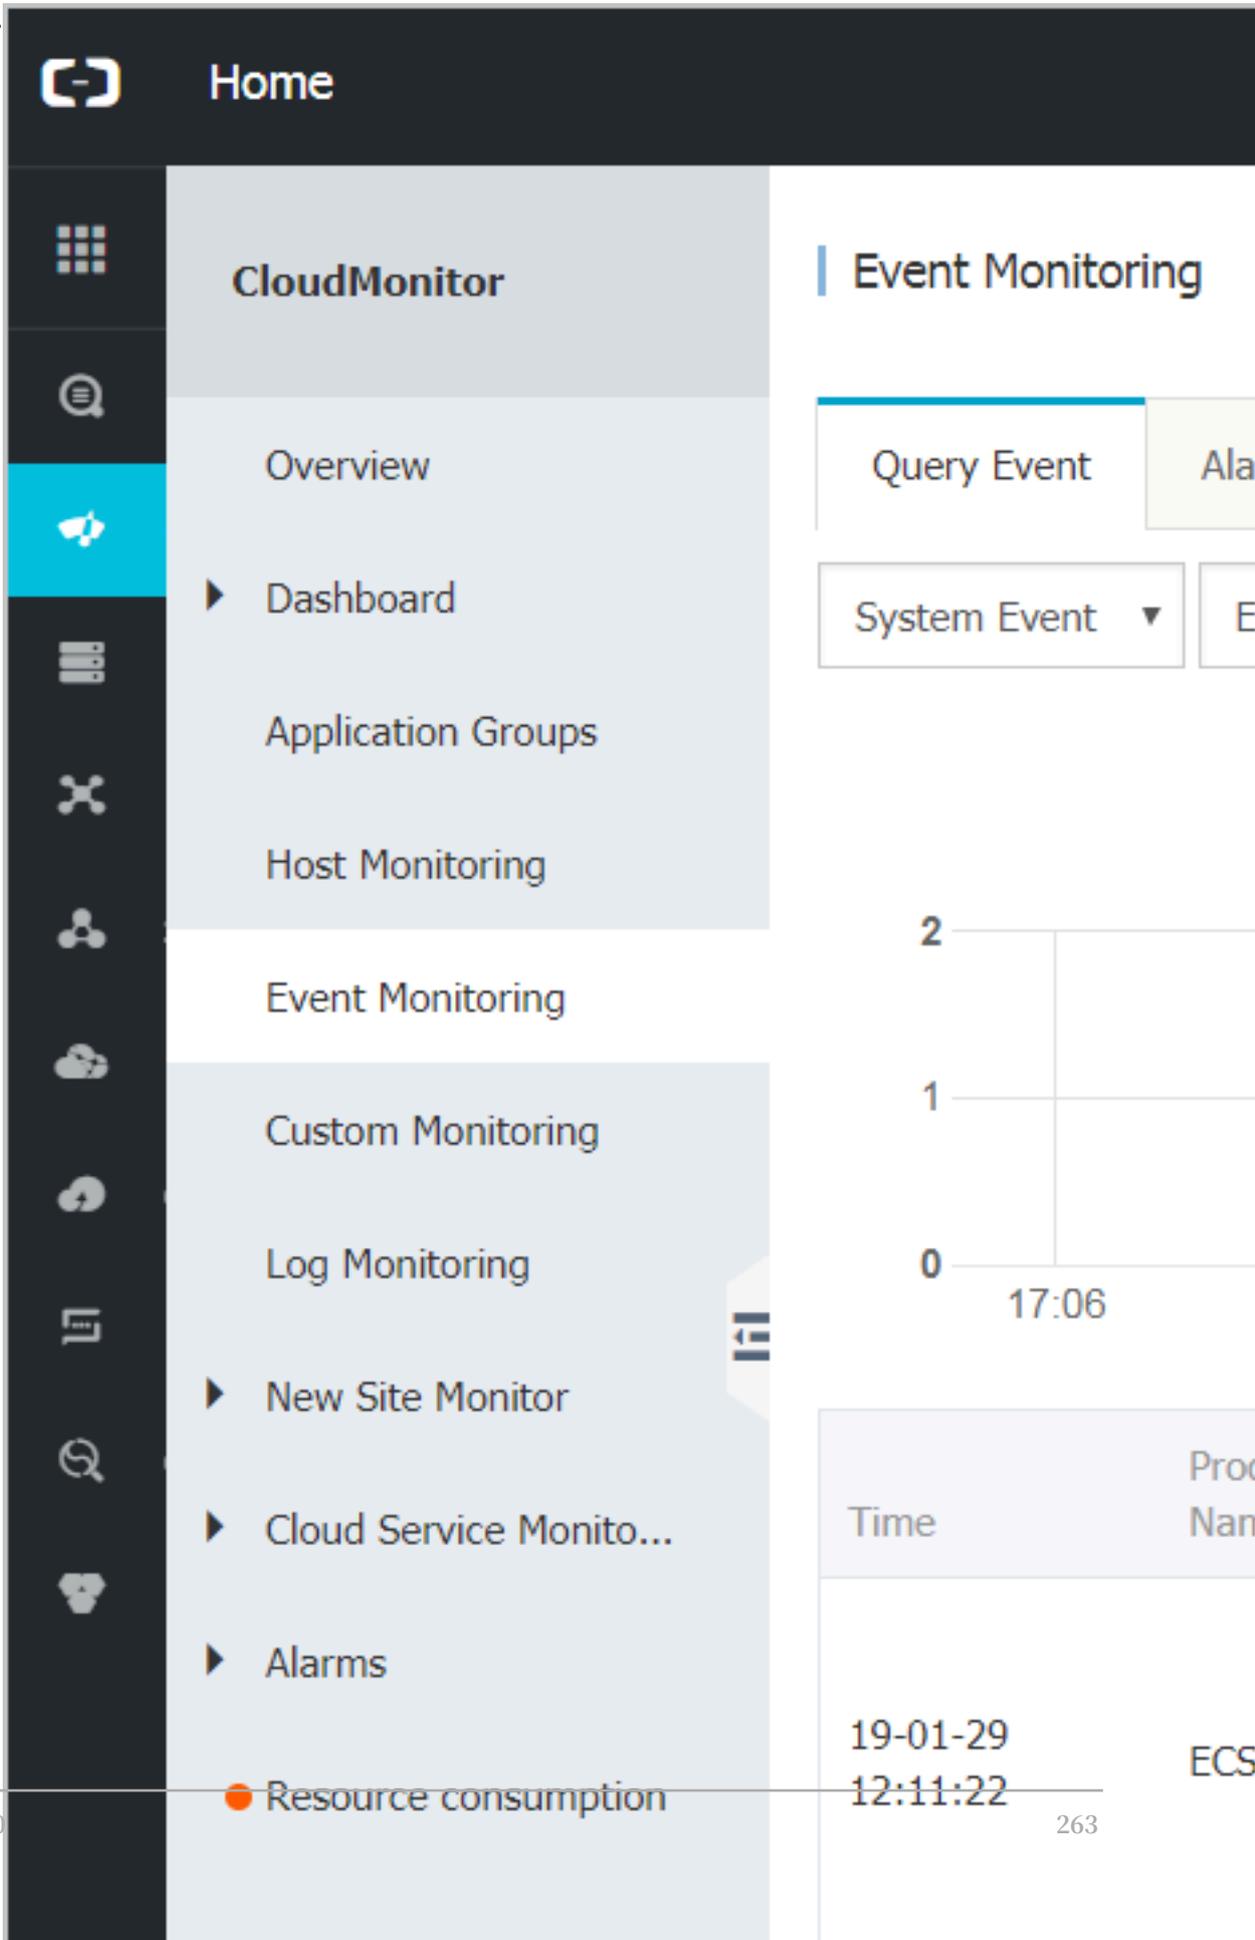
View system events by service

1. Log on to the [CloudMonitor console](#).

2. In the left-side navigation pane, click Event Monitoring. The Event Monitoring page is displayed. Set event type to System Events. Select a service from the service drop-down list and an event from the event drop-down list. Select a time range. The events that occurred within the specified time range are displayed.



3. Click View Details in the Actions column corresponding to an event to view the details of the event.



View system events by group

If you manage your instances by application group, you can also access a specific application group to view the system events related to the instances in the group.

1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups. The Application Groups page is displayed.
3. Click a group name to go to the group details page.
4. On the group details page, click Event Monitoring in the left-side navigation pane. The Event Monitoring page that appears displays the system events related to the instances in the group.

9.2.3 Use the event alert function for Alibaba Cloud services

This topic describes how to use the event alert function for Alibaba Cloud services to enable alerts when a system exception occurs.

Background information

When an Alibaba Cloud service encounters a system exception, event monitoring can provide two types of notification. You can trace the event and automate the handling process.

- Event alert notification can be sent to you through , emails, and DingTalk Chatbot.
- The event is distributed to your MNS queue, Function Compute, Log Service, and URL callback. Then, you can automate the handling process based on your scenario.

Preparations

If you want system events to be distributed to your MNS queue, Function Compute, Log Service, and URL callback, you must enable the corresponding services.

Procedure

You can create an event alert rule. Then, you can use system event testing to check whether the configured MNS queue of the alert rule can receive event alerts in a timely manner, and whether Function Compute can be triggered.

- Create an event alert rule
 1. Log on to the [CloudMonitor console](#).
 2. In the left-side navigation pane, click Event Monitoring.
 3. On the Alarm Rules tab, click Create Event Alarms. The Create/Modify Event Alarms dialog box is displayed.
 4. In the Basic Information section, set Alarm Rule Name.
 5. In the Event Alarm section, set the following parameters:
 - a. Event Type: Select System Event.
 - b. Product Type, Event Level, and Event Name: Set the parameters based on the actual situation.
 - c. Resource Range: If you select All Resources, alerts are sent when any resource-related events occur. If you select Application Groups, alerts are

sent only when events that are related to the resources in the specified group occur.

- 6. Select the Alarm Type. CloudMonitor supports four alert types: alert notification, MNS queue, function service, and URL callback.

Alarm type

Alarm notification

Contact Group

Default Contact Group

Notification Method

Warning (Message+Email ID+ Ali WangWang+DingTalk Robot)

+Add

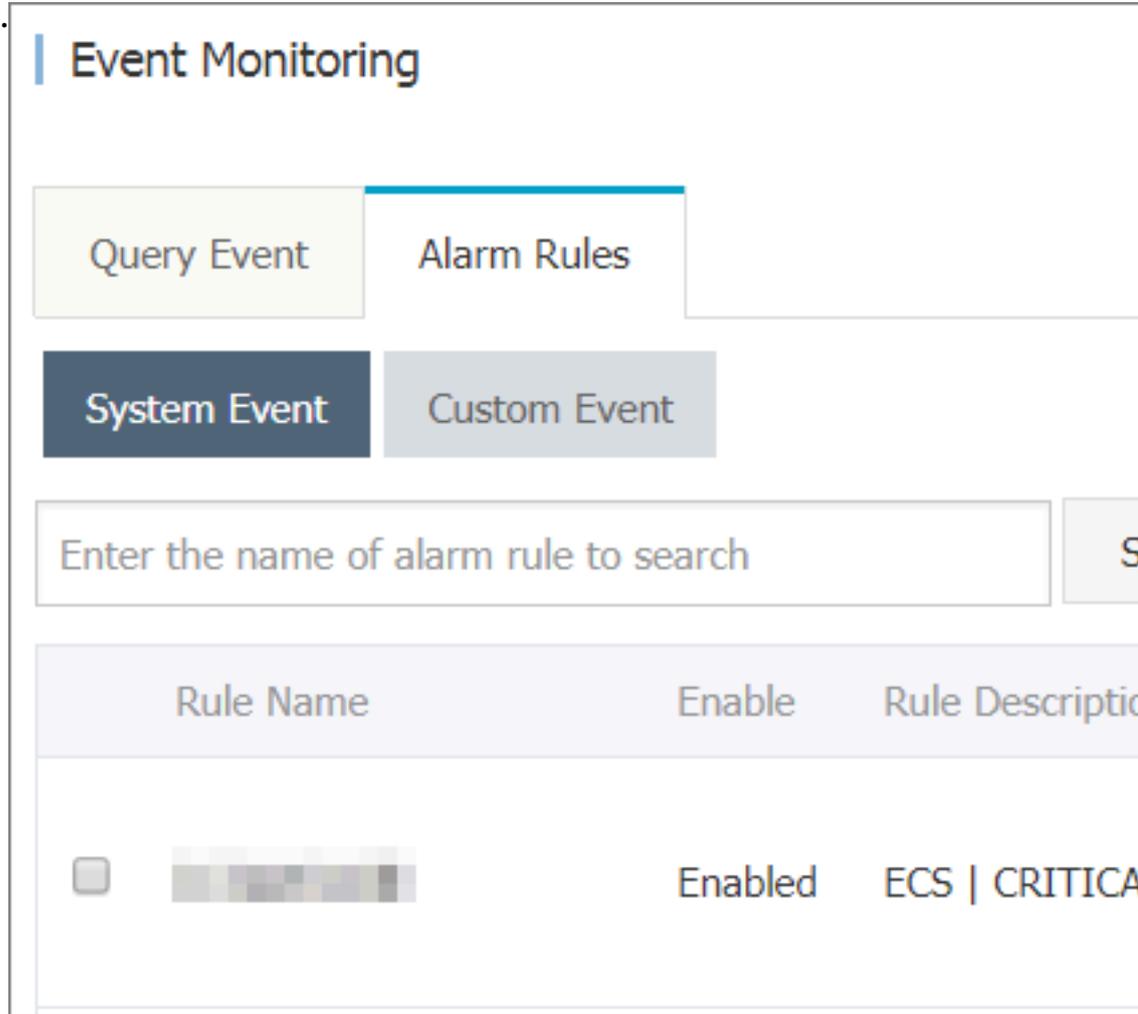
MNS queue

Function service

URL callback

- Test an alert rule

1. Access the Alarm Rules tab for event monitoring.



2. Click Test in the Actions column corresponding to the RAM user.
3. On the test page that appears, select the event to be tested. The corresponding event content will be displayed. You can change content fields such as the instance ID as needed.
4. Click OK. The system will send an event based on the content, triggering alert notification, MNS queue, Function

9.3 Custom events

9.3.1 Report custom event data

Event monitoring provides APIs for reporting custom events. You can use the APIs to collect event-related exceptions and report them to CloudMonitor. You can also configure alert rules to receive alert notification when an event-related exception occurs.

CloudMonitor provides three methods to report data: APIs, Java SDK, and Alibaba Cloud CLI.

Limits

- Each Alibaba Cloud account can send up to 20 report requests per second.
- Each report can contain up to 100 events.
- Each report can contain up to 500 KB of data.

Report data through APIs

- Endpoints

```
https://metrichub-cms-cn-hangzhou.aliyuncs.com
```

- Request syntax

```
POST /event/custom/upload HTTP/1.1
Authorization:<AuthorizationString>
Content-Length:<ContentLength>
Content-MD5:<ContentMD5>
Content-Type:application/json
Date:<GMTDate>
Host:metrichub-cms-cn-hangzhou.aliyuncs.com
x-cms-signature:hmac-sha1
x-cms-api-version:1.0
x-cms-ip:30.27.84.196
User-Agent:cms-java-sdk-v-1.0
[{"content":"EventContent","groupId":GroupId,"name":"
EventName","time":"20171023T14:44:39.948+0800"}]
```

- Request parameters

Name	Type	Required?	Description
name	String	Yes	The name of the event

Name	Type	Required?	Description
groupId	Numerical	Yes	The ID of the application group, to which the event belongs
time	String	Yes	The time when the event occurs
content	String	Yes	The event details

- Request header definition

The following table lists the request headers of event monitoring APIs.

Header	Type	Description
Authorization	String	The authorization string . Content: AccessKeyId: SignString
User-Agent	String	The client description.
Content-MD5	String	The uppercase string obtained after performing MD5 computation on the value of the Body field in the request. If the request does not have the Body field, this request header is not required.
Content-Length	Numerical	The Body field length in the HTTP request as defined in RFC 2616. If the request does not have the Body field, this request header is not required.
Content-Type	String	The Body field type in the HTTP request. Valid values: application and json.

Header	Type	Description
Date	String	The standard timestamp header of the HTTP request. This header follows the RFC 1123 format and uses GMT standard time, such as Mon, 3 Jan 2010 08:33:47 GMT.
Host	String	The full host name of the HTTP request. This header does not include the protocol header such as https://. For example , metrichub-cms-cn-hangzhou.aliyuncs.com.
x-cms-api-version	String	The API version. The current version is 1.0.
x-cms-signature	String	The signature algorithm . Currently, the only supported signature algorithm is HMAC-SHA1.
x-cms-ip	String	The IP address of the machine that reports the event, such as 10.1.1.1.

- Signature algorithm

Currently, the only supported signature algorithm is HMAC-SHA1.

1. Prepare an Alibaba Cloud AccessKey pair.

To generate a digital signature for an API request, you must use an AccessKey pair that is composed of an AccessKey ID and an AccessKey Secret. You can use an existing AccessKey pair or create a new one. The AccessKey pair must be in the Active state.

2. Generate a signature string for the request

An API signature string consists of the Method , Header , and Body fields of the HTTP request.

```
SignString = VERB + "\ n "
            + CONTENT - MD5 + "\ n "
```

```
+ CONTENT - TYPE + "\ n "
+ DATE + "\ n "
+ CanonicalizedHeaders + "\ n "
+ CanonicalizedResource
```

In the preceding formula, \ n indicates the newline escape character and the plus sign (+) indicates the string concatenation operator. The other parts are defined as follows:

Name	Definition	Examples
VERB	The method name of the HTTP request	PUT, GET, and POST
CONTENT-MD5	The MD5 value of the Body field in the HTTP request, which must be an uppercase string	875264590688CA6171F6228AF5BBB3D2
CONTENT-TYPE	The Body field type in the request	application/json
DATE	The standard timestamp header of the HTTP request, which follows the RFC 1123 format and uses GMT standard time	Mon, 3 Jan 2010 08:33:47 GMT
CanonicalizedHeaders	The string constructed by the custom headers prefixed with x-cms and x-acs in the HTTP request	x-cms-api-version:0.1.0\nx-cms-signature

Name	Definition	Examples
CanonicalizedResource	The string constructed by the HTTP request resources, as described in the following section	/event/custom/upload

CanonicalizedHeaders in the preceding table is constructed as follows:

- a. Convert the names of all HTTP request headers prefixed with `x - cms` and `x - acs` to lowercase letters.
- b. Sort the CMS custom request headers obtained in the preceding step in ascending lexicographical order.
- c. Delete any space on either side of a delimiter between the request header and content.
- d. Separate all headers and content with separators (`\ n`) to form the final CanonicalizedHeaders.

CanonicalizedResource in the preceding table is constructed as follows:

- a. Set CanonicalizedResource as an empty string ("").
- b. Place the URI that you want to access, such as `/ event / custom / upload`, between the quotation marks.
- c. If the request contains a query string (`QUERY_STRI NG`), add a question mark (?) and the query string to the end of the CanonicalizedResource string.
 `QUERY_STRI NG` is the lexicographic string of the request parameters included in the URL. Equal signs (=) are used between the names and values of parameters to form a string. The parameter name-value pairs are then

sorted in ascending lexicographical order and connected with ampersands (&) to form a string. The formula is as follows:

```
QUERY_STRING = " KEY1 = VALUE1 " + "&" + " KEY2 = VALUE2 "
```

3. Generate a digital signature for the request

Currently, the only supported signature algorithm is HMAC-SHA1. The following formula is used to generate a signature:

```
Signature = base64 ( hmac - sha1 ( UTF8 - Encoding - Of ( SignString ), AccessKeySecret ) )
```

- **Response elements**

The system returns the HTTP status code 200.

- **Examples**

- **Sample request**

```
POST / event / custom / upload HTTP / 1 . 1
Host : metrichub - cms - cn - hangzhou . aliyuncs . com
x - cms - api - version : 1 . 0
Authorization : YourAccKey : YourAccSecret
Host : metrichub - cms - cn - hangzhou . aliyuncs . com
Date : Mon , 23 Oct 2017 06 : 51 : 11 GMT
Content - Length : 180
x - cms - signature : hmac - sha1
Content - MD5 : E9EF574D1A EAAA370860 FE37856995 CD
x - cms - ip : 30 . 27 . 84 . 196
User - Agent : cms - java - sdk - v - 1 . 0
Content - Type : application / json
[{" content ":" 123 , abc "," groupId ":" 100 ," name ":" Event_0
"," time ":" 20171023T1 44439 . 948 + 0800 "}]
```

- **Sample response**

```
{
  " code ":" 200 ",
  " msg ":" // The returned msg is null when the
reporting is normal .
}
```

Report data through Java SDK

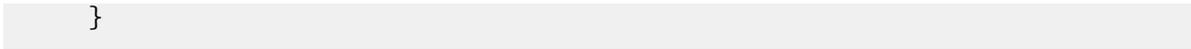
- **Maven dependency**

```
< dependency >
  < groupId > com . aliyun . openservices </ groupId >
  < artifactId > aliyun - cms </ artifactId >
  < version > 0 . 1 . 2 </ version >
```

```
</ dependency >
```

- **Sample code**

```
public void uploadEvent () throws CMSException ,
InterruptedException {
    // Initialize the client .
    CMSClient cmsClient = new CMSClient ( endpoint ,
accKey , secret );
    // Construct two events to be reported .
    CustomEventUploadRequest request = CustomEvent
tUploadRequest . builder ()
        . append ( CustomEvent . builder ()
            . setContent ( " abc , 123 " )
            . setGroupId ( 101l )
            . setName ( " Event001 " ). build () )
        . append ( CustomEvent . builder ()
            . setContent ( " abc , 123 " )
            . setGroupId ( 101l )
            . setName ( " Event002 " ). build () )
        . build ();
    CustomEventUploadResponse response = cmsClient
. putCustomEvent ( request );
    List < CustomEvent > eventList = new ArrayList <
CustomEvent > ();
    eventList . add ( CustomEvent . builder ()
        . setContent ( " abcd , 1234 " )
        . setGroupId ( 101l )
        . setName ( " Event001 " ). build ());
    eventList . add ( CustomEvent . builder ()
        . setContent ( " abcd , 1234 " )
        . setGroupId ( 101l )
        . setName ( " Event002 " ). build ());
    request = CustomEventUploadRequest . builder ()
        . setEventList ( eventList ). build ();
    response = cmsClient . putCustomEvent ( request );
```

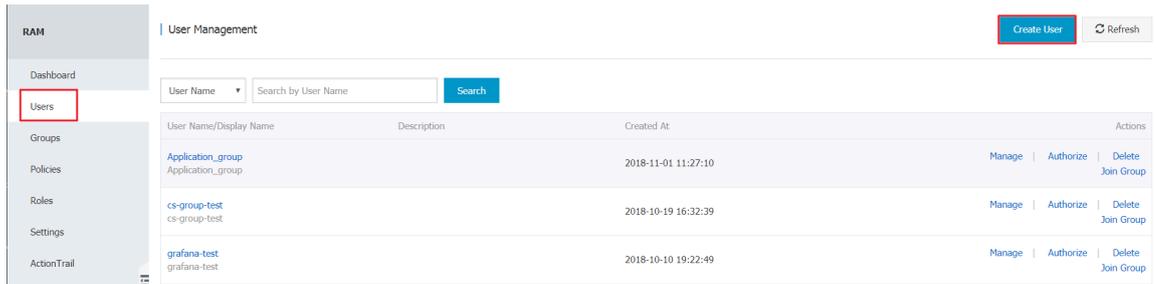


Report data through Alibaba Cloud CLI

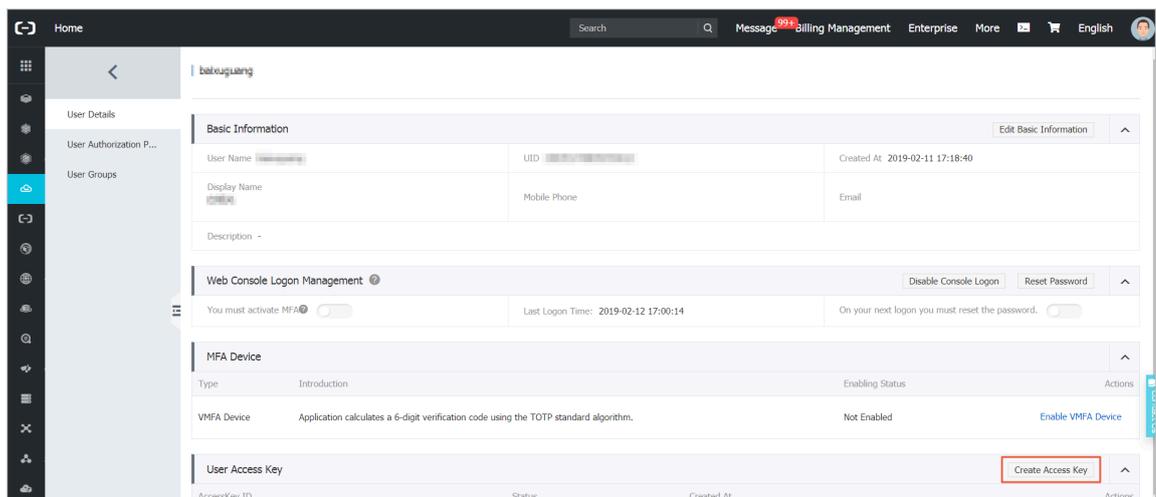
1. Prerequisites

Ensure that you have created an Alibaba Cloud account, created a RAM user for your account, and generated a RAM user AK with CloudMonitor permissions.

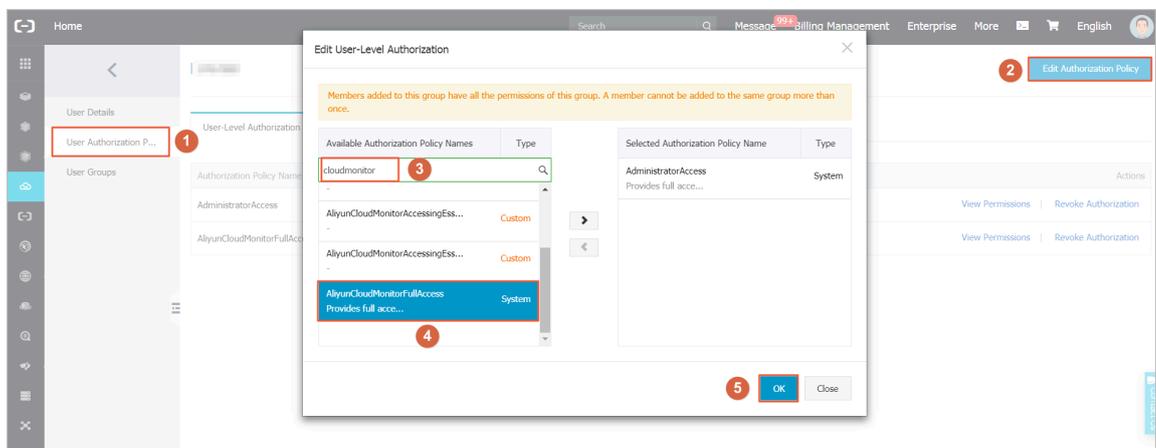
a. Create a RAM user.



b. Generate an AccessKey ID and an AccessKey Secret for the RAM user.



c. Grant CloudMonitor permissions to the RAM user.



2. Install CMS SDK

- The installation method for a Windows system is as follows:

```
cd C:\Python27\Scripts
pip install aliyun-python-sdk-cms
```

Run the following command to update the SDK:

```
pip install --upgrade aliyun-python-sdk-cms
```

- The installation method for a Linux system is as follows:

```
sudo pip install aliyun-python-sdk-cms
```

Run the following command to update the SDK:

```
sudo pip install --upgrade aliyun-python-sdk-cms
```

3. Report monitoring data

Use the `PutEvent` API to report the monitoring data.

- Example for a Windows system:

```
aliyuncli .exe cms PutEvent -- EventInfo "[{' content ':'
helloworld ',' time ':' 20171013T1 70923 . 456 + 0800 ',' name
':' ErrorEvent ',' groupId ':' 27147 '}]"
```

- Example for a Linux system:

```
aliyuncli cms PutEvent -- EventInfo "[{' content ':'
helloworld ',' time ':' 20171023T1 80923 . 456 + 0800 ',' name
':' ErrorEvent ',' groupId ':' 27147 '}]"
```

- If the data is reported successfully, status code 200 is returned.

```
{
  " Code ":" 200 "
}
```

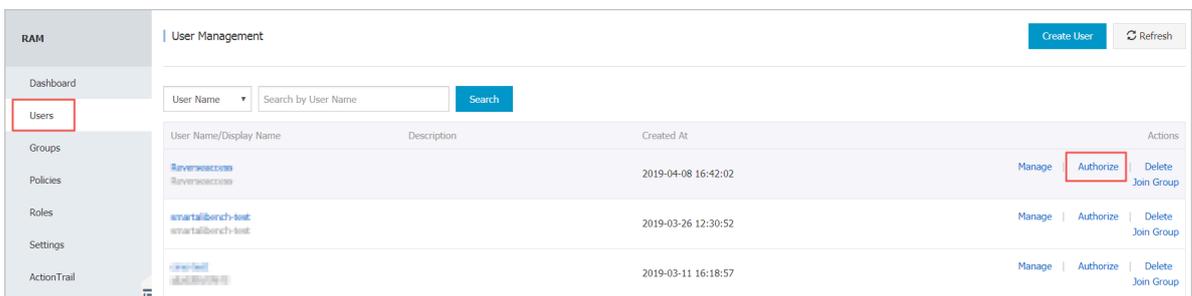
Error codes

Error code	Description
200	Normal
400	Syntax error in the client request
403	Verification failure, speed limit, or not authorized
500	Internal server error

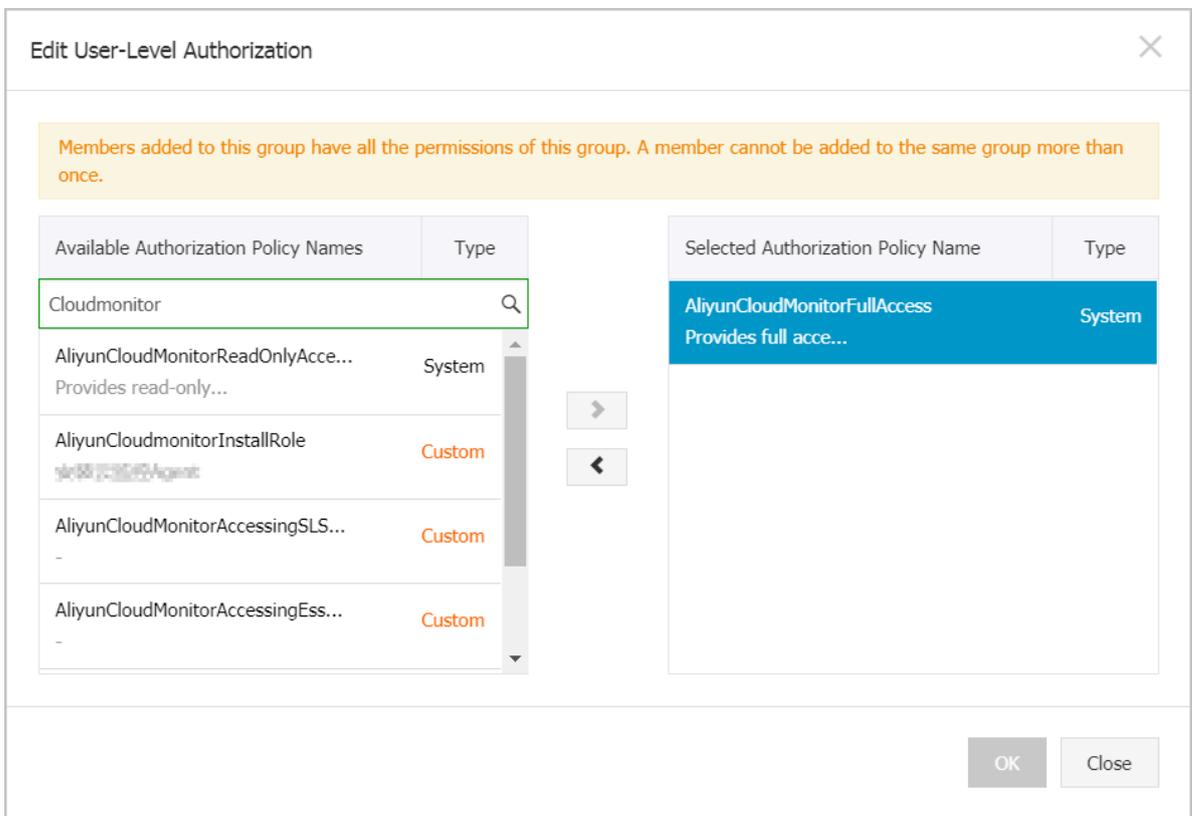
RAM user authorization

You must grant CloudMonitor permissions to the corresponding RAM user before event data can be reported with the RAM user AK. If you do not grant the permissions, when you report data, the prompt "cannot upload event, please use ram to auth" is displayed.

1. Log on to the [RAM console](#).
2. In the left-side navigation pane, click Users.
3. On the Users page that appears, click Authorize in the Actions column corresponding to the RAM user.



4. On the authorization page, select AliyunCloudMonitorFullAccess and click OK.



9.3.2 View custom events

Event monitoring allows you to query data and view statistics related to custom events.

View custom events by event type

1. Log on to the [CloudMonitor Console](#).
2. Choose Event Monitoring > Query Event. Select Custom Event from the first drop-down list. Next, select the target event type from the second one and the specific event from the third one. Then, specify the time period.
3. In the Operation column, click View the Detail.

View custom events by application group

If you manage your instances by using an application group, you can view the custom events of an instance by directly accessing the application group page.

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, click Application Groups.
3. On the Application Groups page, click the name of the target group.
4. On the displayed page, click Event Monitor in the left-side navigation pane. On the displayed page, choose Custom Event from the first drop-down list.

9.3.3 Use the custom event alarm function

This topic describes how to use the custom event alarm function.

Overview

To notify you of data exceptions, the custom event alarm function provides the following two notification methods:

- Notifications sent as e-mails or DingTalk messages
- Notifications sent to your alarm callback URL for scenario-oriented troubleshooting

Procedure

1. Log on to the [CloudMonitor Console](#).
2. Choose Event Monitoring > Alarm Rules.

3. Click Create Event Alerts.

The following figure shows the displayed Create / Modify Event Alerts dialog box.

Create / Modify Event Alerts ✕

Basic Information

Alarm Rule Name

Combination of alphabets, numbers and underscore, in 30 characters

Event alert

Event Type

System Event Custom Event

Application Groups

2149326 / k8s-c61a139b41e144d22a1124ba8159f2f73-worker

Event Name

Enter the name of the reported event

Rule Description

1minutes accumulatively happened for times

Notification Method

Email + DingTalk ?

Email + DingTalk

Email + DingTalk

[Advanced Configuration](#)

4. In the Basic Information area, enter a name for the alarm rule.

5. In the Event alert area, configure the following settings:
 - a. Set Event Type to Custom Event.
 - b. Set Application Groups to the target application group.
 - c. Enter a Event Name.
 - d. Select an option from the Rule Description drop-down list and set the accumulation times.
 - e. Choose your preferred Notification Method.
 - f. In the Advanced Configuration area, set Effective From and Alarm Callback.
 - **Effective From:** Indicates the time from which the alarm rule begins to take effect. The alarm rule checks whether to report alarms for monitoring data exceptions only during the period of time that you specified.
 - **Alarm Callback:** Enter a URL that can be accessed from the Internet. CloudMonitor will then send alarm notifications to the URL using an HTTP POST request.
 - g. Click OK.

When the reported custom event meets the conditions specified by the alarm rule, a notification is sent.

9.3.4 Event monitoring best practices

Use cases

Exceptions may occur when the service is running. Some exceptions can be automatically restored by retry and other methods, while the others cannot. Serious exceptions can even lead to customer business interruption. Therefore, a system is necessary to record these exceptions and trigger alarms when specific conditions are met. The traditional method is to print file logs and collect the logs to specific systems, for example, open-source ELK (ElasticSearch, Logstash, and Kibana). These open-source systems consist of multiple complex distributed systems. The complicated technology and high cost make independent maintenance challenging. CloudMonitor provides the event monitoring feature to effectively solve these problems.

The following examples explain how to use the event monitoring feature.

Case studies

1. Report exceptions

Event monitoring provides two methods for data reporting, namely, Java SDK and Open API. The following describes how to report data by using Java SDK.

a. Add Maven dependency

```
< dependency >
  < groupId > com . aliyun . openservic es </ groupId >
  < artifactId > aliyun - cms </ artifactId >
  < version > 0 . 1 . 2 </ version >
</ dependency >
```

b. Initialize SDK

```
// Here , 118 is the applicatio n grouping ID
of CloudMonit or . Events can be categorize d
by applicatio ns . You can view group IDs in
CloudMonit or applicatio n grouping list .
CMSClientI nit . groupId = 118L ;
// The address is the reporting entry of the
event system , which is currently the public network
address . AccessKey and Secret / key are used for
personal identity verificati on .
CMSClient c = new CMSClient (" https :// metrhub - cms -
cn - hangzhou . aliyuncs . com ", accesskey , secretkey );
```

c. Determine whether to asynchronously report the data.

CloudMonitor event monitoring provides synchronous reporting policy by default. The good thing is that writing code is simple, and the reported events are reliable and free from data loss.

However, such policy also brings some problems as well. Event reporting codes are embedded in business codes, which may block code running and affect the normal business in case of network fluctuations. Many business scenarios do not require events to be 100% reliable, so a simple asynchronous reporting encapsulation is sufficient. Write the event into a `LinkedBlockingQueue` and perform batch reporting on the backend asynchronously using `ScheduledExecutorService`.

```
// Initialize queue and Executors :
private LinkedBloc kingQueue < EventEntry > eventQueue =
new LinkedBloc kingQueue < EventEntry >( 10000 );
private ScheduledE xecutorSer vice schedule = Executors .
newSingleT hreadSched uledExecut or ();
// Report event :
// Every event contains its name and content . The
name is for identifca tion and the content
contains details of the event , in which the full
- text search is supported .
```

```

public void put ( String name , String content ) {
    EventEntry event = new EventEntry ( name , content );
    // When the event queue is full , additional
    events are discarded directly . You can adjust this
    policy as needed .
    boolean b = eventQueue . offer ( event );
    if ( ! b ) {
        logger . warn ( " The event queue is full ,
discard : {}" , event );
    }

// Submit events asynchronously . Initialize scheduled
tasks . Report events in batch by run every second
. You can adjust the reporting interval as needed
.
schedule . scheduleAt FixedRate ( this , 1 , 1 , TimeUnit .
SECONDS );
public void run () {
    do {
        batchPut ();
    } while ( this . eventQueue . size () > 500 );

private void batchPut () {
    // Extract 99 events from the queue for batch
reporting .
    List < CustomEvent > events = new ArrayList <
CustomEvent > ();
    for ( int i = 0 ; i < 99 ; i ++ ) {
        EventEntry e = this . eventQueue . poll ();
        if ( e == null ) {
            break ;

            events . add ( CustomEvent . builder (). setContent ( e
. getContent () ). setName ( e . getName () ). build () );

        if ( events . isEmpty () ) {
            return ;

            // Report events in batch to CloudMonitor . No
retry or retry in SDK is added here . If you
have high requirement for event reliability , add
retry policies .
            try {
                CustomEventUploadRequestBuilder builder =
CustomEventUploadRequest . builder ();
                builder . setEventList ( events );
                CustomEventUploadResponse response = cmsClient .
putCustomEvent ( builder . build () );
                if ( !" 200 " . equals ( response . getErrorCo de () ) ) {
                    logger . warn ( " event reporting error : msg
: {} , rid : {}" , response . getErrorMessage () , response .
getRequest Id () );

                } catch ( Exception e1 ) {
                    logger . error ( " event reporting exception " , e1
);

```

d. Event reporting demo

- Demo1: http Controller exception monitoring

The main purpose is to monitor if a large number of exceptions exist in HTTP requests. If the number of exceptions per minute exceeds a certain limit, an alarm is triggered. The implementation principle is to intercept HTTP requests by using Spring interceptor, servlet filter and other technologies. Logs are created in case of exceptions and alarms are triggered by setting alarm rules.

The event reporting demo is as follows:

```
// Each event should be informative for
// searching and locating. Here, map is used for
// organizing events and converted to json format
// as event content.
Map<String, String> eventContent = new HashMap<
String, String>();
eventContent.put("method", "GET"); // http request
// method
eventContent.put("path", "/users"); // http path
eventContent.put("exception", e.getClass().getName
()); // Exception class name for searching
eventContent.put("error", e.getMessage()); //
// Error message of exception
eventContent.put("stack_trace", ExceptionUtils
.getStackTrace(e)); // Exception stack for
// locating
// Finally submit the events in the preceding
// asynchronous reporting method. Since no retry
// is performed in asynchronous reporting, event
// loss of small probability may happen. However
// it is sufficient for alarms of unknown http
// exceptions.
put("http_error", JsonUtils.toJson(eventContent));
image.png](http://ata2-img.cn-hangzhou.img-pub
.aliyun-inc.com/864cf095977cf61bd340dd1461a0247c
.png)
```

- Demo2: Monitoring of scheduled tasks on the backend and message consumption

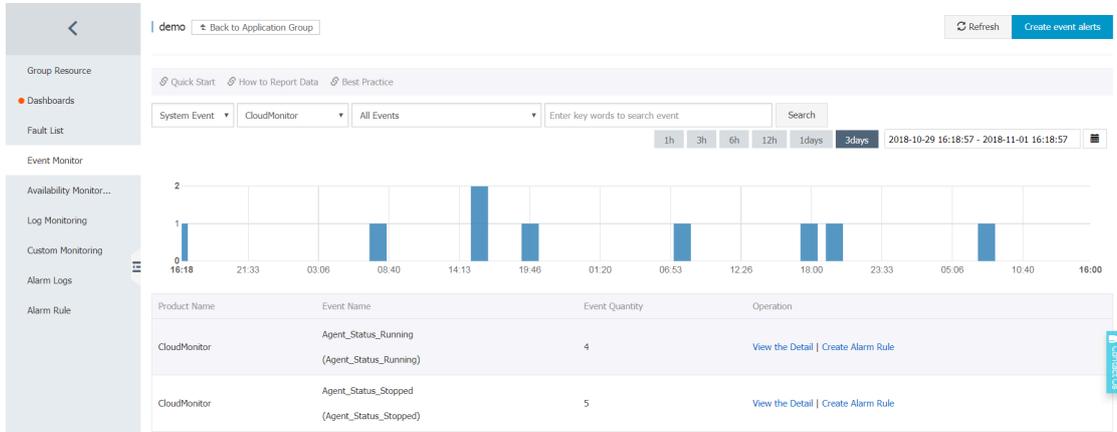
Like the preceding http events, many similar business scenarios require alarms. In the business scenarios such as backend tasks and message queue consumption, the events can be reported by using similar methods to achieve effective monitoring. When any exception occurs, alarms are triggered immediately.

```
// Event organization of the message queue :
```

```
Map < String , String > eventContent = new HashMap <
String , String >();
eventContent . put ( " cid " , consumerId ); // Consumer
ID
eventContent . put ( " mid " , msg . getMsgId ()); //
Message ID
eventContent . put ( " topic " , msg . getTopic ()); //
Message topic
eventContent . put ( " body " , body ); // Message body
eventContent . put ( " reconsume_ times " , String . valueOf
( msg . getReconsumeTimes ()); // The number of
retries after message failure
eventContent . put ( " exception " , e . getClass (). getName
()); // Exception class name in case of exception
eventContent . put ( " error " , e . getMessage ()); //
Exception message
eventContent . put ( " stack_trace " , ExceptionUtils .
getStackTrace ( e )); // Exception stack
// Finally , report the event
```

```
put (" metaq_error ", JsonUtils . toJson ( eventContent ));
```

Check the event after reporting:



- **Set alarms for queue message consumption exceptions:**

- **Demo 3: Record important events**

Another use case of events is to record important actions for later check without sending alarms. For example, operation logs for important business, password change/order change, remote logon, and so on.

demo [Back to Application Group](#) Refresh Create event alerts

[Quick Start](#) [How to Report Data](#) [Best Practice](#)

System Event Enter key words to search event

1h 3h 6h 12h 1days 3days 2018-10-29 16:18:57 - 2018-11-01 16:18:57

Time	Product Name	Event Name	Event Level	Status	Region	Resource	Contents	Close Detail
18-11-01 09:28:38	CloudMonitor	Agent_Status_Running (Agent_Status_Running)	CRITICAL	running	Unknown	acs:ecs:unknown:1270676679546704:instance/host-KQrpxFFRfS	{"ipGroup":"30.25.88.45","tianjimomVersion":"1.3.4"}	
18-10-31 09:32:34	CloudMonitor	Agent_Status_Running (Agent_Status_Running)	CRITICAL	running	Unknown	acs:ecs:unknown:1270676679546704:instance/host-KQrpxFFRfS	{"ipGroup":"30.25.88.37","tianjimomVersion":"1.3.4"}	
18-10-30 17:12:22	CloudMonitor	Agent_Status_Running (Agent_Status_Running)	CRITICAL	running	Unknown	acs:ecs:unknown:1270676679546704:instance/host-68E4vWgrSIY	{"ipGroup":"30.25.88.24","tianjimomVersion":"2.1.48"}	

10 Custom monitoring

10.1 Custom monitoring overview

Application scenarios

Custom monitoring allows you to customize metrics and alarm rules so that you can monitor metrics, report monitoring data, and set alarm rules with your specific requirements in mind.

Custom monitoring is different from event monitoring in that custom monitoring reports and queries time-series data that is collected periodically, whereas event monitoring only reports and queries data that is related to a singular event.

This topic discusses the procedures for operations custom monitoring including reporting, querying, and viewing monitoring data on the console, and how to set alarm rules for custom monitoring.

Procedures

- Report monitoring data.

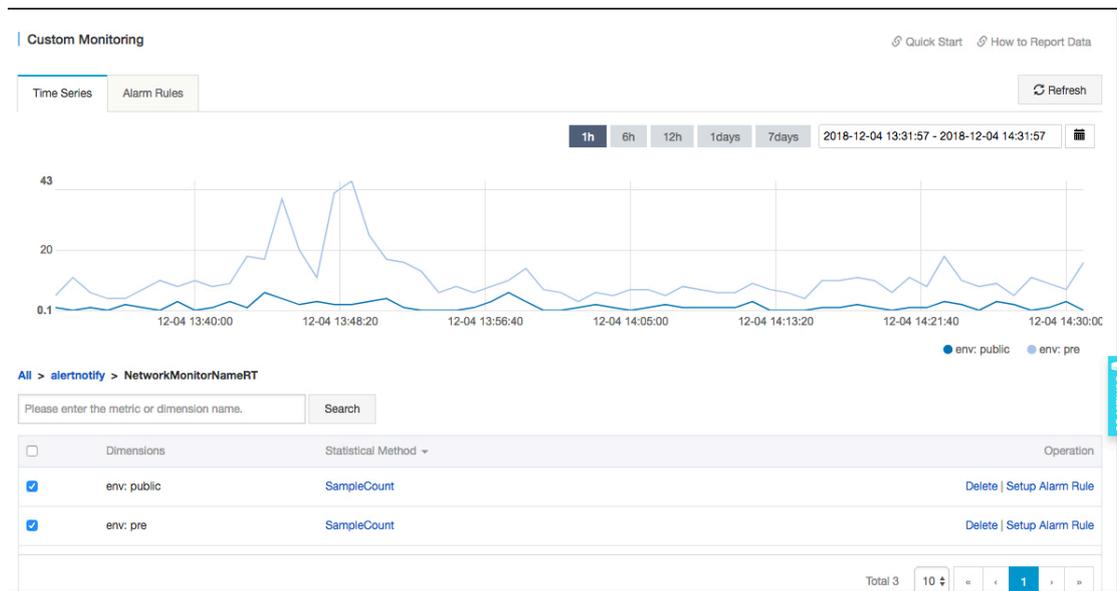
For more information and the specific procedure used, see [Report monitoring data](#).

- Query monitoring data.

After you have reported monitoring data, you can view the reported data in the console. You can choose to view all monitoring data on the custom monitoring page or to view custom monitoring data for one or more application group.

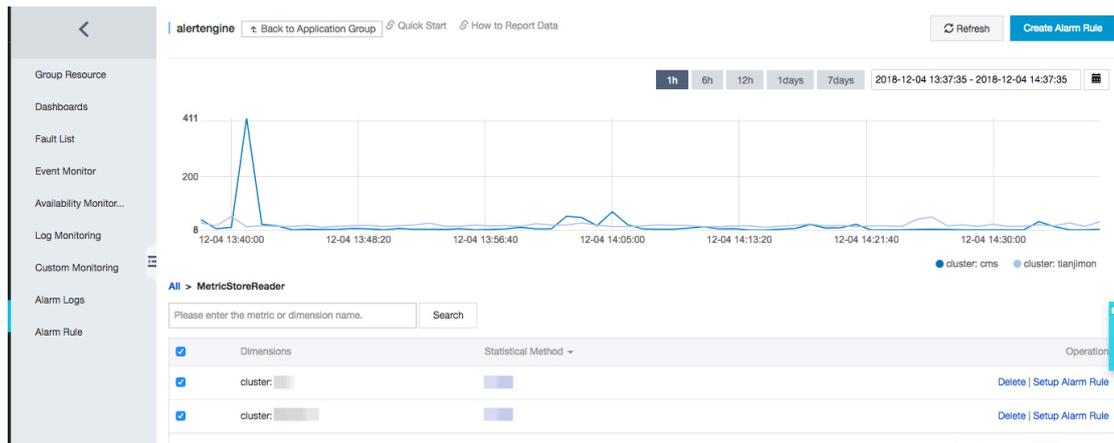
- To view all custom monitoring data, complete the following steps:

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, click Custom Monitoring. The Custom Monitoring page is displayed.
3. Select the corresponding application group and metric to access the Time Series page.
4. Select the time series you want to view.



- To view the custom monitoring data in an application group, complete the following steps:

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, click Application Groups. The Application Groups page is displayed.
3. Select the target application group.
4. Click Custom Monitoring. The Custom Monitoring page is displayed.
5. Select the target metric. The Time Series page is displayed.
6. Select the time series you want to view.



- Set an alarm rule.

Custom monitoring provides an alarm reporting feature. To set an alarm rule, you need to select an application group. When an alarm is triggered, a notification will be sent to the alarm contacts in the application group. To generate alarms for your monitoring data, set the alarm rule using either of the following two methods:

- Method 1:

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, click Custom Monitoring. The Custom Monitoring page is displayed.
3. Select the corresponding application group and metric. The Time Series page is displayed.
4. Select the time series for which you want to create an alarm rule, and then click Setup Alarm Rule in the Operation column.
5. On the Setup Alarm Rule page, enter a name for the alarm rule and set the corresponding alarm policy and notification method.

- Method 2:

1. Log on to the [CloudMonitor Console](#).
2. In the left-side navigation pane, click Application Groups. The Application Groups page is displayed.
3. Select the target application group. The Custom Monitoring page is displayed. Select the time series for which you want to create an alarm rule, and then click Setup Alarm Rule in the Operation column.
4. On the Setup Alarm Rule page, enter a name for the alarm rule and select the corresponding metric, dimension, alarm rule, and notification method.

10.2 Report monitoring data

This topic describes how to report custom monitoring data.

Custom monitoring allows you to customize metrics and alert rules to meet your own business needs. Custom monitoring provides APIs for reporting monitoring data. You can use the APIs to report collected time series data to CloudMonitor. You can also configure alert rules to receive alert notification when an exception occurs.

CloudMonitor provides three methods to report data: APIs, Java SDK, and Alibaba Cloud command line interface (CLI).

Limits

- The number of queries per second (QPS) is limited to 200 in China (Beijing), China (Shanghai), and China (Hangzhou), 100 in China (Zhangjiakou-Beijing Winter Olympics) and China (Shenzhen), and 50 in the other regions.
- A maximum of 100 data entries can be reported at a time. The body size cannot exceed 256 KB.
- The `metricName` field can contain letters, digits, and underscores (`_`). This field must start with a letter. If the starting character is not a letter it is replaced with an uppercase A. Invalid characters are replaced with underscores (`_`).
- The `dimensions` field cannot contain equal signs (`=`), ampersands (`&`), or commas (`,`). Invalid characters are replaced with underscores (`_`).
- The string length of the key or value of both `metricName` and `dimensions` parameters cannot exceed 64 bytes. Otherwise, the key or value string is truncated.
- Reporting raw data is a paid function. The free edition allows you to aggregate data. When you report data, you must pass "1" for the Type field to the request parameter.

Report data through APIs

After you report the raw data through APIs, CloudMonitor uses the following statistical methods to calculate the statistics at 1-minute and 5-minute intervals:

- **Average:** the average value
- **Maximum:** the maximum value
- **Minimum:** the minimum value
- **Sum:** the sum value

- **SampleCount:** the count
- **SumPerSecond:** the sum divided by the total number of seconds of the corresponding aggregation period. You can also use the moving average calculation.
- **CountPerSecond:** the count divided by the total number of seconds of the corresponding aggregation period. You can also use the moving average calculation.
- **LastValue:** the last sampled value in the aggregation period. It is similar to gauge.
- **P10:** The value of the 10th percentile. It is greater than 10% of all data in the aggregation period.
- **P20:** The value of the 20th percentile. It is greater than 20% of all data in the aggregation period.
- **P30:** The value of the 30th percentile. It is greater than 30% of all data in the aggregation period.
- **P40:** The value of the 40th percentile. It is greater than 40% of all data in the aggregation period.
- **P50:** The value of the 50th percentile. It is the median and greater than 50% of all data in the aggregation period.
- **P60:** The value of the 60th percentile. It is greater than 60% of all data in the aggregation period.
- **P70:** The value of the 70th percentile. It is greater than 70% of all data in the aggregation period.
- **P75:** The value of the 75th percentile. It is greater than 75% of all data in the aggregation period.
- **P80:** The value of the 80th percentile. It is greater than 80% of all data in the aggregation period.
- **P90:** The value of the 90th percentile. It is greater than 90% of all data in the aggregation period.
- **P95:** The value of the 95th percentile. It is greater than 95% of all data in the aggregation period.
- **P98:** The value of the 98th percentile. It is greater than 98% of all data in the aggregation period.
- **P99:** The value of the 99th percentile. It is greater than 99% of all data in the aggregation period.

- Endpoints

Internet endpoint: `https://metrichub-cms-cn-hangzhou.aliyuncs.com`

The following table lists the intranet endpoints.

Region	Region ID	Endpoint
China (Hangzhou)	cn-hangzhou	http://metrichub-cn-hangzhou.aliyun.com
China (Zhangjiakou-Beijing Winter Olympics)	cn-zhangjiakou	http://metrichub-cn-zhangjiakou.aliyun.com
China (Shanghai)	cn-shanghai	http://metrichub-cn-shanghai.aliyun.com
China (Beijing)	cn-beijing	http://metrichub-cn-beijing.aliyun.com
China (Qingdao)	cn-qingdao	http://metrichub-cn-qingdao.aliyun.com
China (Shenzhen)	cn-shenzhen	http://metrichub-cn-shenzhen.aliyun.com
Hong Kong	cn-hongkong	http://metrichub-cn-hongkong.aliyun.com
China (Hohhot)	cn-huhehaote	http://metrichub-cn-huhehaote.aliyun.com
UAE (Dubai)	me-east-1	http://metrichub-me-east-1.aliyun.com
US (Silicon Valley)	us-west-1	http://metrichub-us-west-1.aliyun.com
US (Virginia)	us-east-1	http://metrichub-us-east-1.aliyun.com
Japan (Tokyo)	ap-northeast-1	http://metrichub-ap-northeast-1.aliyun.com
Germany (Frankfurt)	eu-central-1	http://metrichub-eu-central-1.aliyun.com
Australia (Sydney)	ap-southeast-2	http://metrichub-ap-southeast-2.aliyun.com
Singapore	ap-southeast-1	http://metrichub-ap-southeast-1.aliyun.com

Region	Region ID	Endpoint
Malaysia (Kuala Lumpur)	ap-southeast-3	http://metrichub-ap-southeast-3.aliyun.com
India (Mumbai)	ap-south-1	http://metrichub-ap-south-1.aliyuncs.com

- Request syntax

```
POST / metric / custom / upload HTTP / 1 . 1
Authorizat ion :< Authorizat ionString >
Content - Length :< Content Length >
Content - MD5 :< Content MD5 >
Content - Type applicatio n / json
Date :< GMT Date >
Host : metrichub - cms - cn - hangzhou . aliyuncs . com
x - cms - signature : hmac - sha1
x - cms - api - version : 1 . 0
x - cms - ip : 30 . 27 . 84 . 196
User - Agent : cms - java - sdk - v - 1 . 0
[{" groupId " : 101 , " metricName " : "" , " dimensions " : {" sampleName
1 " : " value1 " , " sampleName 2 " : " value2 " } , " time " : "" , " type " :
0 , " period " : 60 , " values " : {" value " : 10 . 5 , " Sum " : 100 } } ]
```

- Signature algorithm

Currently, the only supported signature algorithm is HMAC-SHA1.

1. Prepare an Alibaba Cloud AccessKey pair.

To generate a digital signature for an API request, you must use an AccessKey pair that is composed of an AccessKey ID and an AccessKey Secret. You can use an existing AccessKey pair or create a new one. The AccessKey pair must be in the Active state.

2. Generate a signature string for the request

An API signature string consists of the Method , Header , and Body fields of an HTTP request.

```
SignString = VERB + "\ n "
            + CONTENT - MD5 + "\ n "
            + CONTENT - TYPE + "\ n "
            + DATE + "\ n "
            + Canonicali zedHeaders + "\ n "
```

+ CanonicalizedResource

In the preceding formula, \n indicates the newline escape character and the plus sign (+) indicates the string concatenation operator. The other parts are defined as follows:

Name	Definition	Examples
VERB	The method name of the HTTP request	PUT, GET, and POST
CONTENT-MD5	The MD5 value of the Body field in the HTTP request, which must be an uppercase string	875264590688CA6171F6228AF5BBB3D2
CONTENT-TYPE	The type of the Body field in the request	application/json
DATE	The standard timestamp header of the HTTP request, which follows the RFC 1123 format and uses GMT standard time	Mon, 3 Jan 2010 08:33:47 GMT
CanonicalizedHeaders	The string constructed by the custom headers prefixed with x-cms and x-acs in the HTTP request	x-cms-api-version:0.1.0\nx-cms-signature

Name	Definition	Examples
CanonicalizedResource	The string constructed by the HTTP request resources, as described in the following section	/event/custom/upload

CanonicalizedHeaders in the preceding table is constructed as follows:

- a. Convert the names of all HTTP request headers prefixed with `x - acs` and `x - acs` to lowercase letters.
- b. Sort the CMS custom request headers obtained in the preceding step in ascending lexicographical order.
- c. Delete any spaces on either side of a delimiter between the request header and content.
- d. Separate all headers and content with separators (`\ n`) to form the final CanonicalizedHeaders.

CanonicalizedResource in the preceding table is constructed as follows:

- a. Set CanonicalizedResource as an empty string ("").
- b. Place the URI that you want to access, such as `/ event / custom / upload`, between the quotation marks.
- c. If the request contains a query string (`QUERY_STRI NG`), add a question mark (?) and the query string to the end of the CanonicalizedResource string. `QUERY_STRI NG` is the lexicographic string of request parameters included in the URL. Equal signs (=) are placed between the names and values of parameters to form a string. The key-value pairs are then sorted in ascending

lexicographical order and connected with ampersands (&) to form a string.

The formula is as follows:

```
QUERY_STRING = " KEY1 = VALUE1 " + "&" + " KEY2 = VALUE2 "
```

3. Generate a digital signature for the request

Currently, the only supported signature algorithm is HMAC-SHA1. The following formula is used to generate a signature:

```
Signature = base16 ( hmac - sha1 ( UTF8 - Encoding - Of ( SignString ), AccessKeySecret ))
```

• Request parameters

Name	Type	Required?	Description
groupId	Long	Yes	The ID of an application group.
metricName	String	Yes	The name of a metric. A metric name can contain letters, digits, and connectors such as underscores (_), hyphens (-), periods (.), forward slashes (/), and backslashes (\). Other characters are invalid. The maximum length is 64 bytes. Excess characters will be truncated from the string.

Name	Type	Required?	Description
dimensions	Object	Yes	<p>The dimension map. All key-value pairs are strings. A string can contain letters, digits, and connectors such as underscores (<code>_</code>), hyphens (<code>-</code>), periods (<code>.</code>), forward slashes (<code>/</code>), and backslashes (<code>\</code>). The maximum number of key-value pairs is 10. The maximum length of a key is 64 bytes. The maximum length of a value is 64 bytes. Excess characters will be truncated from the string.</p>
time	String	Yes	<p>The time at which the metric value is generated. It supports timestamps in the "yyyyMMdd'T'HHmmss.SSSZ" format or long format, such as 20171012T132456.888+0800 or 1508136760000.</p>

Name	Type	Required?	Description
type	Integer	Yes	<p>The data type of the reported value. 0 indicates raw data and 1 indicates aggregate data.</p> <p>When you report aggregate data, we recommend that you report data for both 60s and 300s aggregation periods. Otherwise, you will not be able to query monitoring data that is older than seven days.</p>
period	String	No	<p>The aggregation period. Unit: second.</p> <p>If the preceding Type parameter is set to 1, this field is required. Valid values: 60 or 300.</p>
values	Object	Yes	<p>The collection of metric values. If type is 0, the key must be "value" and raw data is reported. CloudMonitor aggregates raw data over the aggregation period into several data types, such as maximum, count, and sum.</p>

Report data through the Java SDK

- Install Java SDK

When you install Java SDK through Maven, the following dependencies must be added:

```
< dependency >
  < groupId > com . aliyun . openservic es </ groupId >
  < artifactId > aliyun - cms </ artifactId >
  < version > 0 . 2 . 4 </ version >
</ dependency >
```

- Response elements

The system returns the HTTP status code 200.

- Examples

- Sample request

```
POST / metric / custom / upload HTTP / 1 . 1
Host : metrichub - cms - cn - hangzhou . aliyuncs . com
x - cms - api - version : 1 . 0
Authorizat ion : yourAccess KeyId : yourAccess KeySecret
Host : metrichub - cms - cn - hangzhou . aliyuncs . com "
Date : Mon , 23 Oct 2017 06 : 51 : 11 GMT
Content - Length : 180
x - cms - signature : hmac - sha1
Content - MD5 : E9EF574D1A EAAA370860 FE37856995 CD
x - cms - ip : 30 . 27 . 84 . 196
User - Agent : cms - java - sdk - v - 1 . 0
Content - Type applicatio n / json
[{" groupId " : 101 , " metricName " : "" , " dimensions " : {"
sampleName 1 " : " value1 " , " sampleName 2 " : " value2 " } , " time
" : "" , " type " : 0 , " period " : 60 , " values " : {" value " : 10 . 5
, " Sum " : 100 } } ]
```

- Sample response

```
{
  " code " : " 200 " ,
  " msg " : "" // The returned msg is null when the
reporting is normal .
}
```

- Sample code

- Report raw data

```
CMSClientI nit . groupId = 101L ; // Set a common group
ID .
CMSClient cmsClient = new CMSClient ( endpoint ,
accKey , secret ); // Initialize the client .
CustomMetr icUploadRe quest request = CustomMetr
icUploadRe quest . builder ()
. append ( CustomMetr ic . builder ()
. setMetricN ame ( " testMetric " ) // The
metric name
```

```

group ID .
        . setGroupId ( 102L )// Set a custom
        . setTime ( new Date () )
        . setType ( CustomMetric . TYPE_VALUE )//
The type is raw data .
, if )// The original value . Key must be an original
value .
        . appendValue ( MetricAttribute . VALUE
Add a dimension .
        . appendDimension ( " key ", " value " )//
1 " )// Add a dimension .
        . build ()
        . build ();
        CustomMetricUploadResponse response = cmsClient .
putCustomMetric ( request );// Report data .
        System . out . println ( JSONObject . toString (
response ) );
    
```

- Automatically report aggregate data for multiple aggregation periods

SDK supports data reporting after local aggregation. Data can be aggregated in periods of 1 minute and 5 minutes.

Data type	Description	Aggregate value	Memory consumption (excluding the name, dimension, individual time series, and individual aggregation periods)
value	Typical value type	All properties except LastValue	About 4 KB
gauge	Sample value	LastValue	4 bytes
meter	Sum and speed	Sum, SumPerSecond	50 bytes
counter	Count	SampleCount	10 bytes
timer	Computing time	SampleCount, CountPerSecond, Average, Maximum, Minimum, PXX(P10-P99)	About 4 KB

Data type	Description	Aggregate value	Memory consumption (excluding the name, dimension, individual time series, and individual aggregation periods)
histogram	Distribution	SampleCount, Average, Maximum, Minimum, PXX(P10-P99)	About 4 KB

```
// Initialization
    CMSClientInit groupId = 0L;
    CMSClient cmsClient = new CMSClient ( accKey ,
secret , endpoint );// Create a client .
    CMSMetricRegistryBuilder builder = new
CMSMetricRegistryBuilder ();
    builder . setCmsClient ( cmsClient );
    final MetricRegistry registry = builder . build
();// Create a registry , which includes two aggregation
n periods .
    // or has final MetricRegistry registry as
builder . build ( RecordLevel . _60S );. // Create a
registry that only includes aggregate data with an
aggregation period of 1 minute .
// Use value .
ValueWrapper value = registry . value ( MetricName . build ( "
value "));
value . update ( 6 . 5 );
// Use meter .
MeterWrapper meter = registry . meter ( MetricName . build ( "
meter "));
meter . update ( 7 . 2 );
// Use counter .
CounterWrapper counter = registry . counter ( MetricName .
build ( " counter "));
counter . inc ( 20 );
counter . dec ( 5 );
// Use timer .
TimerWrapper timer = registry . timer ( MetricName . build ( "
timer "));
timer . update ( 30 , TimeUnit . MILLISECON DS );
// Use histogram .
HistogramWrapper histogram = registry . histogram (
MetricName . build ( " histogram "));
histogram . update ( 20 );
// Use gauge .
final List list = new ArrayList ();
registry . gauge ( MetricName . build ( " gauge " ), new Gauge ()
{
```

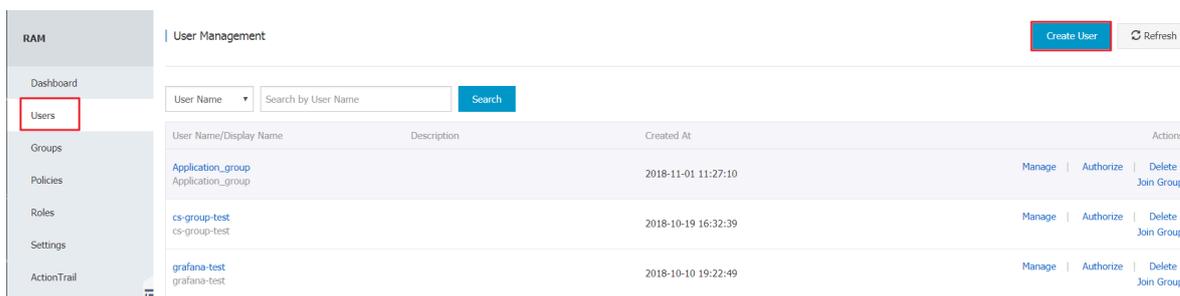
```
@ Override  
public Number getValue () {  
    return list . size ();  
}  
});
```

Report data through Alibaba Cloud CLI

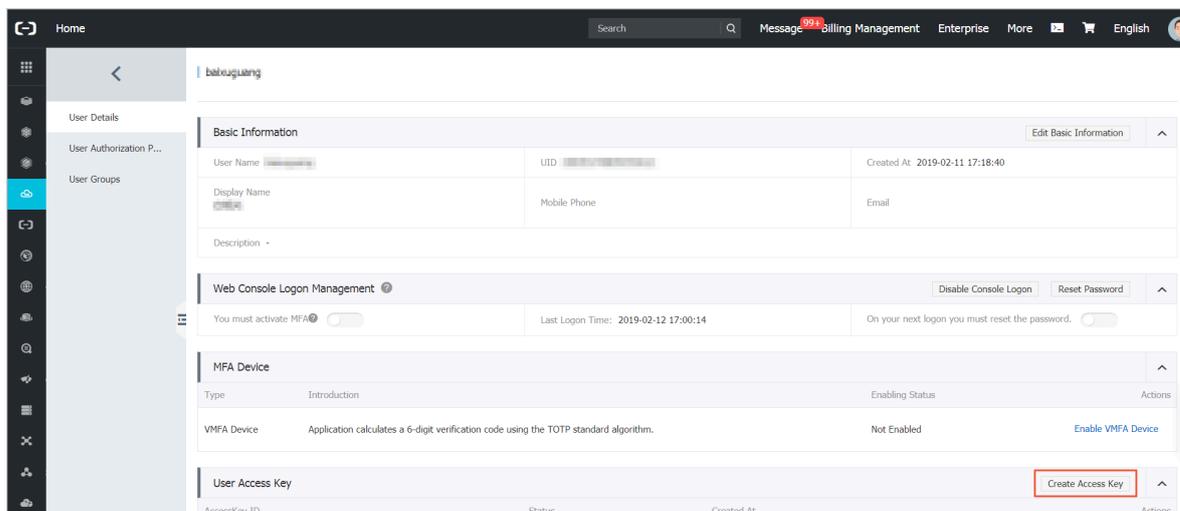
Prepare an Alibaba Cloud account

Ensure that you have created an Alibaba Cloud account, created a RAM user for your account, and generated a RAM user AccessKey (AK) with CloudMonitor permissions.

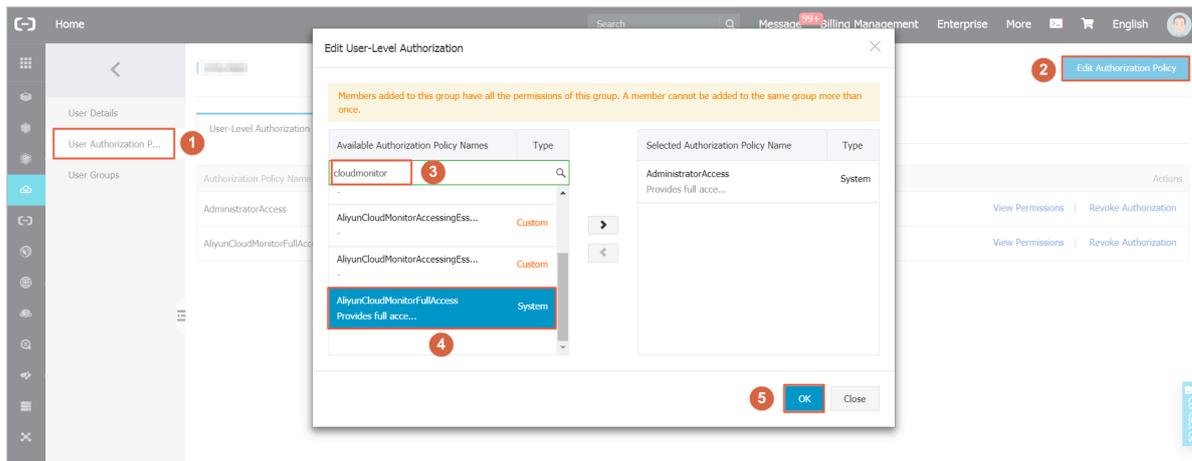
- Create a RAM user



- Generate an AccessKey ID and an AccessKey Secret for the RAM user.



- Grant CloudMonitor permissions to the RAM user.



Install Alibaba Cloud CLI

System requirement: Linux, Unix, or MacOS.

Environment requirement: Python 2.7.x has been installed.

1. Install Python.

- Skip this step if Python 2.7.x is already installed.
- Otherwise, run the following command on the command line to install Python:



Notice:

Ensure that wget is installed on your device.

```
wget https://www.python.org/ftp/python/2.7.8/Python-2.7.8.tgz (or download it in some other way and put it in a certain path)
tar -zxvf Python-2.7.8.tgz
cd Python-2.7.8
./configure
make
sudo make install
```

2. Install pip.

- Skip this step if pip is already installed on your device.
- Otherwise, run the following command on the command line to install pip:

```
curl "https://bootstrap.pypa.io/get-pip.py" -o "pip-install.py"
```

```
sudo python pip - install . py
```

- If the system displays the following or similar information, the installation was successful.

```
Successful ly installed pip - 7 . 1 . 2 setuptools - 18 . 7 wheel - 0 . 26 . 0
```

3. Install Alibaba Cloud CLI.

Make sure that you are using pip V7.x or later. Earlier versions of pip will cause the installation of Alibaba Cloud CLI to fail. You can run the following command to upgrade pip. Skip this step if you are using the latest version of pip.

- a. Run the following command on the command line to upgrade pip:

```
sudo pip install - U pip
```

If the system displays the following or similar information, the upgrade was successful.

```
Successful ly uninstalle d pip - 7 . 1 . 2
Successful ly installed pip - 8 . 1 . 2
```

- b. Run the following command to install Alibaba Cloud CLI:

```
sudo pip install aliyuncli
```

If the system displays the following or similar information, the installation is successful.

```
Successful ly installed aliyuncli - 2 . 1 . 2 colorama - 0 . 3 . 3 jmespath - 0 . 7 . 1
```

4. Configure Alibaba Cloud CLI.

```
~ sudo aliyuncli configure
Aliyun Access Key ID [***** a ]: youraccess
keyid
Aliyun Access Key Secret [***** b ]:
youraccess keysecret
Default Region Id [ cn - hangzhou ]: cn - hangzhou
Default output format [ json ]: json
```

Install CMS SDK

- The installation method for a Windows system is as follows:

```
cd C :\ Python27 \ Scripts
```

```
pip install aliyun-python-sdk-cms
```

- Run the following command to update the SDK:

```
pip install --upgrade aliyun-python-sdk-cms
```

- The installation method for a Linux system is as follows:

```
sudo pip install aliyun-python-sdk-cms
```

- Run the following command to update the SDK:

```
sudo pip install --upgrade aliyun-python-sdk-cms
```

Report monitoring data

Use the `PutCustomMetric` API to report the monitoring data.

- Example for a Windows system:

```
aliyuncli.exe cms PutCustomMetric -- MetricList "[{' groupId ': 1, ' metricName ': ' testMetric ', ' dimensions ': {' sampleName 1 ': ' value1 ', ' sampleName 2 ': ' value2 ' }, ' type ': 0, ' values ': {' value ': 10.5 } }]"
```

- Example for a Linux system:

```
aliyuncli cms PutCustomMetric -- MetricList "[{' groupId ': 1, ' metricName ': ' testMetric ', ' dimensions ': {' sampleName 1 ': ' value1 ', ' sampleName 2 ': ' value2 ' }, ' type ': 0, ' values ': {' value ': 10.5 } }]"
```

- If the data is reported successfully, status code 200 is returned.

```
{
  " Code ": " 200 "
}
```

Error codes

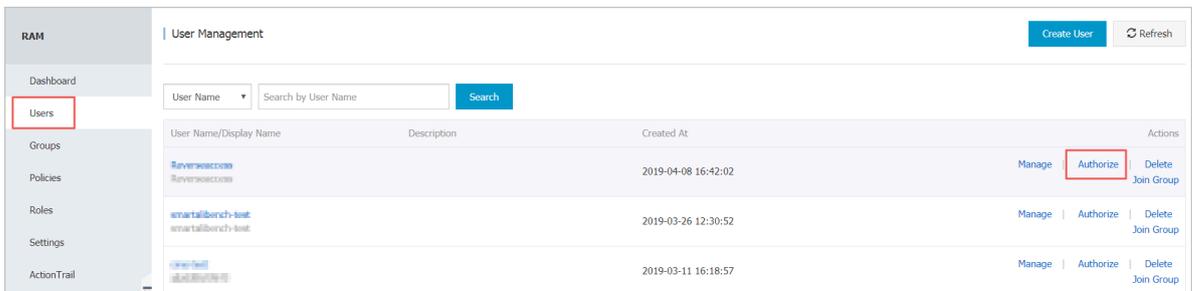
Error code	Description
200	Normal.

Error code	Description
206	Partially successful. If "reach Max time series num" is returned, you have reached the maximum amount of time series. We recommend that you purchase a higher quota or remove unnecessary time series . If "not allowed original value, please upgrade service" is returned, you are using the free edition, which cannot support raw data. If "type is invalid" is returned, the value of the type parameter is invalid. Ensure that the value of this parameter is 0 or 1.
400	Syntax error in a client request.
403	Verification failure, speed limit, or not authorized.
500	Internal server error

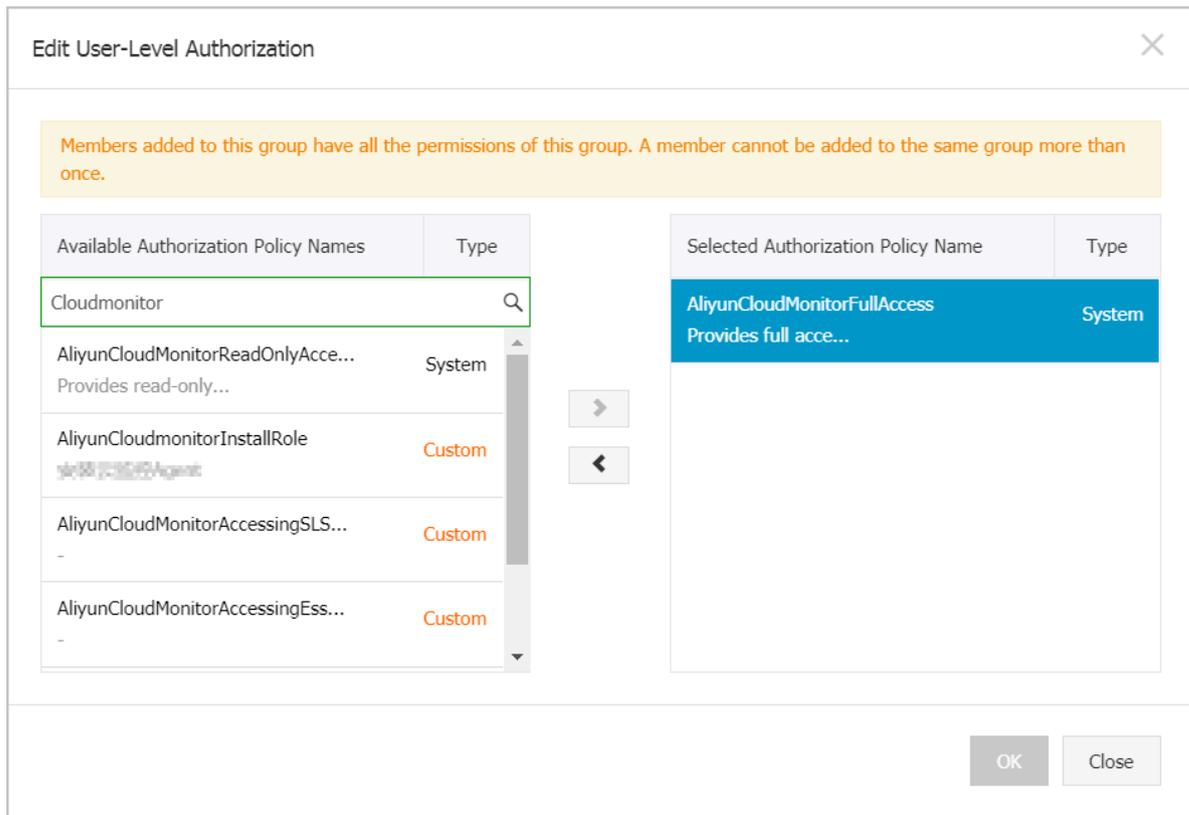
RAM user authorization

You must grant CloudMonitor permissions to the corresponding RAM user before event data can be reported with the RAM user AK. If these permissions are not granted, the prompt "cannot upload, please use ram to auth" will be displayed when you report data.

1. Log on to the [RAM console](#).
2. In the left-side navigation pane, click Users to go to the Users page.
3. Click Authorize in the Actions column corresponding to the RAM user.



4. On the authorization page, select `AliyunCloudMonitorFullAccess` and click OK.



10.3 View custom monitoring charts

This topic describes how to create a monitoring dashboard and add charts to view the custom monitoring data.

Background information

CloudMonitor allows you to customize what monitoring data is reported, and process and display that data in charts on the dashboard.

Prerequisites

Monitoring data is reported. For more information, see [Report monitoring data](#).

Procedure

Create a dashboard

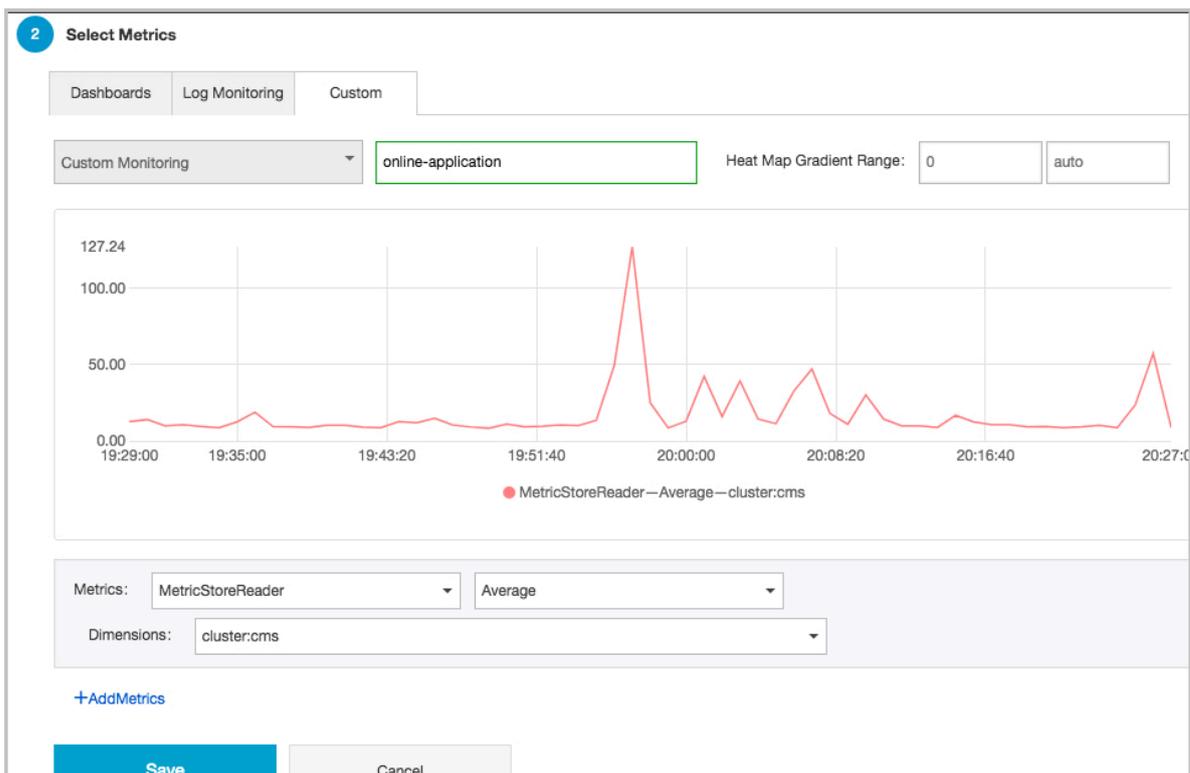
1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, choose Dashboard > Custom Dashboard. The Dashboards page is displayed.

3. Click **Create Dashboard** in the upper-right corner. In the dialog box that appears, enter a dashboard name, and click **Create**.



Add a chart

1. On the **Dashboards** page, click **Add View** in the upper-right corner to go to the chart configuration page.
2. Select a chart type from line chart, area chart, TopN table, heatmap, and pie chart.
3. Click the **Custom** tab.



4. On the **Custom** tab that appears, enter a chart name. Select the metrics, statistical methods, and dimensions to be displayed.

5. Click Save. After saving these configurations, you can view the custom monitoring chart.

