# Alibaba Cloud
# Cloud Monitor

## User Guide

Issue: 20190904

# 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:<br>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:<br>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:<br>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:<br>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 `cd / d C :/ windows` command to enter the Windows system folder. |
| *Italics* | It is used for parameters and variables. | `bae log list --instanceid` *Instance_ID* |
| [] or [a\|b] | It indicates that it is a optional value, and only one item can be selected. | `ipconfig` *[-all\|-t]* |

| Style | Description | Example |
| --- | --- | --- |
| {} or {a\|b} | It indicates that it is a required value, and only one item can be selected. | `swich {stand \| slave}` |

# Contents

# 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 applicatio n 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.

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

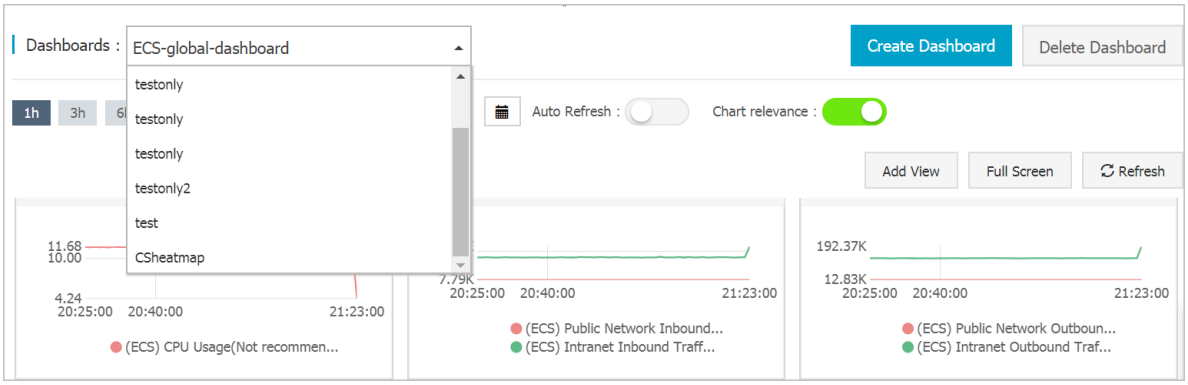


4. To view the dashboard in full screen, click Full Screen in the upper-right corner of the page.
5. Select a time range. Click the time range button at the top of the page. From there , you can quickly select the time range shown in the charts of the dashboard. The time range you select apply to all the charts on the dashboard.
6. 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.
7. The units of the metrics measured are displayed in parentheses for the chart name.
8. 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

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 page, click Create Dashboard.

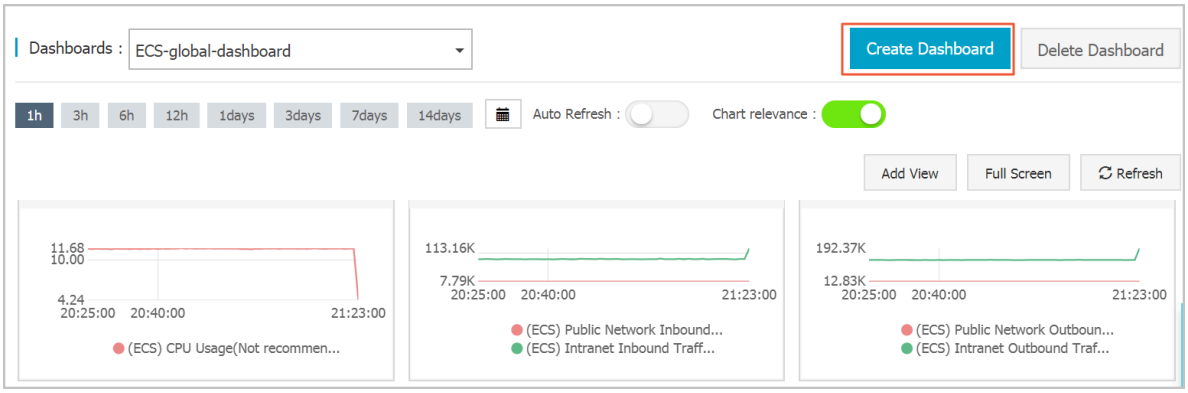


4. Enter the name of the dashboard.


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.

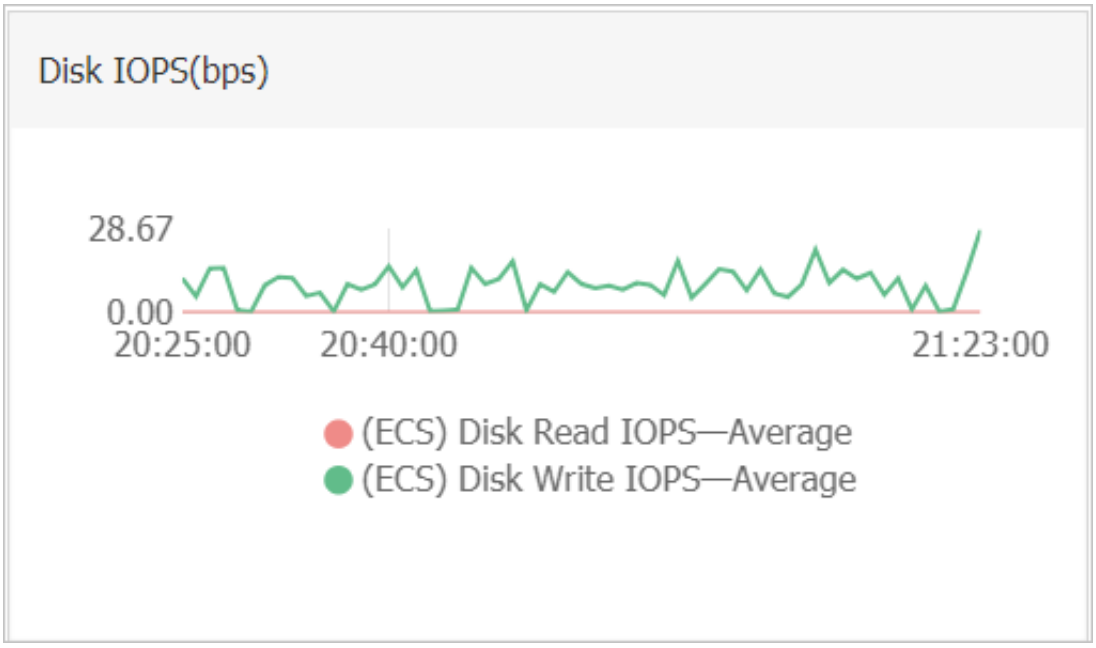

- Area chart: Displays monitoring data on a basis of time series. Multiple metrics can be added.

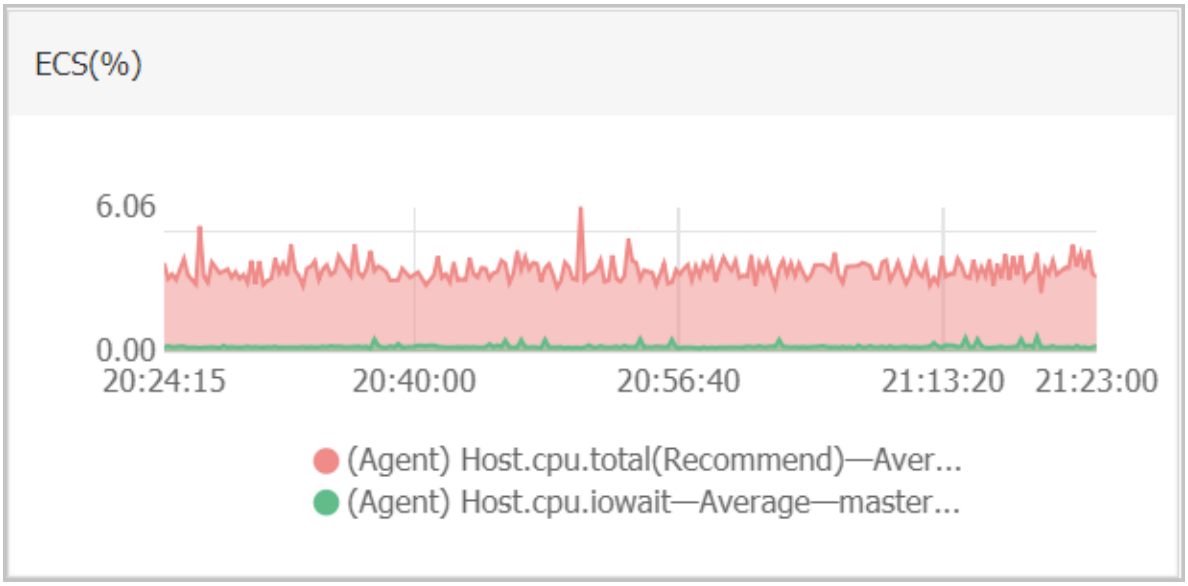


- Table: Displays real-time metric data in descending order. Each table displays up to 1,000 data records, which are either the first 1,000 records or the last 1,000 records. Only one metric can be added.

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.

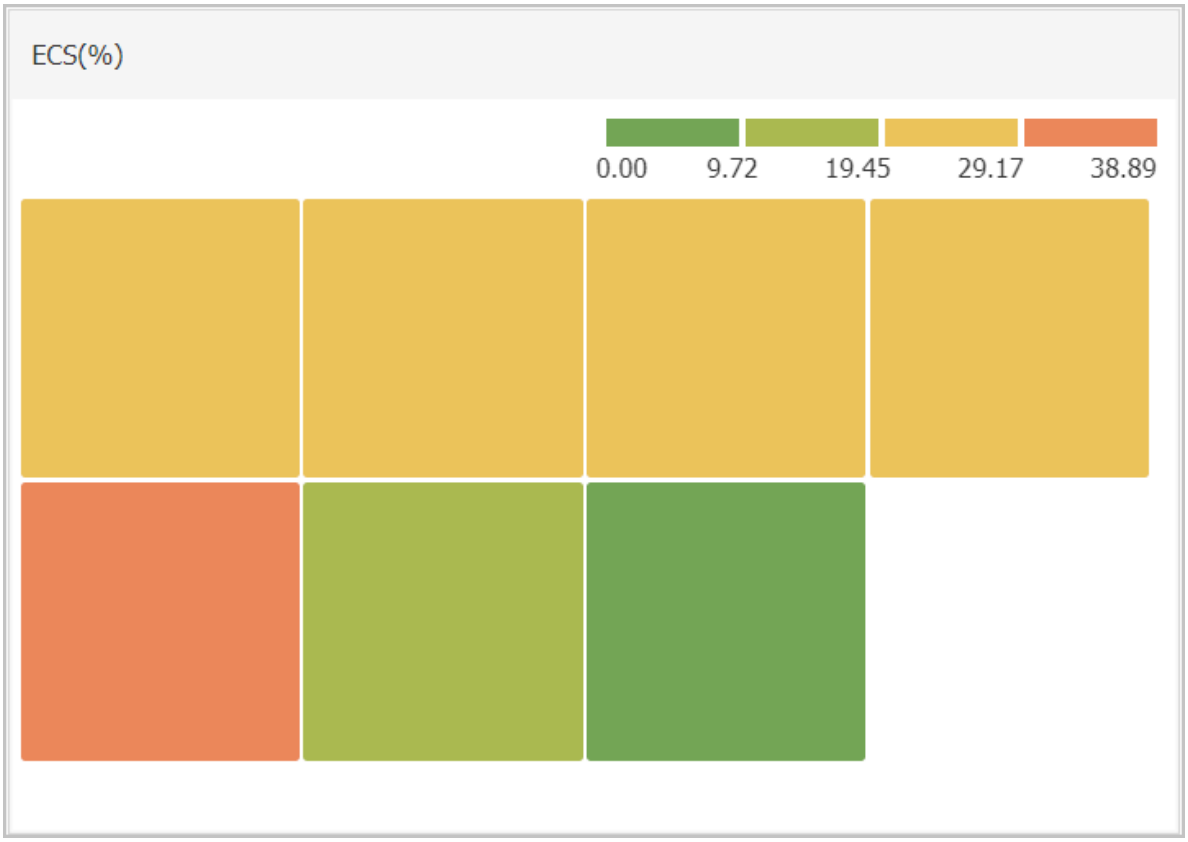


- Pie chart: Displays real-time metric data and can be used for data comparisons. Only one metric can be added.

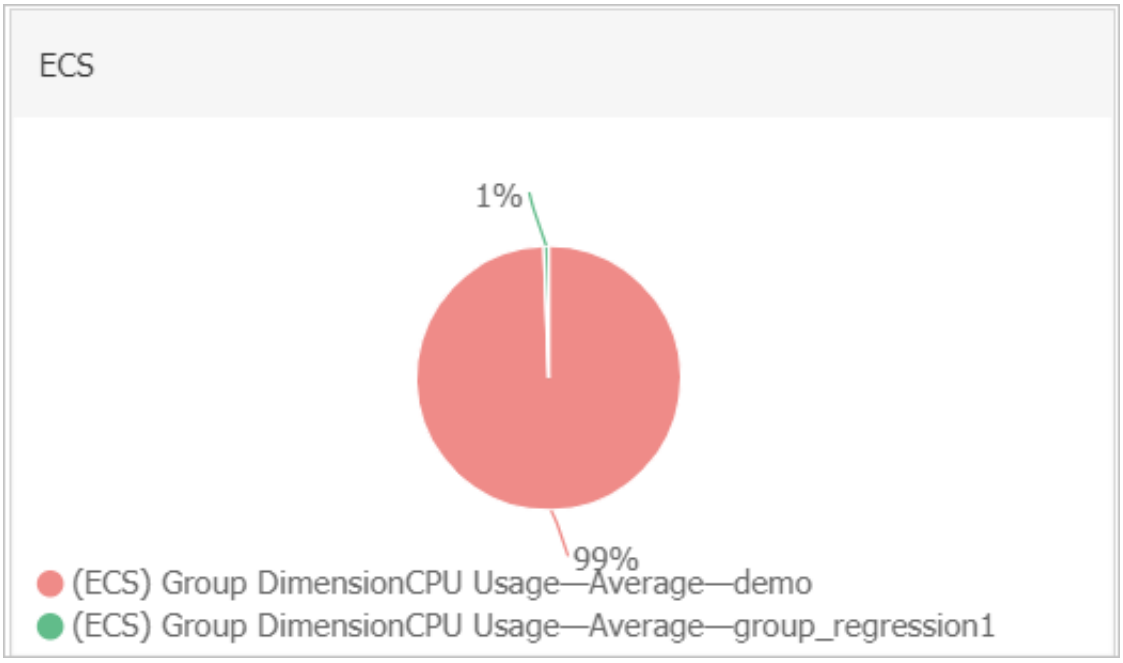

## Add a chart

 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.

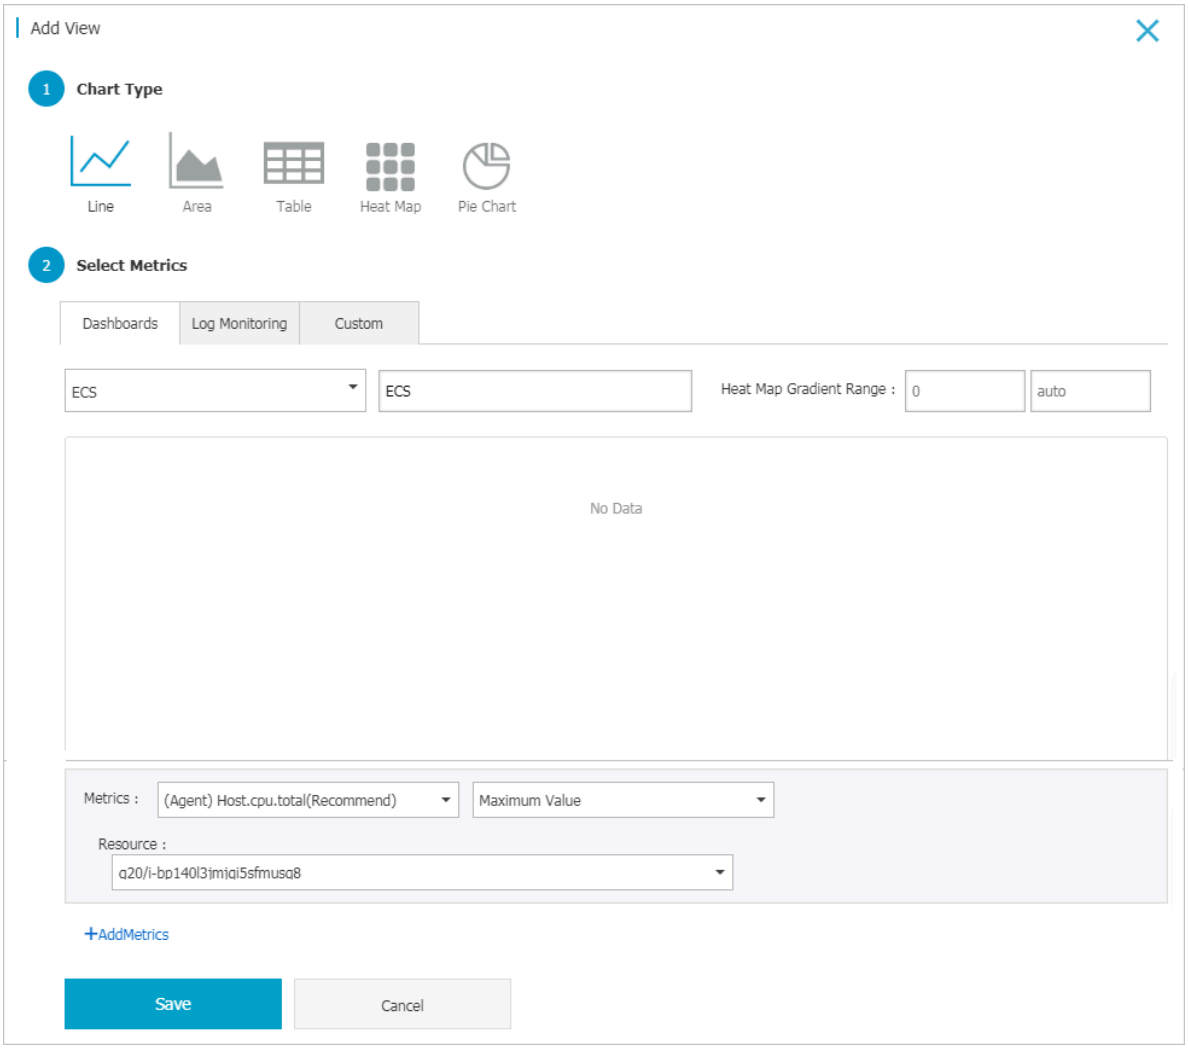


4. Select a chart type.
5. Choose from Dashboards, Log Monitoring, and Custom tab pages. In this example, click the Dashboards tab.
6. Select the target Alibaba Cloud product and enter a name for the chart.
7. Select the metric, the statistical method, and the resources.
   - Select the metric you want 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 localinsta ll 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.<br>Type: Select CMS Grafana Service. |
| HTTP | URL: `http :// metrics . cn - shanghai . aliyuncs . com` is used as an example. For more information, see Endpoints.<br>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.<br>The AK of your RAM user account is recommended. |

The following figure shows the configuration items.

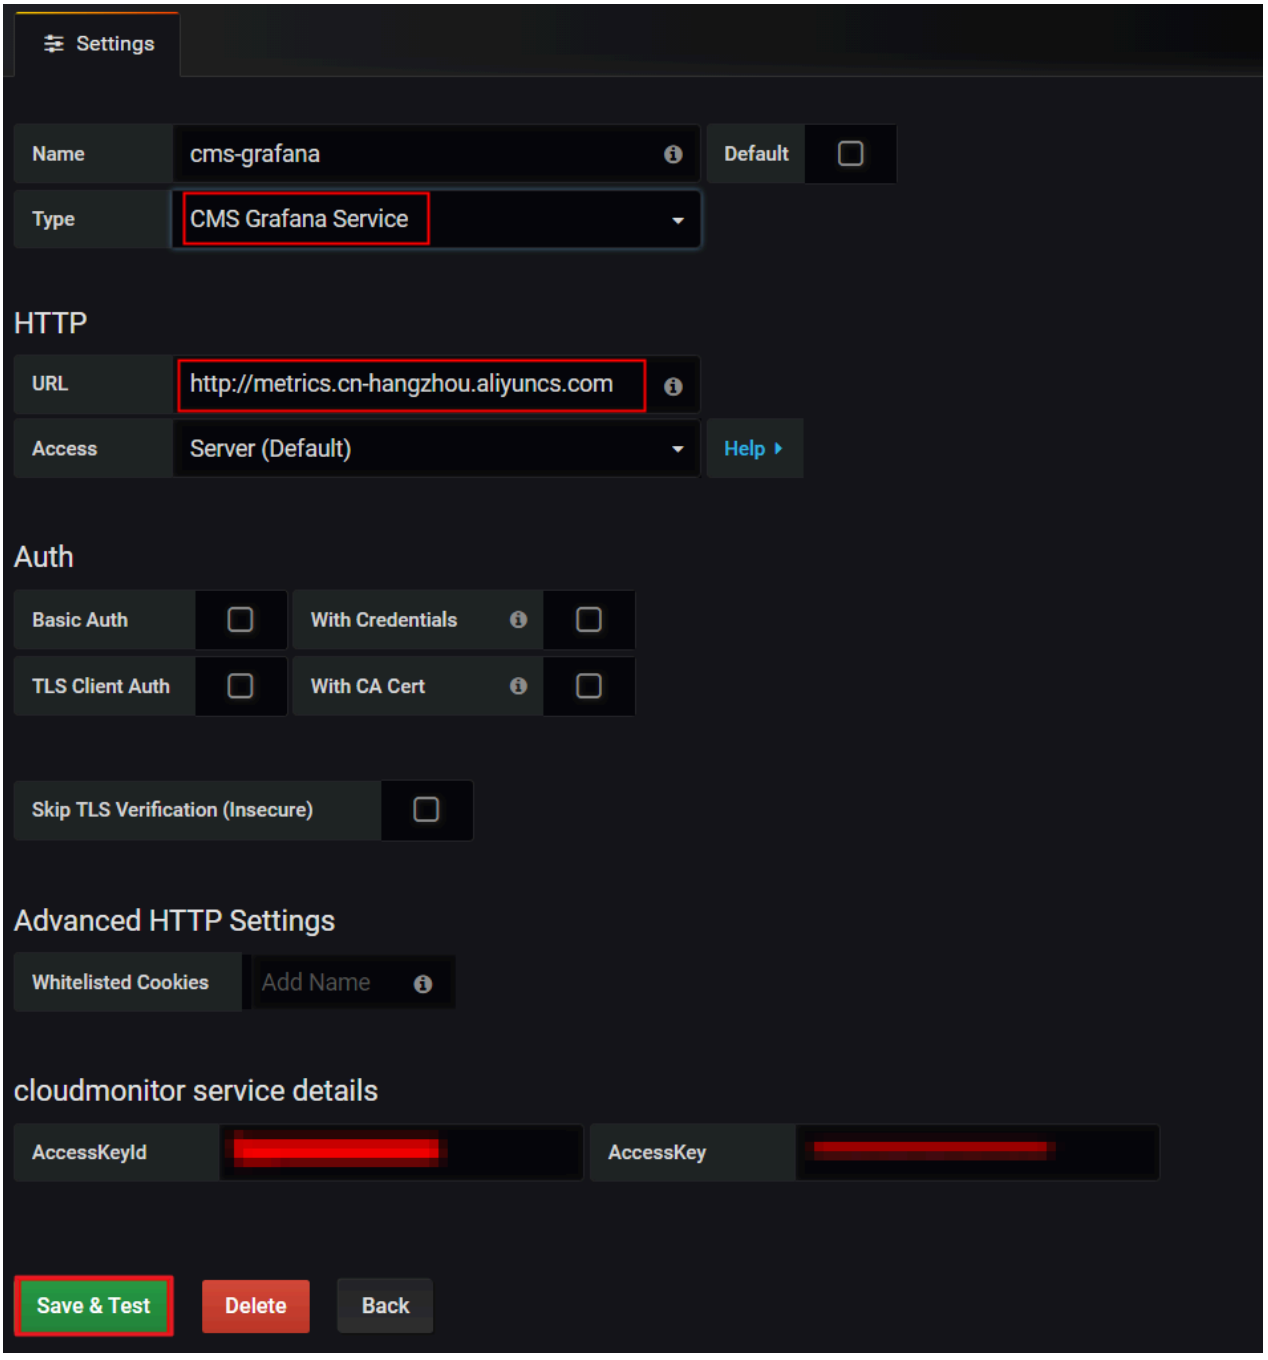


d. Click Save & Test.

3. Create a dashboard.

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

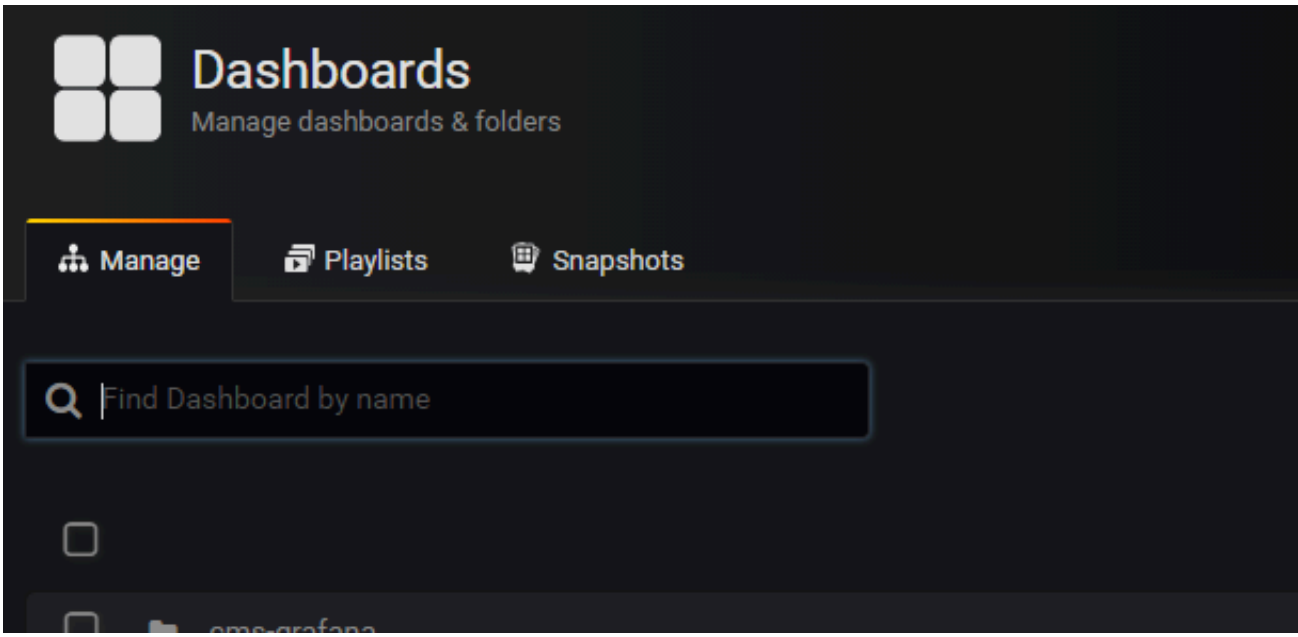


   b. Create a dashboard by using any of the following conditions:

      - 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` , `SampleCoun t` , `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.

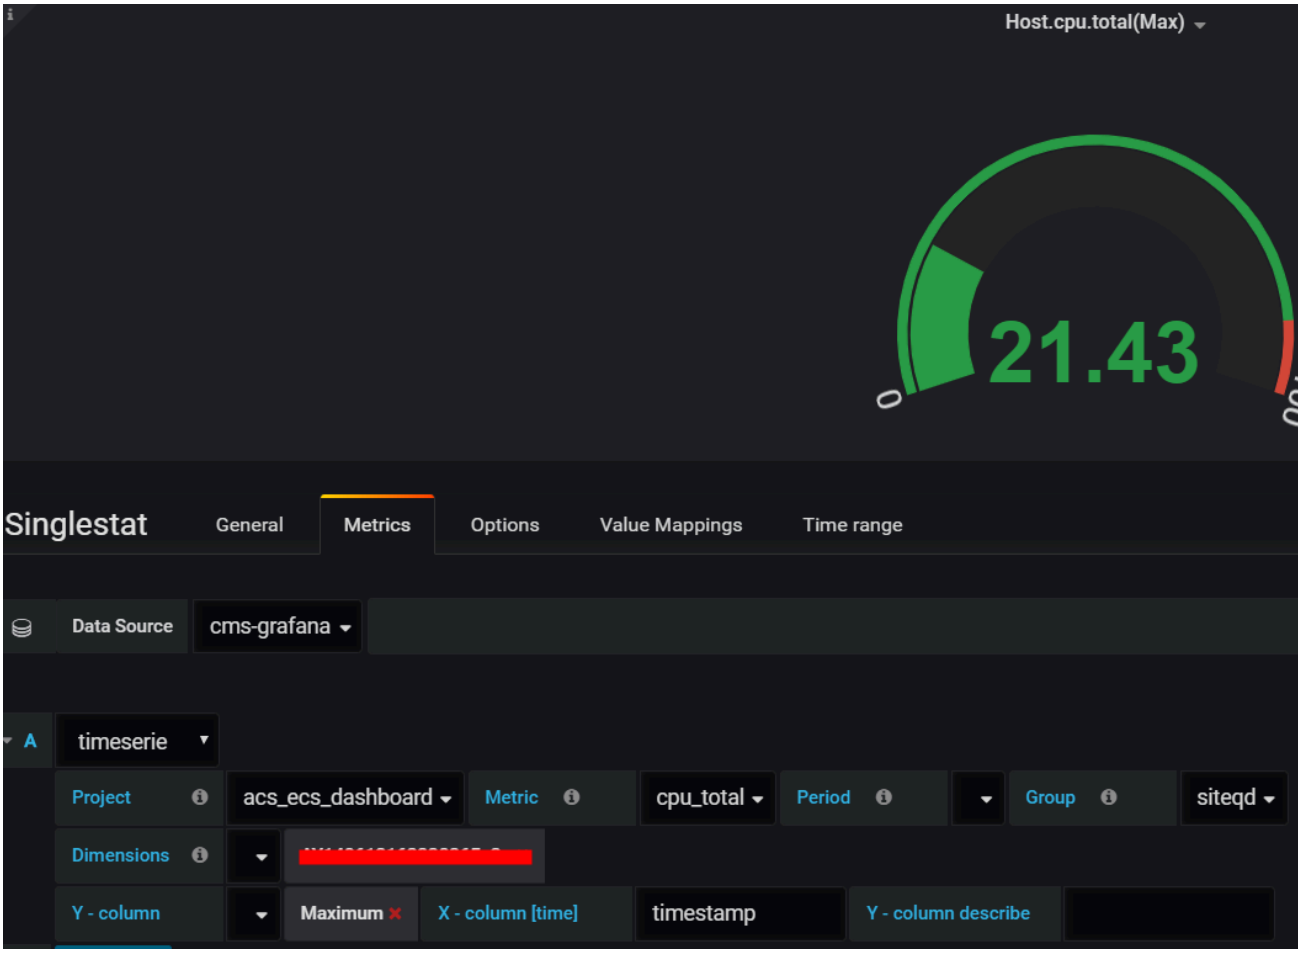


      For more information, see Singlestat.

6. View monitoring results.

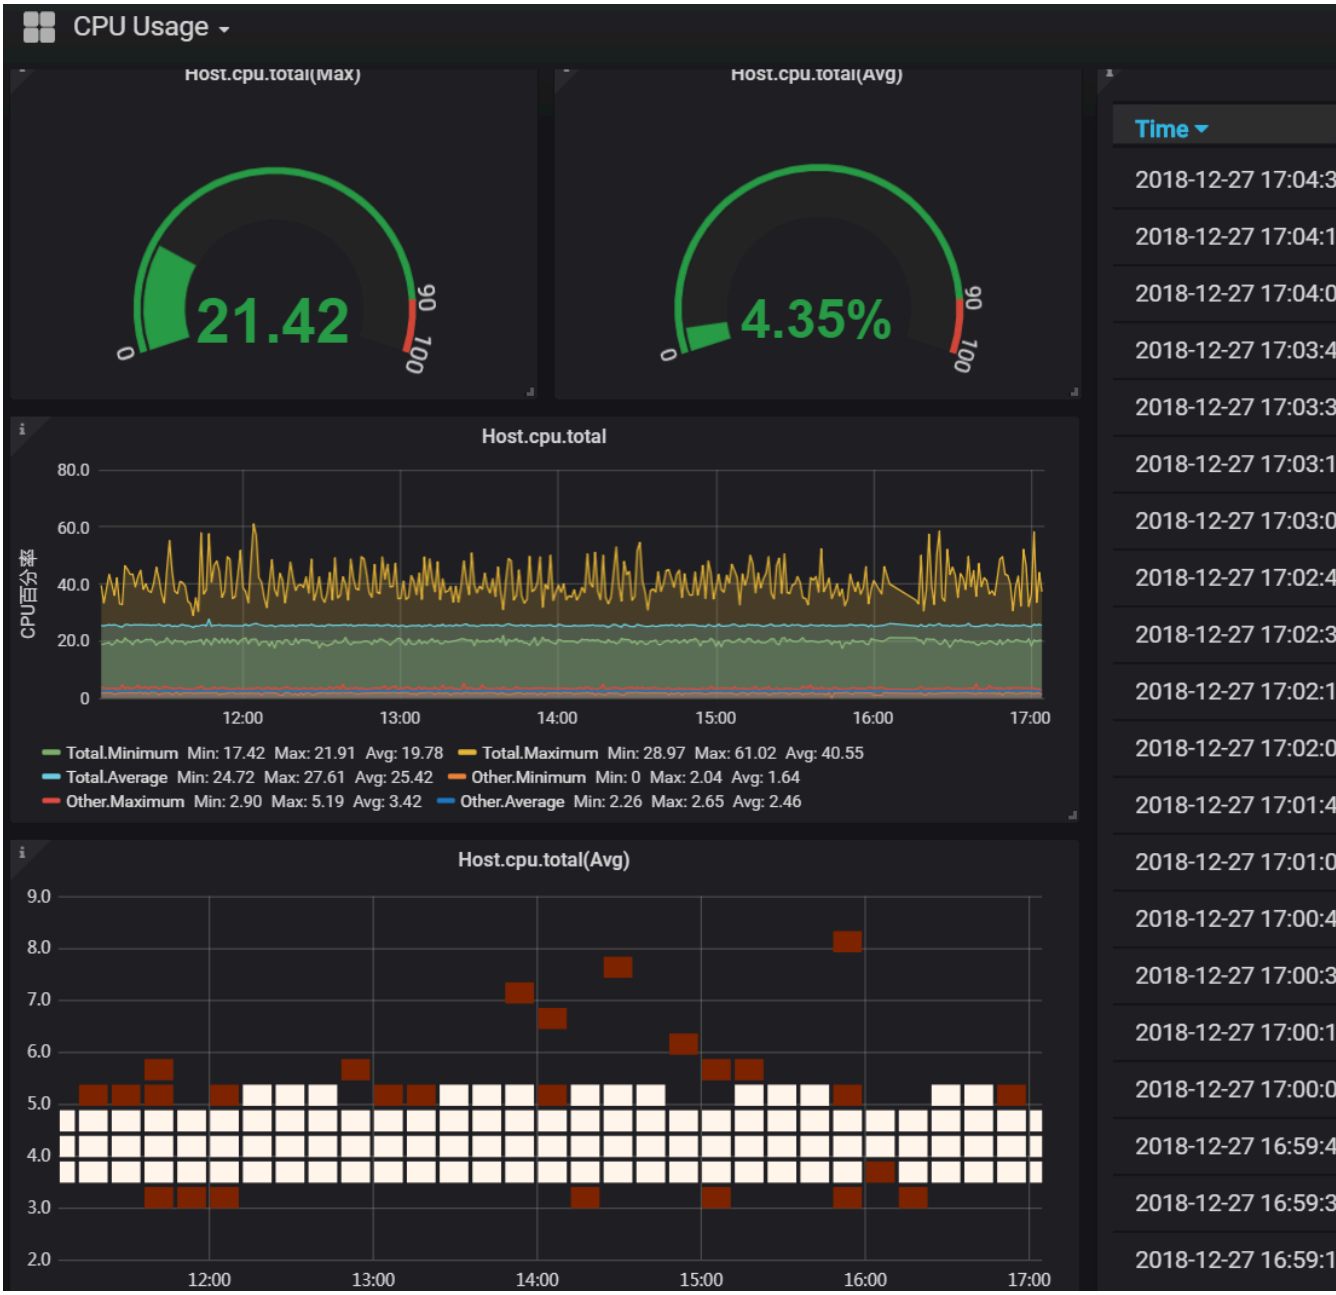

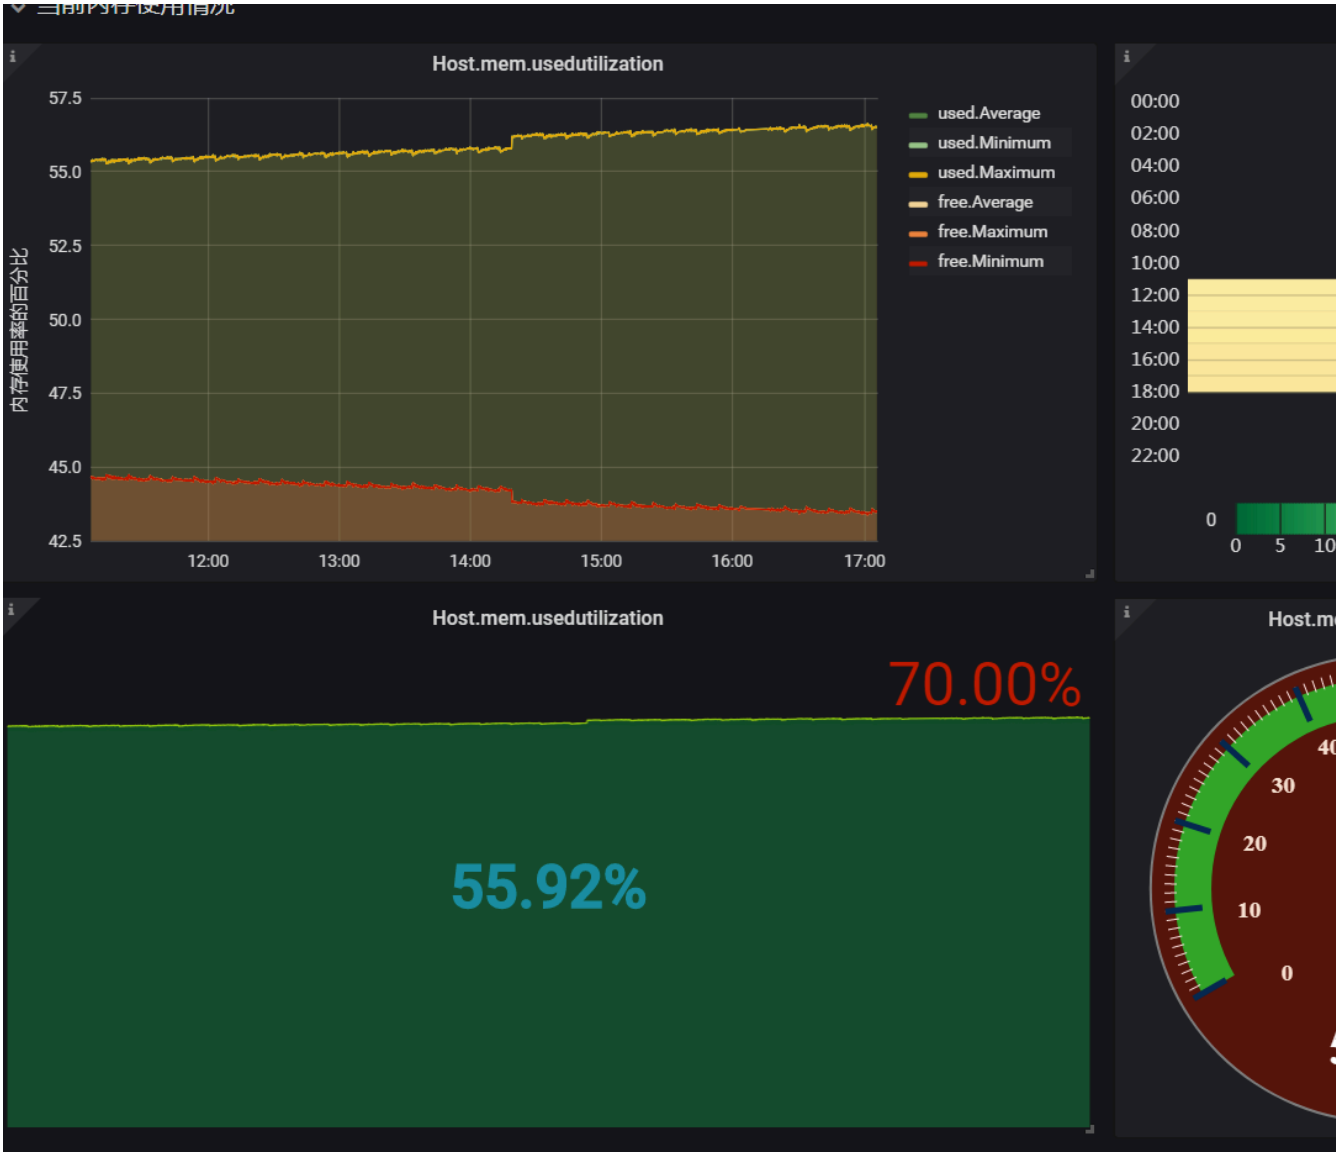
Host.mem.usedutilization
57.5
55.0
52.5
50.0
47.5
45.0
42.5
内存使用率的百分比
12:00
13:00
14:00
15:00
16:00
17:00
used.Average
used.Minimum
used.Maximum
free.Average
free.Maximum
free.Minimum
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
0
0 5 10
Host.mem.usedutilization
70.00%
55.92%
Host.m
40
30
20
10
0

# 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 `cloudmonitor/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\cloudmonitor\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.

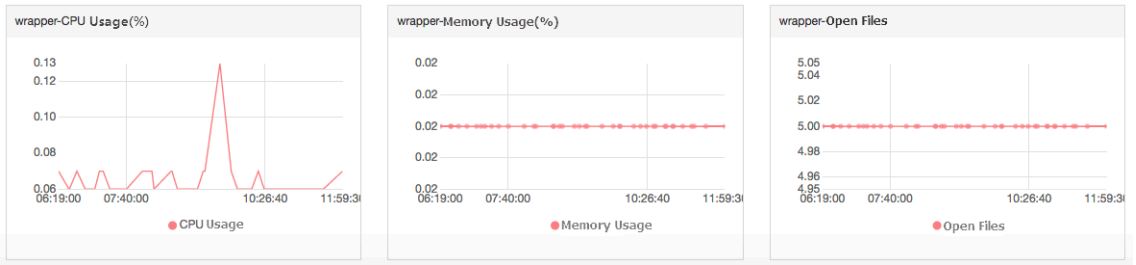


- The following figure shows the charts of the java process. The data points in the charts are dense and continuous. This means that the process continuously ranks in the top five processes with the most CPU usage.

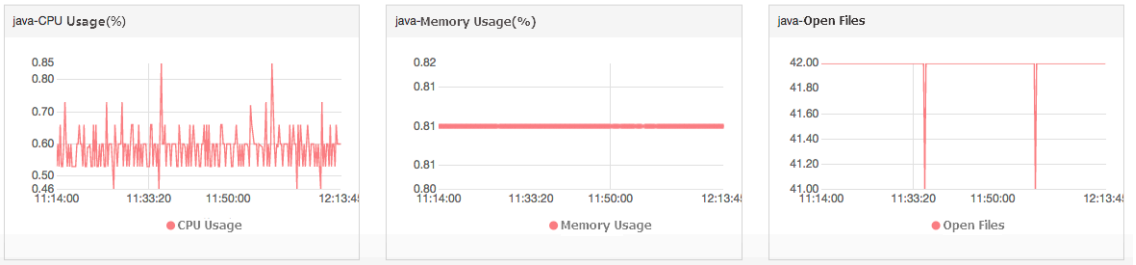


### Monitor the number of specified processes

You can lean 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 / ect / 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.

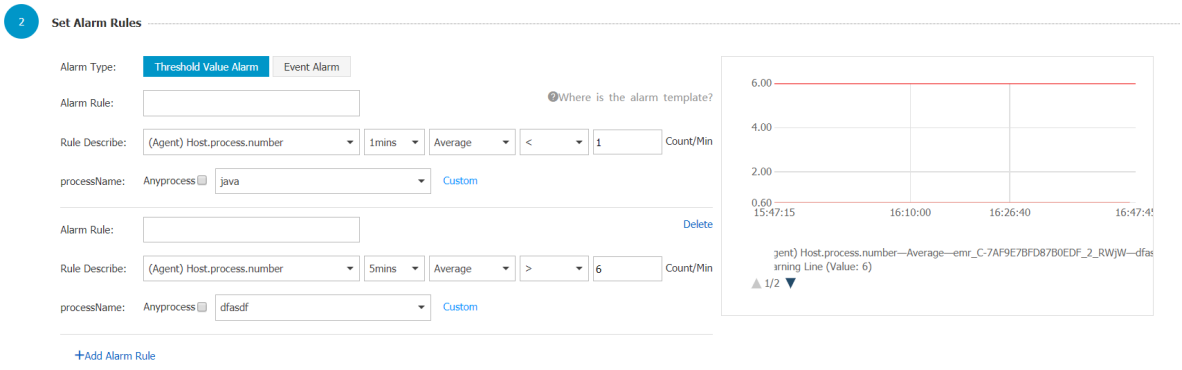


## 2.3 GPU monitoring

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_usedutilization | % | 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 | ℃ | 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 | ℃ | GPU temperature at the instance level | instanceId |
| instance_gpu_gpu_dutilization | % | 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_usedutilization | % | 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 | ℃ | GPU temperature at the application group level | groupId |
| group_gpu_gpu_usedutilization | % | 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_usedutilization | % | GPU memory usage at the application group level | groupId |
| group_gpu_power_readings_power_draw | W | GPU power at the application group level | groupId |

## Query GPU monitoring data in the console

After you have purchased an ECS instance of the GPU Compute type, you need to install the GPU driver and a CloudMonitor agent to be able to view and configure GPU monitoring charts and set alarm rules.

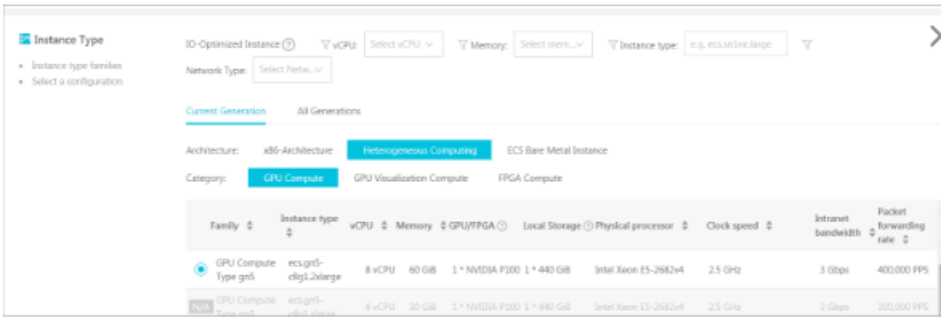


### View monitoring charts

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, click Host Monitoring.
3. On the Instances tab page, find the target instance and click the instance name.
4. Click the GPUMonitor tab to view the GPU monitoring charts.

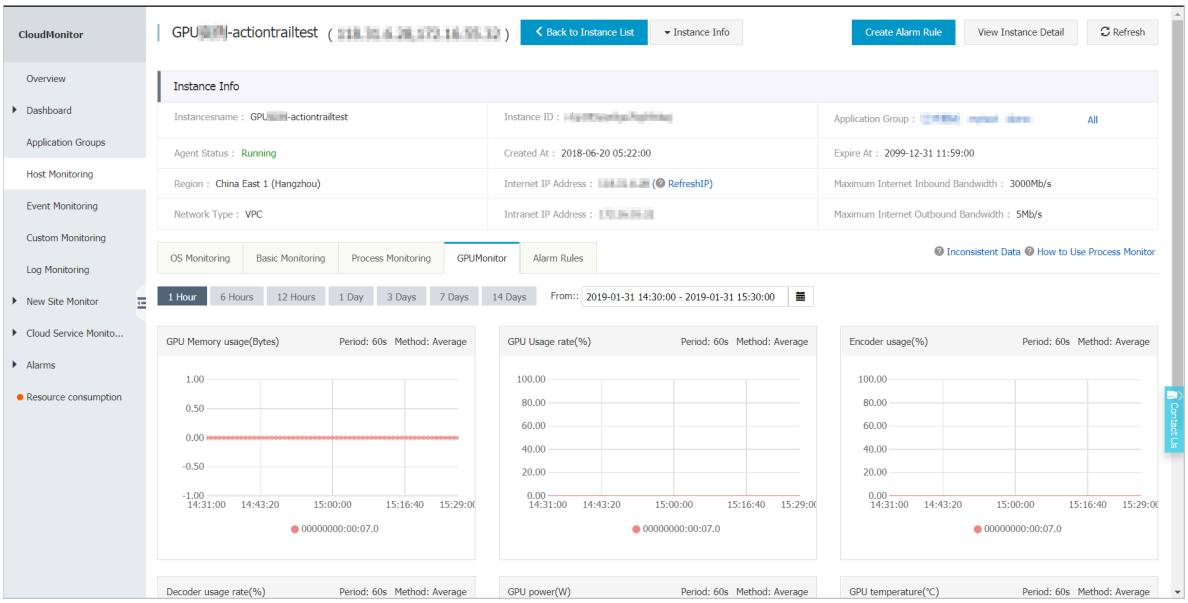


### Configure monitoring charts

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Dashboard > Custom Dashboard.
3. In the upper-right corner, click Create Dashboard.
4. In the displayed dialog box, enter a name for the dashboard and click Create.
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.

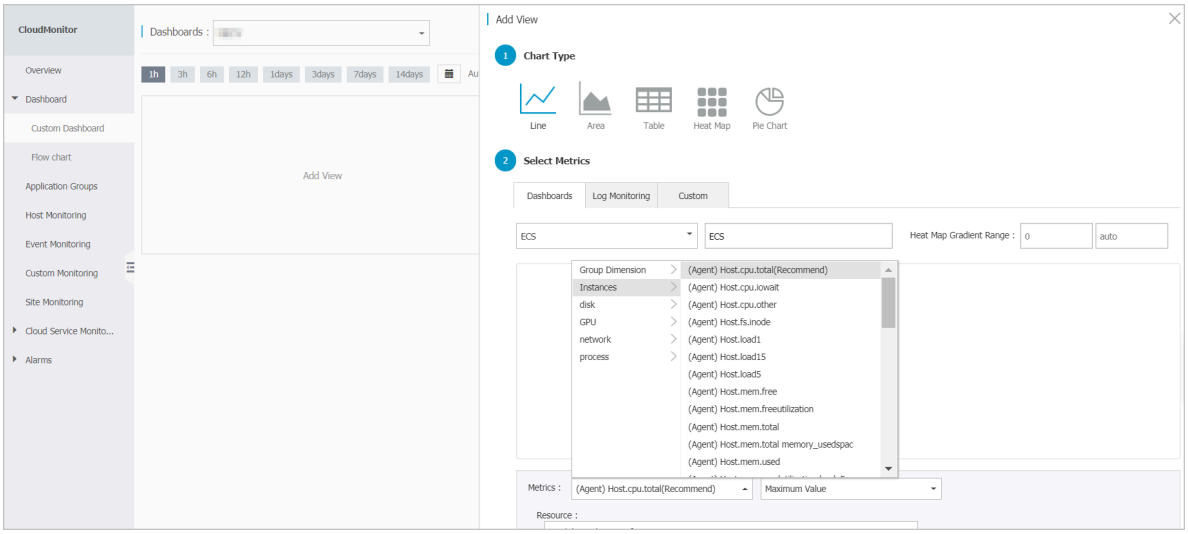


7. Click Save.

Set alarm rules

We recommended 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 GUP 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 Metrics

Host monitoring metrics include the metrics that a CloudMonitor agent monitors and the ECS basic metrics. The CloudMonitor agent collects monitoring data at a 15-second interval. CloudMonitor also collects monitoring data of ECS basic metrics at a 1-minute interval.

**Note:**
The ECS basic monitoring data may be different from the monitoring data that the CloudMonitor agent collects due to 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 for ECS basic monitoring data is one minute. The statistical period for monitoring data that the agent collects is 15 seconds. In the case of large monitoring data fluctuations, the value of ECS basic monitoring data is smaller than that of monitoring data that the agent collects due to load shifting.
- Different monitoring targets: the network traffic data collected by means of ECS basic metrics monitoring is used for billing. The data does not include the free traffic between ECS instances and Server Load Balancer (SLB) instances. However, the network traffic data collected by the agent indicates the actual network traffic of each network interface card (NIC). Therefore, the network traffic data collected by the agent is greater than that collected by means of ECS basic metrics monitoring. In this case, the data that the agent collects is greater than the actually purchased bandwidth or traffic quota.

Metrics collected by the CloudMonitor agent

- CPU metrics

  You can run the top command in Linux to understand the metrics listed in the following table.

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| Host.cpu.idle | The percentage of CPU currently not utilized. | % | Indicates the percentage of CPU currently in the idle status. |
| Host.cpu.system | The percentage of CPU currently occupied by the kernel. | % | Measures the CPU occupied by a system context switch. A high value indicates that many processes or threads are running on the host. |

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| 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. This metric is usually used for alarm purposes. |

- Memory metrics

  You can run the free command in Linux to understand the metrics listed in the following table. Data source: `/ proc / meminfo` . CloudMonitor uses the GlobalMemoryStatusExAPI function to collect Windows system data.

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| Host.mem.total | Total memory. | Byte | The total memory of the host.<br>Data source: MemTotal in the `/ proc / meminfo` directory. |

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| Host.mem.free | The amount of unused memory. | Byte | The amount of available memory in the system.<br>Data source: MemFree in the /proc/meminfo directory. |
| Host.mem.used | The amount of memory in use. | Byte | The amount of used memory in the system.<br>Calculation method: total - free. |

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| Host.mem.actualused | The memory actually used by the user. | Byte | Calculation method:<br>- When MemAvailable exists in the `/proc/meminfo` directory: total - MemAvailable.<br>- When MemAvailable does not exist in the `/proc/meminfo` directory: used - buffers - cached.<br>CentOS 7.2, Ubuntu 16.04, and later versions use the new Linux kernel. These versions provide more accurate memory estimation. For more information about the description of the MemAvailable column about the kernel, see commit . |

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| Host.mem.freeutilization | The percentage of available memory. | % | Calculation method:<br>- When MemAvailable exists in the `/proc/meminfo` directory: MemAvailable/total × 100%.<br>- When MemAvailable does not exist in the `/proc/meminfo` directory: (total - actualused)/total × 100%. |
| Host.mem.usedutilization | The memory usage. | % | Calculation method:<br>- When MemAvailable exists in the `/proc/meminfo` directory: (total - MemAvailable)/total × 100%.<br>- When MemAvailable does not exist in the `/proc/meminfo` directory: (total - free - buffers - cached)/total × 100%. |

- Metrics of average system loads

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

| Metric | Description | 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 run the df command in Linux to understand the disk usage and inode usage metrics.
  - You can run the iostat command in Linux to understand the disk read/write metrics.

| Metric | Description | 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.diskussage.total | The total disk storage. | Byte |
| Host.disk.readbytes | The number of bytes per second read from the disk . | Byte/s |

| Metric | Description | Unit |
| --- | --- | --- |
| Host.disk.writebytes | The number of bytes per second written to the disk . | Byte/s |
| Host.disk.readiops | The number of read requests per second received by the disk. | Request/s |
| Host.disk.writeiops | The number of write requests per second received by the disk. | Request/s |

- File system metrics

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| 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 you have used up inode numbers, you cannot create new files even if some disk space is available . Therefore, CloudMonitor must monitor the inode usage . The number of inodes indicates the number of files . A large number of small files can cause a high inode usage. |

- Network metrics
  - You can run the iftop command in Linux to understand the network metrics . You can run the ss command in Linux to understand the metrics of TCP connection data.
  - The system collects the following TCP connection data by default: TCP_TOTAL (the total number of connections), ESTABLISHED (the number of established connections), and NON_ESTABLISHED (the number of connections not in established status). To obtain such data, follow these steps:

    ■ Linux

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

    ■ Windows

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

| Metric | Description | Unit |
| --- | --- | --- |
| Host.netin.rate | The number of bits per second received by the NIC. This is the upstream bandwidth of the NIC. | Bit/s |
| Host.netout.rate | The number of bits per second sent by the NIC. This is the downstream bandwidth of the NIC. | Bit/s |
| Host.netin.packages | The number of packets per second received by the NIC. | Packet/s |
| Host.netout.packages | The number of packets per second sent by the NIC. | Packet/s |
| Host.netin.errorpackage | The number of incoming error packets detected by the drive. | Packet/s |

| Metric | Description | Unit |
| --- | --- | --- |
| Host.netout.errorpackages | The number of outgoing error packets detected by the drive. | Packet/s |
| Host.tcpconnection | The number of TCP connections in various statuses, including LISTEN, SYN_SENT, ESTABLISHED, SYN_RECV, FIN_WAIT1, CLOSE_WAIT, FIN_WAIT2, LAST_ACK, TIME_WAIT, CLOSING, and CLOSED. | |

- Process metrics
  - You can run the top command in Linux to understand the CPU usage and memory usage of processes. The CPU usage indicates the consumption of multi-core CPUs.
  - You can run the lsof command in Linux to understand Host.process.openfile.
  - You can run the ps aux |grep ‘Keyword’ command to understand Host.process.number.

| Metric | Description | Unit | Remarks |
| --- | --- | --- | --- |
| Host.process.cpu | The CPU usage of a process. | % | You cannot specify alarms for this metric. |
| Host.process.memory | The memory usage of a process. | % | You cannot specify alarms for this metric. |
| Host.process.openfile | The number of files opened by the current process. | File | You cannot specify alarms for this metric. |
| Host.process.number | The number of processes that match the specified keyword. | Process | You cannot specify alarms for this metric. |

### ECS basic metrics

If your host is an ECS instance, CloudMonitor automatically monitors the metrics listed in the following table after you purchase the ECS instance. You do not need to install the CloudMonitor agent to monitor these metrics. CloudMonitor collects ECS basic metrics at a 1-minute interval.

| Metric | Description | Unit |
| --- | --- | --- |
| ECS.CPUUtilization | CPU usage. | % |
| 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 per second read from the system disk. | Byte/s |
| ECS.SystemDiskWritebps | The number of bytes per second written to the system disk. | Byte/s |
| ECS.SystemDiskReadOps | The number of reads per second from the system disk. | Time/s |
| ECS.SystemDiskWriteOps | The number of writes per second to the system disk. | Time/s |
| ECS.InternetIn | Inbound Internet traffic. | Byte |
| ECS.InternetOut | Outbound Internet traffic. | Byte |
| ECS.IntranetIn | Inbound intranet traffic. | Byte |
| ECS.IntranetOut | Outbound intranet 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 #unique_21.

### 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 #unique_22.
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 #unique_24.

### 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/cloudmonitor/wrapper/bin/cloudmonitor.sh status

# Start
/usr/local/cloudmonitor/wrapper/bin/cloudmonitor.sh start

# Stop
/usr/local/cloudmonitor/wrapper/bin/cloudmonitor.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_linu  x . sh )"
```

**China (Hong Kong) (Hong Kong) 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_linu  x . 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_linu  x . 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_linu  x . 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_linu  x . 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_linu  x . 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_linu  x . 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 - shenzhen - 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 / UninstallA pp - 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 |
| | 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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| China (Hong Kong ) (Hong Kong) 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 |
| | 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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| 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 |
| | 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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| 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 |
| | 100.100.80.13:80<br>[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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| | 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<br>[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<br>[47.91.8.7:443] | Egress | Used to collect monitoring data to CloudMonitor |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| 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 #unique_27.

### 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.

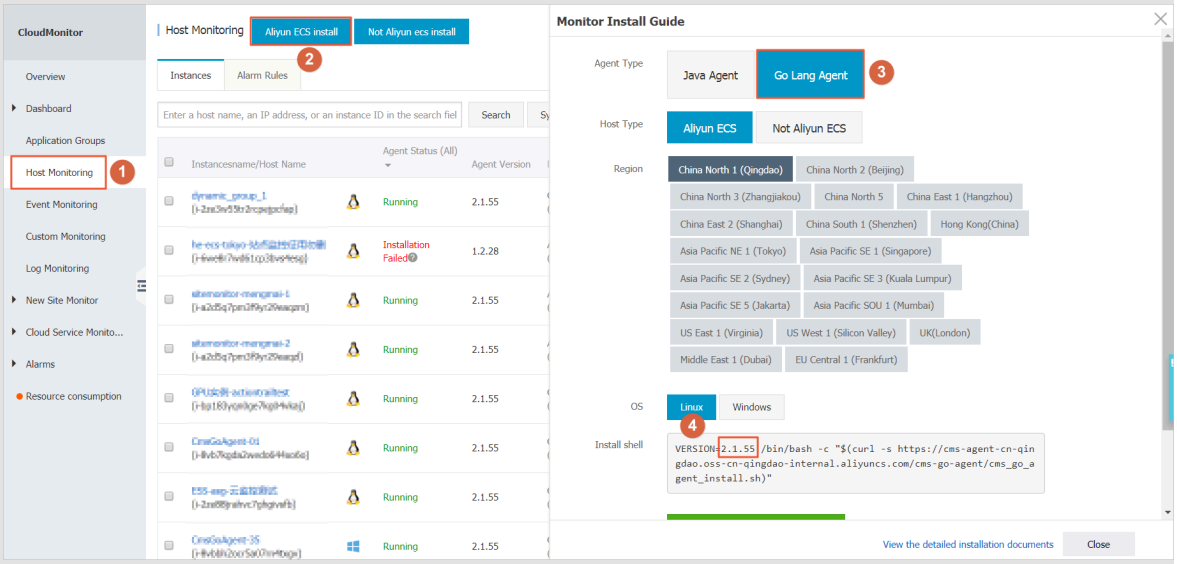


## Frequently used commands

```
#  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 )"
```

China (Hong Kong) (Hong Kong) 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_age  nt_install . sh )"
```

**China South 1 Finance Cloud (Shenzhen) cn-shenzen-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_age  nt_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.

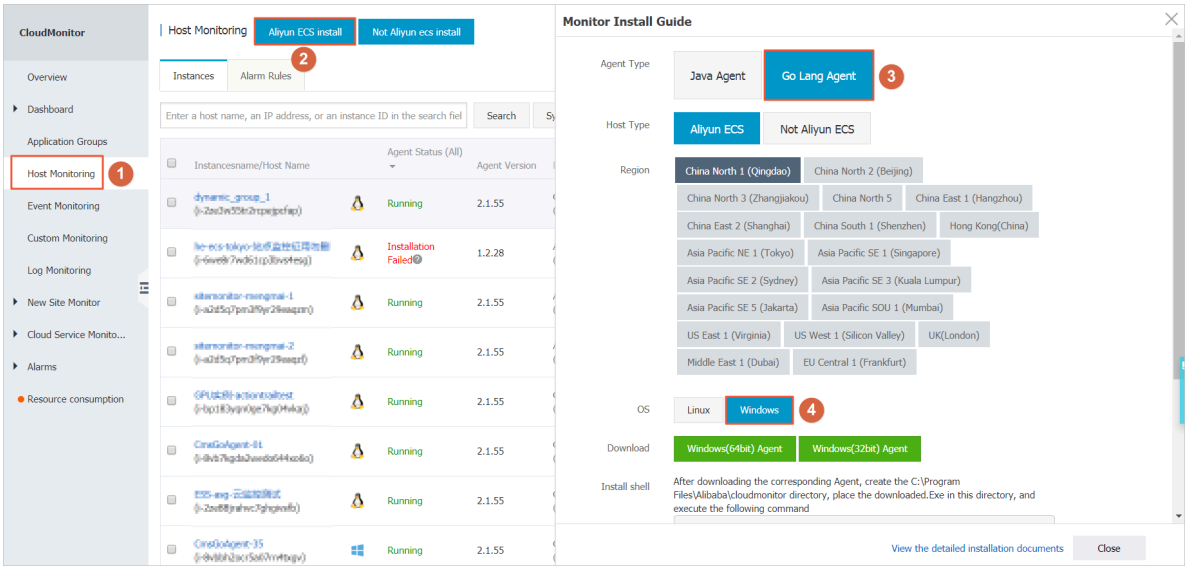


2. Open the Command Prompt as an administrator.
3. Run the following command:

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

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

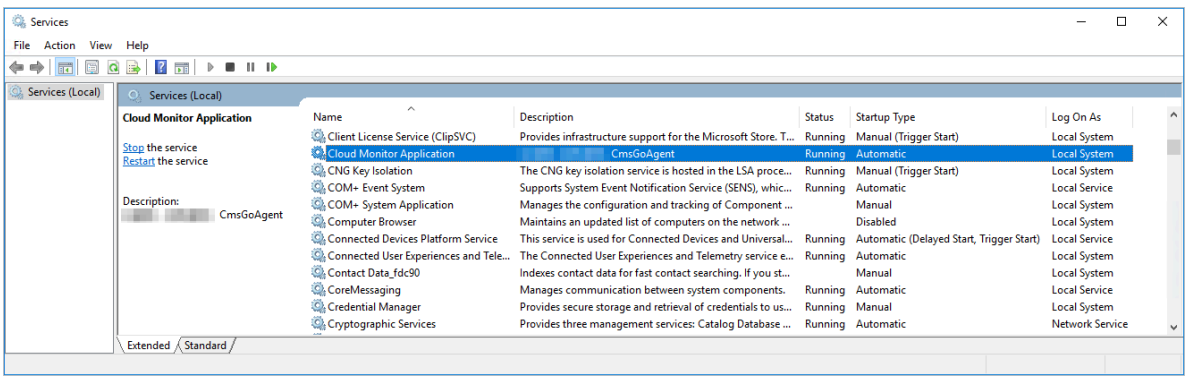


### Uninstall procedure

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

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

3. 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 |
| | 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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| China (Hong Kong ) (Hong Kong) 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 |
| | 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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| 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 |
| | 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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| 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 |
| | 100.100.80.13:80<br>[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 |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| | 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<br>[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<br>[47.91.8.7:443] | Egress | Used to collect monitoring data to CloudMonitor |

| Region | IP | Direction | Description |
| --- | --- | --- | --- |
| 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 #unique_30.

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 MemAvailab le` 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

## 3.1 Overview

This topic provides an overview of site monitoring and relevant application scenarios.

### Scenarios

Site monitoring is a function that simulates user access requests to help you better analyze the behavior of users to your services. This function is available in all Alibaba Cloud regions. The typical application scenarios of site monitoring are as follows.

### Analyze service performance

You can create a site monitoring task to obtain data, such as the time for the Domain Name Server (DNS) to resolve a domain name, the time when a connection is established, the time when an endpoint receives the first packet after sending a request, and the time when packets start to download. The data can be helpful for you to discover issues in your services, allowing you to achieve better performance.

### Compare your service performance with that of your peers

You can add the sites you service and those of your peers in the CloudMonitor console as monitoring items, and specify the probe points to detect network quality and service performance at the sites that you service and those of your peers.

### Get probe coverage

Site monitoring is available in all Alibaba Cloud regions. These regions can simulate user behavior and send access requests.

Detection protocol types

| Detection type | Function |
| --- | --- |
| HTTP or HTTPS | Sends an HTTP or HTTPS request to a specific URL or IP address to obtain the availability metrics, response time, and status code. In advanced settings, you can select a request method ( `GET` , `POST` , or `HEAD` ), set cookie and header information, and determine whether the page content matches the specified content. |
| PING | Sends an ICMP request that carries the `ping` command to a specific URL or IP address. This allows you to obtain the availability metrics, response time, and packet loss rate. |
| TCP | Sends a TCP request to a specific port to obtain the availability metrics, response time, and status code. In advanced settings, you can set the TCP request content and the match response content. |
| UDP | Sends a UDP request to a specific port to obtain the availability metrics, response time, and status code. In advanced settings, you can set the UDP request content and the match response content. |
| DNS | Sends a DNS request to a specific domain to obtain the availability metrics, response time, and status code. In advanced settings, you can specify the type of record that you want to query: A, MX, NS, CNAME, TXT, or ANY. |
| POP3 | Sends a POP3 request to a specific URL or IP address to obtain the availability metrics, response time, and status code. In advanced settings, you can set the port , username, password, and whether to establish a secured connection. |

| Detection type | Function |
| --- | --- |
| SMTP | Sends an SMTP request to a specific URL or IP address to obtain the availability metrics, response time, and status code. In advanced settings, you can set the port, username, password, and whether to establish a secured connection. |
| FTP | Sends an FTP request to a specific URL or IP address to obtain the availability metrics, response time, and status code. In advanced settings, you can set the port and whether to establish a secured connection. |

## 3.2 Create a site monitoring task

This topic describes how to create a site monitoring task. It can be used to analyze network quality and service performance.

### Background information

Site monitoring is a function that simulates user access requests to help you better analyze the behavior of users to your services. This function is available in all Alibaba Cloud regions. By using site monitoring, you can perform the following actions:

- Create a site monitoring task to obtain data such as the time for the Domain Name Server (DNS) to resolve a domain name, the time when a connection is established, the time when an endpoint receives the first packet after sending a request, and the time when packets start to download. The data can be helpful for you to discover issues in your services, allowing you to achieve better performance.
- Add the sites you service and those of your peers on the CloudMonitor console as monitoring items, and specify the probe points to detect network quality and service performance at the sites that you service and those of your peers.
- Simulate user behavior and send access requests from all Alibaba Cloud regions.

### Before you begin

- If you want to set alarm rules when creating a site monitoring task, we recommend that you create contacts and contact groups first. You can select the contact groups when setting alarm rules, and the contact groups will receive notifications when

alarms are reported. For information about how to create contacts and contact groups, see Create an alert contact and an alert contact group.

- If you want to enable the alarm callback function when setting alarm rules, you need to provide a callback URL that is accessible via the Internet. In addition, enable the URL callback function in your O&M system or message system.

## Procedure

> Note:
> When creating a site monitoring task, you can choose to set alarm rules or not as needed.

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose New Site Monitor > Site Manage.
3. On the Site Monitoring page, click New Monitoring Task.

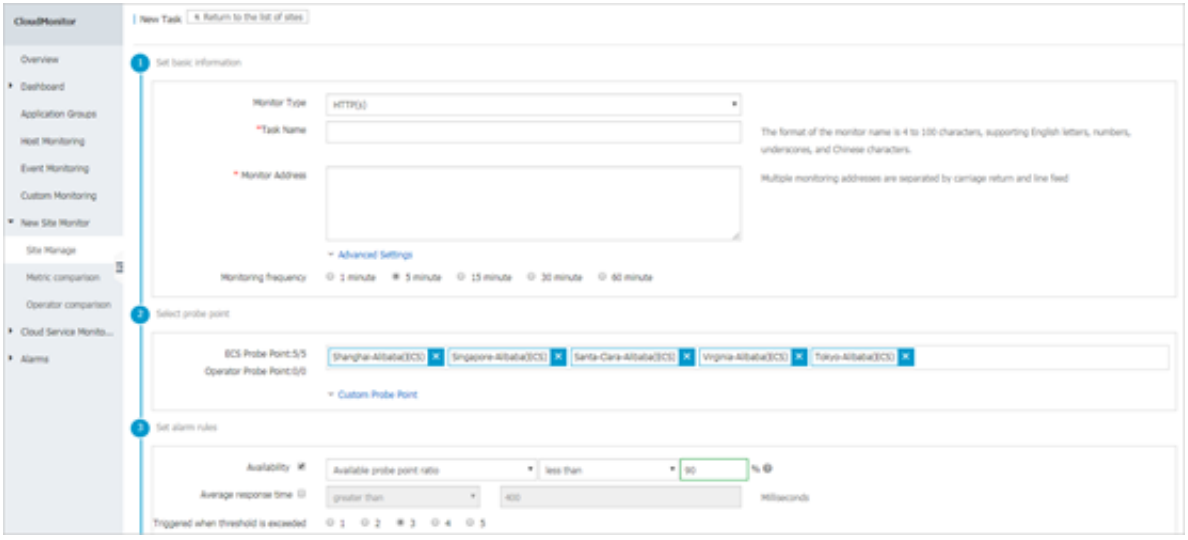

4. On the New Task page, set basic information for the site monitoring task.

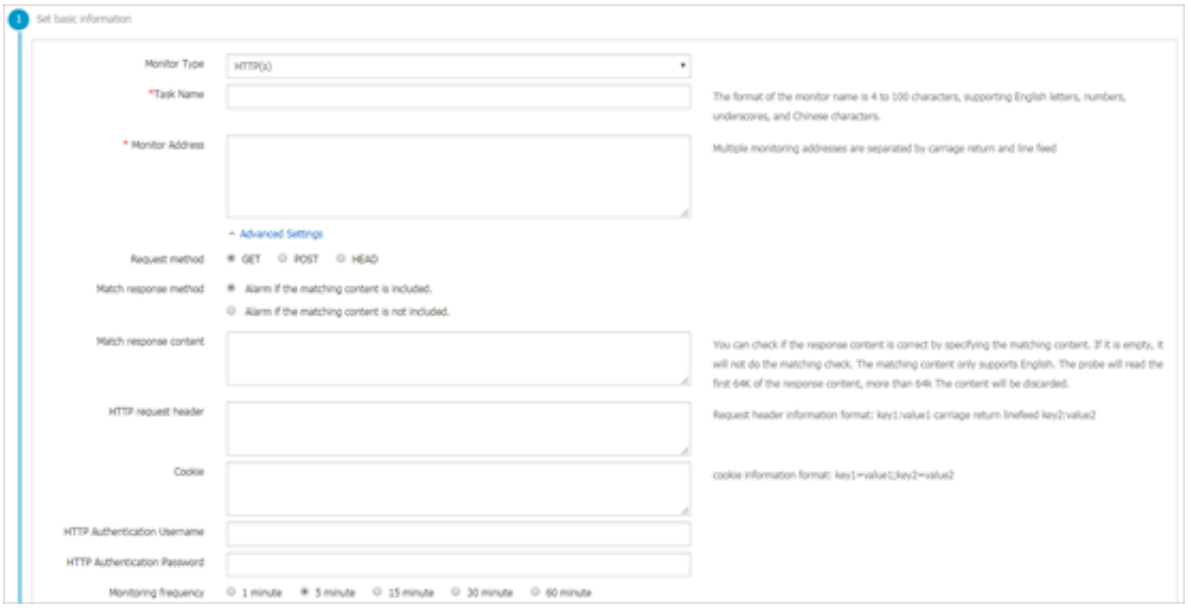


- Monitor Type: Select a monitoring protocol. Valid values: HTTP(s) | PING | TCP | UDP | DNS | SMTP | POP3 | FTP.
- Task Name: the name of the task. The task name contains 4 to 100 characters including letters, numbers, and underscores (_).
- Monitor Address: the destination address that you want to monitor. Separate multiple addresses into new lines. When you save the task settings, each address is saved as a job.
- Monitoring frequency: the interval at which the destination address is monitored regularly. Valid values: 1 minute | 5 minute | 15 minute | 30 minute | 60 minute. For example, if you set Monitoring frequency to 1 minute, each probe point monitors the destination addresses at 1-minute intervals.
- Advanced Settings: The advanced settings available vary depending on the protocol specified by the Monitor Type parameter. For more information, see Advanced settings in Monitoring Type.

5. Select or customize probe points.

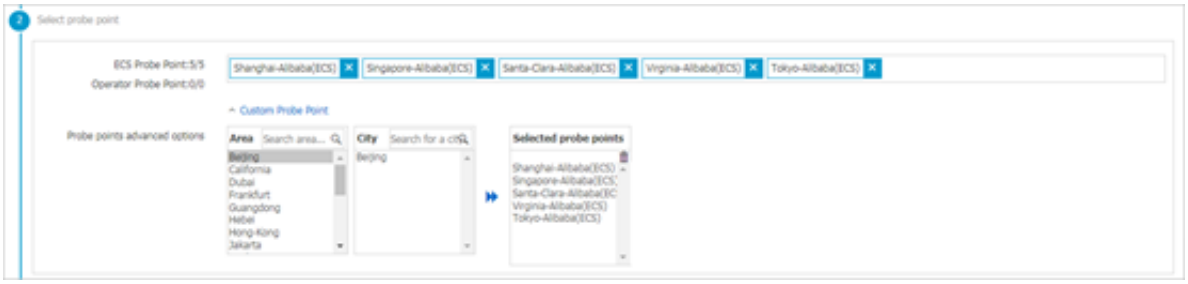


- ECS Probe Point: Select probe points.
- Probe points advanced options: Customize probe points.

6. Optional. Set alarm rules.

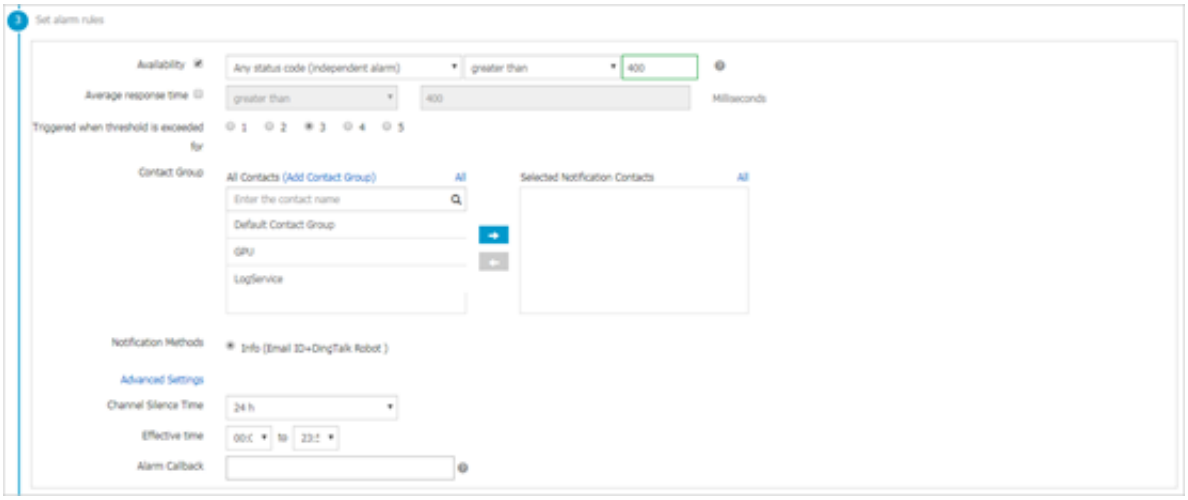


- Availability: Includes four options: Available probe point ratio, Number of probe points available, Any status code (independent alarm), and All status code (independent alarm). If the status code for a destination address in the detection results is greater than 399, the destination address is inaccessible. The number of probe points available equals the number of detection results with a status

code less than 400 within a monitoring period. The proportion of probe points available is calculated as follows:

Proportion of probe points available = (Number of detection results with a status code less than 400 within a monitoring period/Total number of detection results within the same monitoring period) × 100

- Average response time: the time taken by all the selected probe points to respond on average within a monitoring period.
- Triggered when threshold is exceeded for: the number of consecutive times that a metric exceeds its threshold before an alarm is reported. This parameter is used to detect the occasional volatility of the monitoring data.
- Contact Group: the group of contacts to receive alarm notifications.
- Notification Methods: the method for receiving alarm notifications.
- Advanced Settings: Includes three options: Channel Silence Time, Effective time, and Alarm Callback.
  - Channel Silence Time: the interval at which a notification is sent regularly before the reported alarm is cleared.
  - Effective time: the period of time during which the alarm rule is effective. CloudMonitor sends alarm notifications only within the specified period. After the specified period elapses, CloudMonitor records each reported alarm but does not send notifications.
  - Alarm Callback: Enter a URL that is accessible via the Internet. CloudMonitor sends POST requests that carry alarm information to this URL. You must enter a URL based on the HTTP protocol.

7. Click Create.

### Advanced settings in Monitoring Type

- Advanced settings for HTTP(s)

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Monitor Address | URL | Yes | We recommend that you include a schema into every entered address, for example, https://www.alibabacloud.com.<br>If you enter an address that does not contain a schema, CloudMonitor adds http as a schema to this address. |
| Request content | Form data or JSON object | No | If the request content is in the JSON format , ensure that the entered JSON objects are included in braces ({}. If you do not include JSON objects into braces ({}), CloudMonitor regards them to be form data. |
| Request method | A selected option | Yes | Valid values: GET \| POST \| HEAD.<br>Default value: GET. |
| Match response method | A selected option | Yes | When the match response content is specified, CloudMonitor reads the first 64 KB in the body of the response message that is sent from the HTTP server to search for the specified content. The result is one of the following:<br>1. The response contains the specified content.<br>2. The response does not contain the specified content.<br>CloudMonitor determines whether to trigger an alarm based on the specified match response method.<br>Alibaba Cloud probe points support match response content in English. |
| Match response content | Text | No | |

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| HTTP request header | Lines of text | No | The format of HTTP request header information is as follows: `key1 : value1` carriage return linefeed `key2 : value2`. CloudMonitor presets the following item in the request header:<br>Host: $ {Domain name specified in Monitor Address}<br>Pragma: no-cache<br>Cache-Control: no-cache<br>User-Agent: Chrome/57<br>Accept: */*<br>If the request content is a form, the request header may contain the following item:<br>Content-Type: application/x-www-form-urlencoded;charset=UTF-8<br>If your HTTP request header contains one or more of the preceding items, these items are overwritten by your settings.<br>According to the HTTP protocol, CloudMonit or converts the keys in the request header into an MIME header in the canonical format :<br>1. The first letter and the letter that follows a hyphen (-) in a key are capitalized. For example, `accept - encoding` is converted into `Accept - Encoding`.<br>2. If a key contains spaces or other invalid characters, it remains unchanged. |
| Cookie | N/A | No | Enter cookie text based on the HTTP protocol . |

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| HTTP Authentication Username | Username | No | This authentication refers to the basic authentication by the HTTP protocol. |
| HTTP Authentication Password | Password | No | |

- Advanced settings for PING

| Parameter | Value | Required |
| --- | --- | --- |
| Monitor Address | Domain name or IP address | Yes |
| Number of ping packets | Positive integer | Yes |

Note:

The Number of ping packets parameter indicates the number of times that the `ping` command is initiated. Value range: `1` to `40`. Default value: `20`.

- Advanced settings for TCP and UDP

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Monitor Address | Domain name or IP address | Yes | None |
| Request content format | A selected option | Yes | Valid when the request content is specified. Valid values: Hexadecimal Format \| Text. |

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Request content | Text or hexadecimal text | No | `Text` : a string of visible text characters. The text format does not support escape characters. That is, `\ n` is not converted into a new line entered, but rather the system regards it as two characters: a backward slash (\) and a letter *n*.<br>`Hexadecimal Format` : When the request content is a byte string that cannot be represented by a text string, you can convert the byte string into a hexadecimal string. The conversion rules are as follows: Each byte is converted into a 2-byte hexadecimal string. For example, (byte)1 is converted into a hexadecimal string 01, and (byte)27 is converted into a hexadecimal string 1B.<br>According to the conversion rules, the binary array "{(byte)1, (byte)27}" in Java format is converted into the following hexadecimal string: 011b or 011B. CloudMonitor does not distinguish between uppercase and lowercase letters for hexadecimal strings. Enter the string 011B in the Request content field and set Request content format to Hexadecimal Format. |
| match response content format | A selected option | Yes | Valid when the match response content is specified. Valid values: Hexadecimal Format \| Text. |
| Match response content | Text or hexadecimal text | No | For more information, see the Request content parameter. |

- Advanced settings for DNS

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Monitor Domain Name | Domain name | Yes | |
| DNS query type | A selected option | Yes | Valid values: A \| MX \| NS \| CNAME \| TXT \| ANY.<br>Default value: A. |
| DNS server | DNS IP address | No | If this parameter is unspecified, CloudMonit or uses the default DNS IP address. You can enter a domain name or an IP address. |
| Expected to resolve IP | Lines of text | No | Enter a list of domain names or IP addresses that you want to resolve. Each line represents a domain name or an IP address.<br>The detection is successful only when the specified list is a subset of the DNS list. |

- Advanced settings for POP3

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Monitor Address | URL | Yes | If you select the POP3(s) protocol, every address that you enter must contain a schema, for example, pop3s://pop3.aliyun.com.<br>If you enter an address that does not contain a schema, CloudMonitor adds pop3 as a schema to this address.<br>POP3(s) encrypts data by using TLS. |

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| username | Text | Yes | Your account is authenticated by using the `USER` and `PASS` commands.<br>Make sure that you enter a valid username and password. CloudMonitor detects the Internet at the intervals specified by the Monitoring frequency parameter. If the username and password are invalid, frequent detection to a party may cause this party to block your account. |
| Password | Text | Yes | |

- Advanced settings for SMTP

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Monitor Address | URL | Yes | Every address that you enter must contain a schema, for example, smtp://smtp.aliyun.com.<br>If you enter an address that does not contain a schema, CloudMonitor adds smtp as a schema to this address.<br>SMTP uses the `STARTTLS` command to negotiate with the server on encryption. When a secured connection is used, the authentication information is also encrypted. |
| username | Text | Yes | Your account is authenticated by using the `PLAIN` command.<br>Make sure that you enter a valid username and password. CloudMonitor detects the Internet at the intervals specified by the Monitoring frequency parameter. If the username and password are invalid, frequent detection to a party may cause this party to block your account. |
| Password | Text | Yes | |

- Advanced settings for FTP

| Parameter | Value | Required | Description |
| --- | --- | --- | --- |
| Monitor Address | URL | Yes | Example: ftp://smtp.aliyun.com. |
| Are you anonymous login | A selected option | Yes | Valid values: Anonymous Logon \| Authentication Required. Default value: Anonymous Logon.<br>If you select Authentication Required, you must enter a valid username and password. |
| username | Text | | The username and password used for FTP authentication.<br>If you select Anonymous Logon, the username and password are `anonymous` and `ftp @ example . com` , respectively. |
| Password | Text | | |

## 3.3 Manage a site monitoring task

This topic describes how to modify, delete, enable, and disable a site monitoring task.

### Modify a site monitoring task

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose New Site Monitor > Site Manage.
3. On the Site Monitoring page, click Modify in the Action column for the site monitoring task.
4. On the displayed page, modify the task settings and click Modify.

### Delete a site monitoring task

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose New Site Monitor > Site Manage.
3. On the Site Monitoring page, click Delete in the Action column for the site monitoring task.

Note:

After a site monitoring task is deleted, the related alarm rules are also deleted.

### Enable or disable a site monitoring task

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose New Site Monitor > Site Manage.
3. On the Site Monitoring page, click Enable or Disable in the Action column for the site monitoring task.

## 3.4 View site monitoring data

This topic describes how to view site monitoring data.

### Overview

This page provides data about access to the current site from four dimensions: availability, real-time response time, error distribution, and average response time.

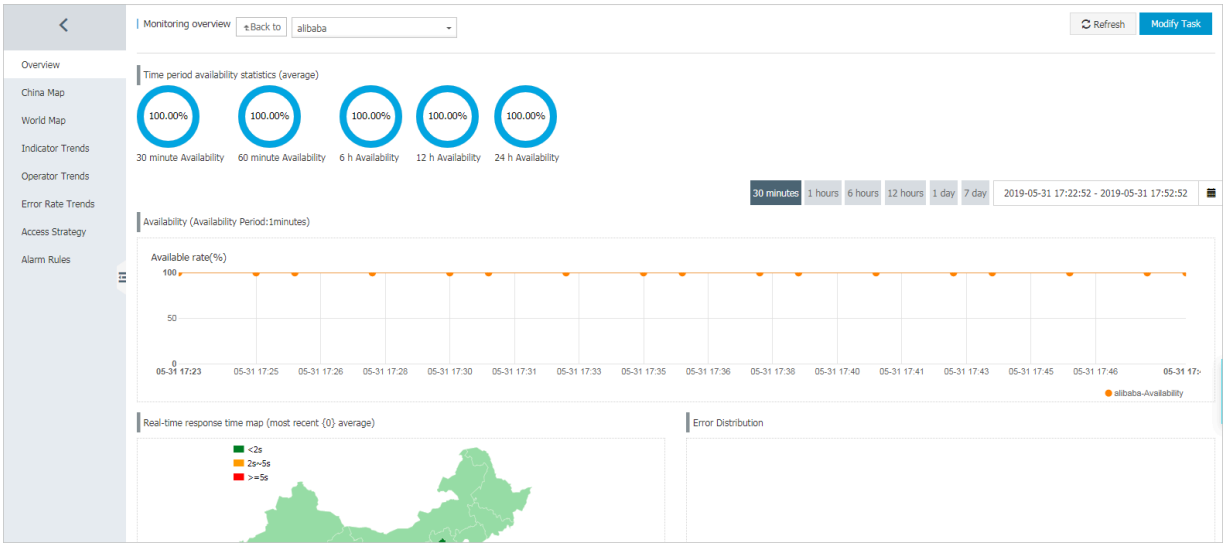


The error distribution shows the number of detection results with a status code greater than 399 for each carrier in each region within a specified period of time. You can click on the chart to view error details.

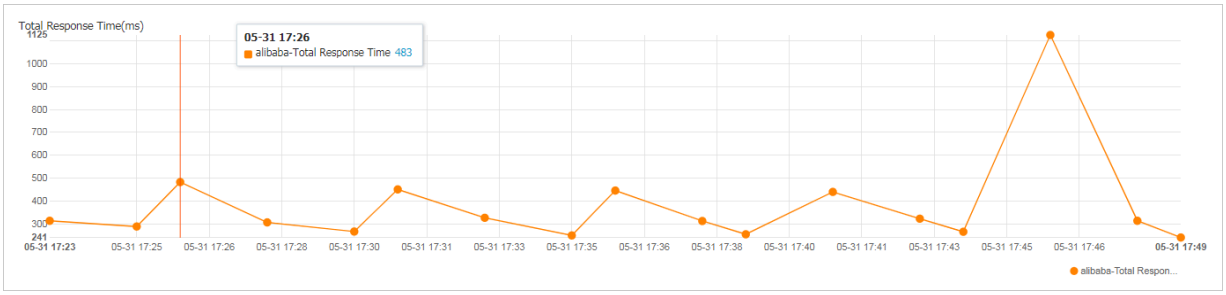


## China Map

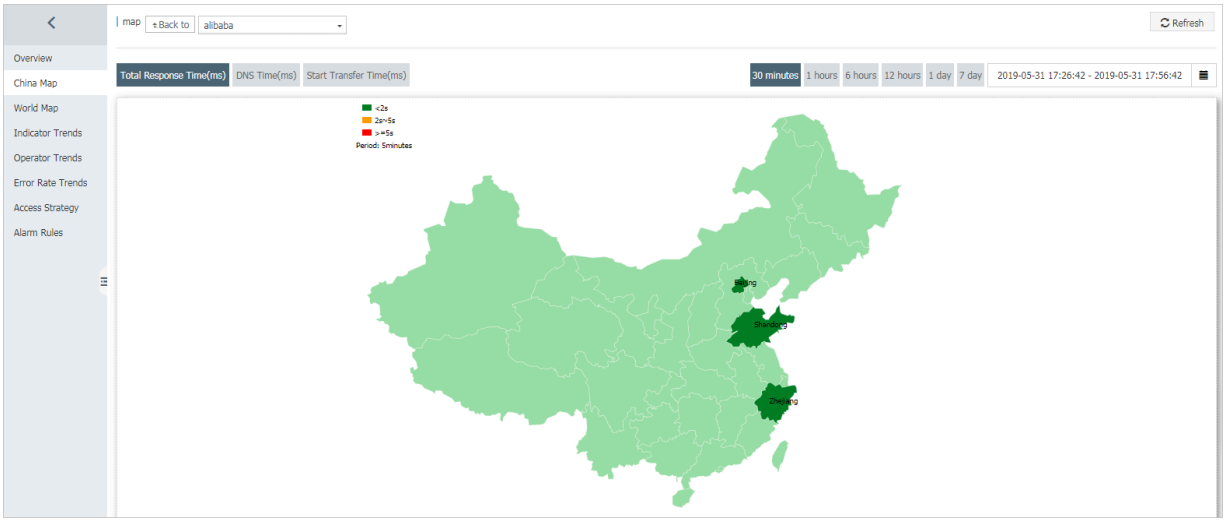


**If you click a province on the map, the second-level administrative divisions in the provide appear.**

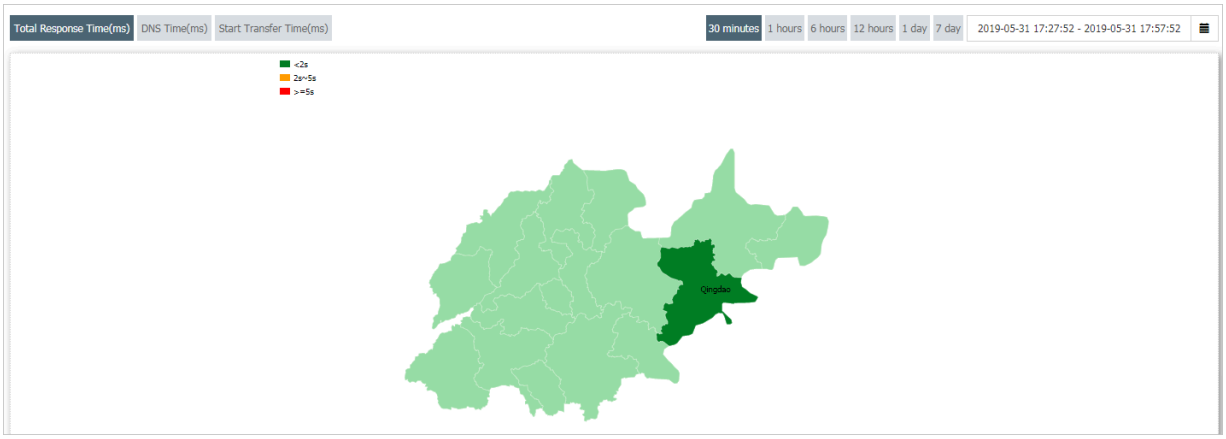


**Detailed monitoring data is shown below the map.**

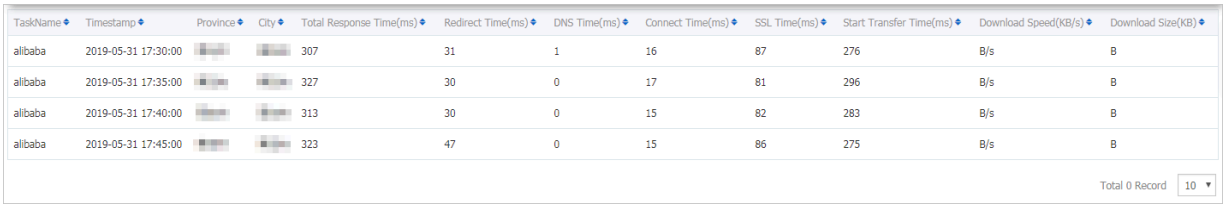

| TaskName | Timestamp | Province | City | Total Response Time(ms) | Redirect Time(ms) | DNS Time(ms) | Connect Time(ms) | SSL Time(ms) | Start Transfer Time(ms) | Download Speed(KB/s) | Download Size(KB) |
|---|---|---|---|---|---|---|---|---|---|---|---|
| alibaba | 2019-05-31 17:30:00 | [redacted] | [redacted] | 307 | 31 | 1 | 16 | 87 | 276 | B/s | B |
| alibaba | 2019-05-31 17:35:00 | [redacted] | [redacted] | 327 | 30 | 0 | 17 | 81 | 296 | B/s | B |
| alibaba | 2019-05-31 17:40:00 | [redacted] | [redacted] | 313 | 30 | 0 | 15 | 82 | 283 | B/s | B |
| alibaba | 2019-05-31 17:45:00 | [redacted] | [redacted] | 323 | 47 | 0 | 15 | 86 | 275 | B/s | B |

Total 0 Record 10

## World Map

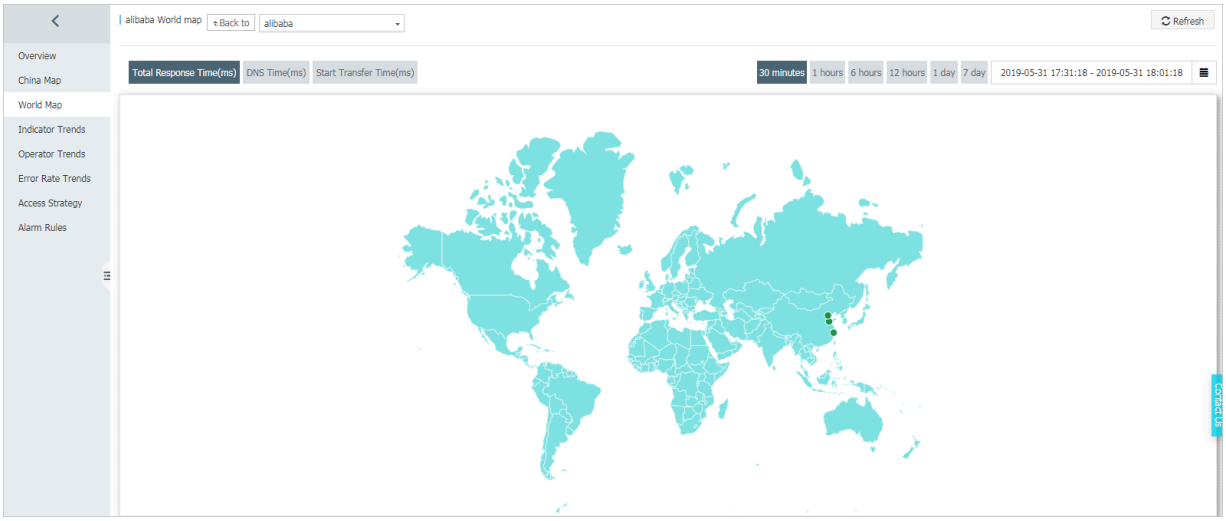


## Indicator Trends

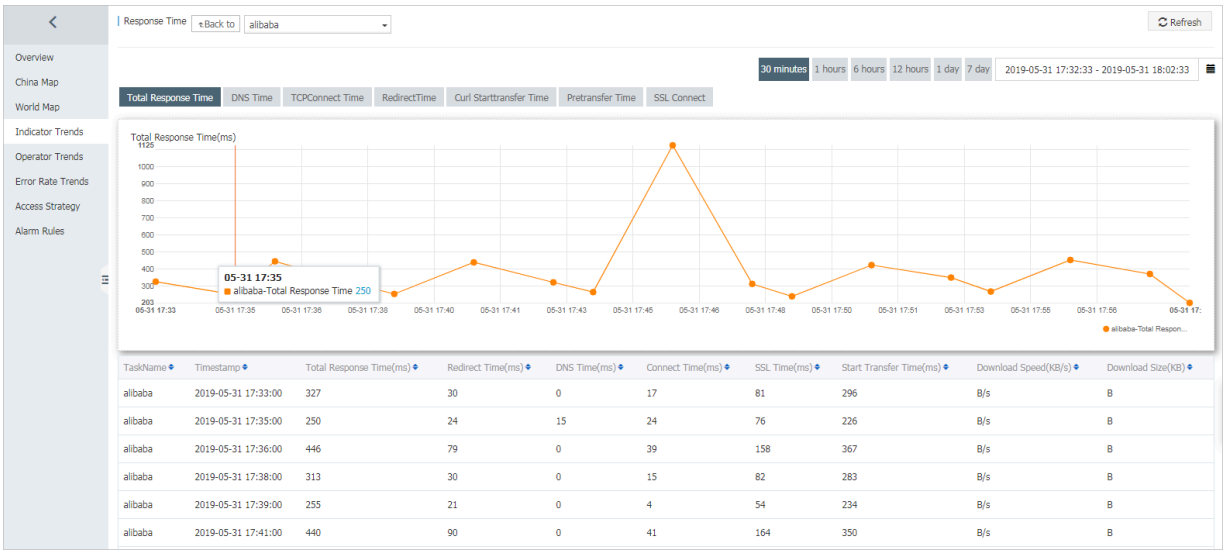


## Operator Trends

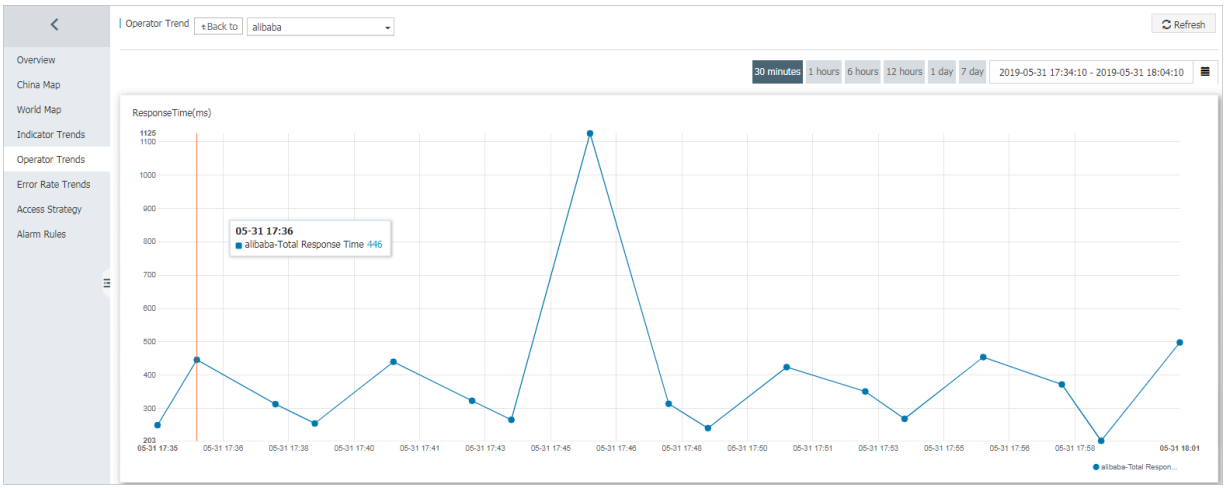

## Error Rate Trends

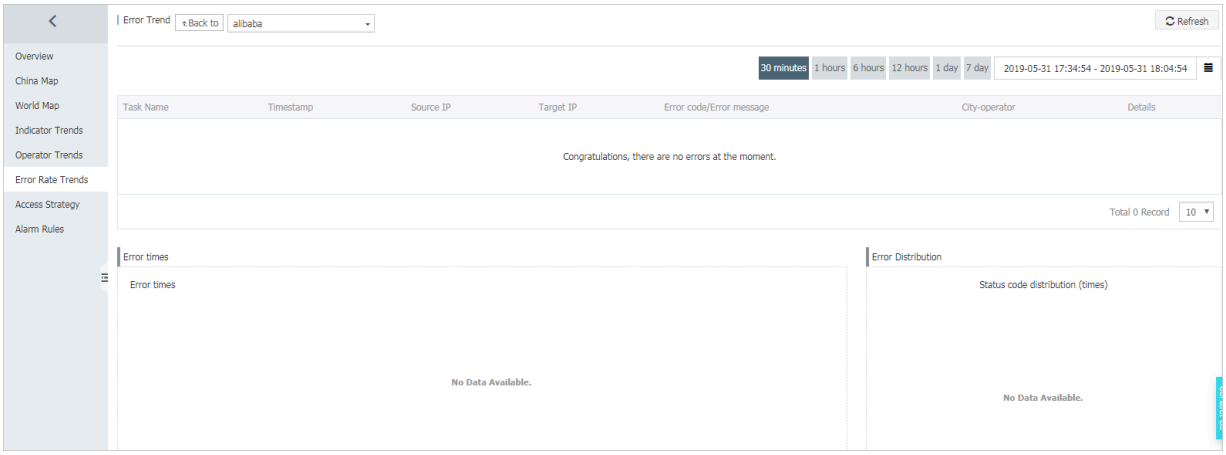


You can click More in the Details column for a task to view the error details for a specific carrier in a specific city.

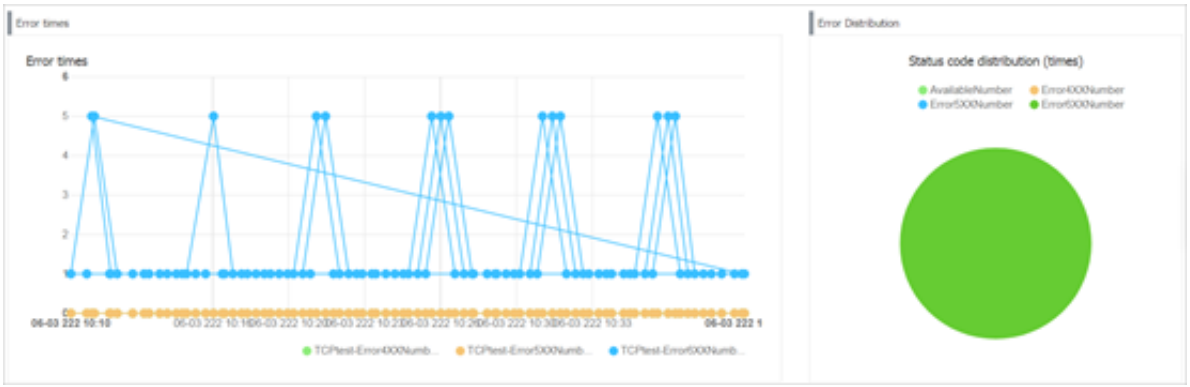


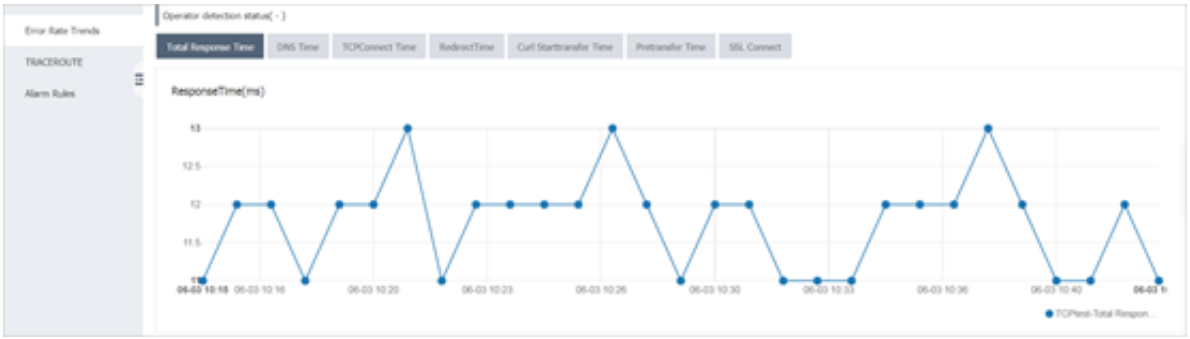


## Access Strategy

This page provides you with detailed detection results for each carrier in each region within a specified period of time.

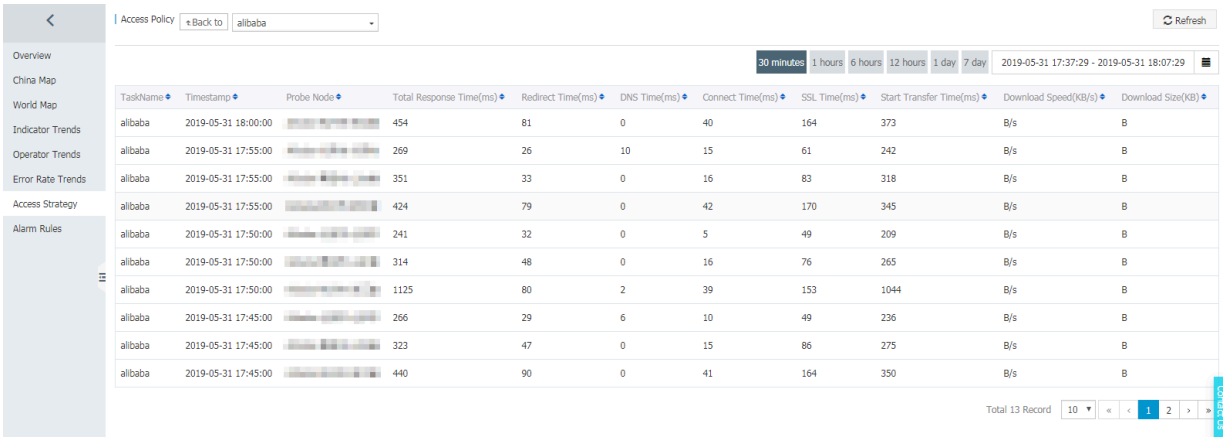


## Procedure

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose New Site Monitor > Site Manage.

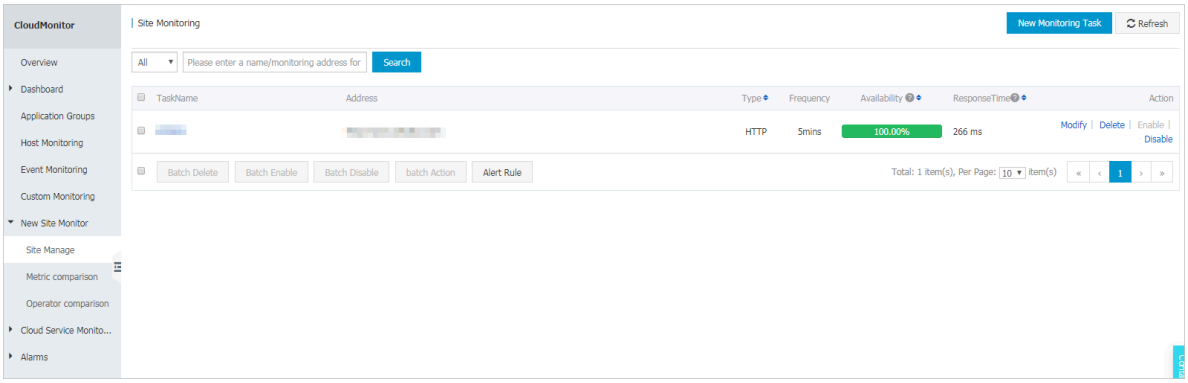


3. On the Site Monitoring page, click a site name to open the Monitoring overview page. In the left-side navigation pane, click a menu item and then in the main workspace view the corresponding monitoring data.

## Terms

| Term | Description |
| --- | --- |
| Availability | Number of detection results with a status code less than 400 from all probe points within a specified period of time/Total number of detection results × 100% |
| Total Response Time | The time taken to receive the first byte of an HTTP response after a probe point initiates detection. If the detection request is redirected, the time also includes the time spent to redirect the page. |

| Term | Description |
| --- | --- |
| DNS Time | The time for the Domain Name Server ( DNS) to resolve the domain name. Unit: millisecond. |
| TCPConnect Time | The time taken to write an HTTP request message after a probe point initiates detection. The time does not include the time for the DNS to resolve the domain name. |
| RedirectTime | The time taken to send the first detection request that is not redirected after a probe point initiates detection. |
| Start Transfer Time | The time taken to receive the first byte of an HTTP response after a probe point initiates detection. |
| Pretransfer Time | The time taken to write an HTTP request message after a probe point initiates detection. |
| SSL Connect | The time spent on SSL authentication after a probe point initiates detection. |
| Download Speed | The speed at which probe points read HTTP responses. |
| Download Size | The size of an HTTP response. If the HTTP response contains the `Content - Length` field, the download size equals the value of this field. If the HTTP response does not contain this field, the download size equals the number of bytes that are read from the HTTP response. |

## 3.5 Status code description

Each site monitoring protocol returns a status code during the detection process. Common status codes are described as follows.

Definitions of custom status codes of CloudMonitor

| Protocols | Status code | Description |
| --- | --- | --- |
| HTTP | 610 | Timeout (connection timeout, SSL Certificate exchange timeout, 30 s) |
| HTTP | 613 | DNS resolution Error |
| HTTP | 615 | The content does not match |
| HTTP | 616 | An error occurred while performing the authentication. |
| HTTP | 611 | Detection failure due to other reasons |
| HTTP | 617 | Maximum jump count exceeded<br>The Max number of 3xx Redirects allowed at the ECS probe point is 5 times<br>The maximum number of 3xx redirected jumps allowed by the carrier probe point is 2 |
| HTTP | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |
| Ping | 550 | Network is not available |

| Protocols | Status code | Description |
| --- | --- | --- |
| Ping | 610 | All sent packets receive no response in 2 seconds despite stable network condition. |
| Ping | 613 | \|Failed IP address resolution based on the host file\| |
| Ping | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |
| TCP | 550 | \|Failed to enable the socket. A typical cause is that the system sources have run out.\| |
| TCP | 610 | \|Failed response reception (time-out or no response)\| |
| TCP | 611 | Failed connection (time-out or rejected by the peer end)\| |
| TCP | 615 | The content does not match |
| TCP | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |
| UDP | 550 | \|Failed to enable the socket. A typical cause is that the system sources have run out.\| |
| UDP | 611 | Failed connection (failed resolution based on the host file)\| |

| Protocols | Status code | Description |
| --- | --- | --- |
| UDP | 610 | \|Failed sending or reception\| |
| UDP | 615 | The content does not match |
| UDP | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |
| DNS | 610 | DNS resolution failed |
| DNS | 613 | \|DNS query communicat ion error\| |
| DNS | 615 | The content does not match |
| DNS | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |
| SMTP | 610 | Connection time-out |
| SMTP | 611 | Your site could not be accessed successfully, failure Reasons include , but are not limited to, DNS resolution failure, Incorrect email format, failed to initialize SMTP client, and so on |
| SMTP | 616 | Login denied |
| SMTP | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |

| Protocols | Status code | Description |
| --- | --- | --- |
| POP3 | 611 | Unable to successfully access your site |
| POP3 | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |
| FTP | 610 | FTP Transfer failed |
| FTP | 611 | Failure caused by other factors, such as failure of DNS resolution, failure of TCP connection, etc. |
| FTP | 616 | A RAM user fails to log on to the console |
| FTP | 703 | The internal network detection of the server is prohibited. The internal network probe can use the availability monitoring function. |

Definitions of standard HTTP status codes

| Status code | Description | Note |
| --- | --- | --- |
| 200 | Request completed | Status Codes 2XX indicate that the service is normal. |
| 300 | Multiple choices | The server can perform multiple operations based on the request. The server selects one operation to perform based on the user agent, or provides a list of operations for the user agent to choose. |

| Status code | Description | Note |
| --- | --- | --- |
| 301 | Moved permanently | The requested webpage has been permanently moved to a new location. When the server returns Status Code 301 (in response to a GET or HEAD request), it automatically redirects the user agent to the new location. You must use this status code to notify Googlebot that a webpage or website has been permanentl y transferred to a new location. |
| 302 | Moved temporarily | The server returns the response from a webpage in a different location, but the user agent must use the original location for subsequent requests . Similar to Code 301 in response to a GET or HEAD request, this code means that the server automatica lly redirects the user agent to a different location. |
| 303 | See other | The server returns this code when the user agent must send GET requests separately for different locations for response retrieval. For all requests except HEAD requests, the server automatically jumps to other locations. |

| Status code | Description | Note |
| --- | --- | --- |
| 304 | Not modified | The requested webpage has not been modified since the last request. The webpage content is not returned when the server returns Status Code 304. |
| 305 | Use proxy | The user agent can access the requested webpage only by proxy. If the server returns this code, it also specifies the proxy that the user agent must use. |
| 400 | Bad request | The server does not understand the syntax of the request. |
| 401 | Unauthorized | Authentication is required for the request. The server may return Status Code 401 in response to the webpage access request after logon. |
| 403 | Forbidden | The server rejects the request. |
| 404 | Not found | The server cannot find the requested webpage. For example, if the requested webpage does not exist on the server, the server returns Status Code 404. |
| 405 | Method not allowed | The method specified in the request is forbidden. |
| 406 | Not acceptable | The content characteristics of the request cannot be used to respond to the webpage access request. |

| Status code | Description | Note |
| --- | --- | --- |
| 407 | Proxy authentication required | This status code is similar to 401 (unauthorized), but it specifies that the user agent must use a proxy for authentication. If the server returns this code, it also specifies the proxy that the user agent must use. |
| 408 | The request times out. | The server timed out waiting for the request. |
| 409 | Conflict | A conflict occurred when the server completed the request. The server must include the conflict information in the response packet. The server may return Status Code 409 and provide a list of differences between two conflicting requests when responding to the PUT request that conflicts with the previous request. |
| 411 | Length required | The server does not accept the request that contains header fields of invalid content length. |
| 412 | Precondition failed | The server does not meet one of the preconditions that the user agent sets in the request. |
| 413 | Request entity too large | The server cannot process the request because the request entity is too large and exceeds the server's processing capability. |

| Status code | Description | Note |
| --- | --- | --- |
| 414 | Requested URI too long | The server cannot process the request because the requested URI (usually the URL of the target website) is too long. |
| 415 | Unsupported media type | The request format is not supported by the requested webpage. |
| 416 | Range not satisfiable | The server returns this code if the request is out of the valid range of the requested webpage. |
| 417 | Expectation failed | The server does not meet the requirements of the ** Expect** request-header field. |
| 499 | Client closed request | This status code is returned when the client closes the connection because it takes a long time for the server to process the request. |
| 500 | Internal server error | The request cannot be completed due to a server error. |
| 501 | Not implemented | This code is returned when the server does not have the function to complete the request. For example, when the server cannot identify the request method, it may return this status code. |
| 502 | Bad gateway | The gateway or proxy server receives an invalid response from the upstream server. |

| Status code | Description | Note |
| --- | --- | --- |
| 503 | Service unavailable | The server is currently unavailable (due to overload or shutdown for maintenance). The unavailable state is temporary. |
| 504 | Gateway time-out | The gateway or proxy server failed to receive requests from the upstream server in time. |
| 505 | HTTP version not supported | The server does not support the HTTP version in the request. |

# 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 #unique_42.

## Create an alarm template

Note:

- Alarm temples 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.

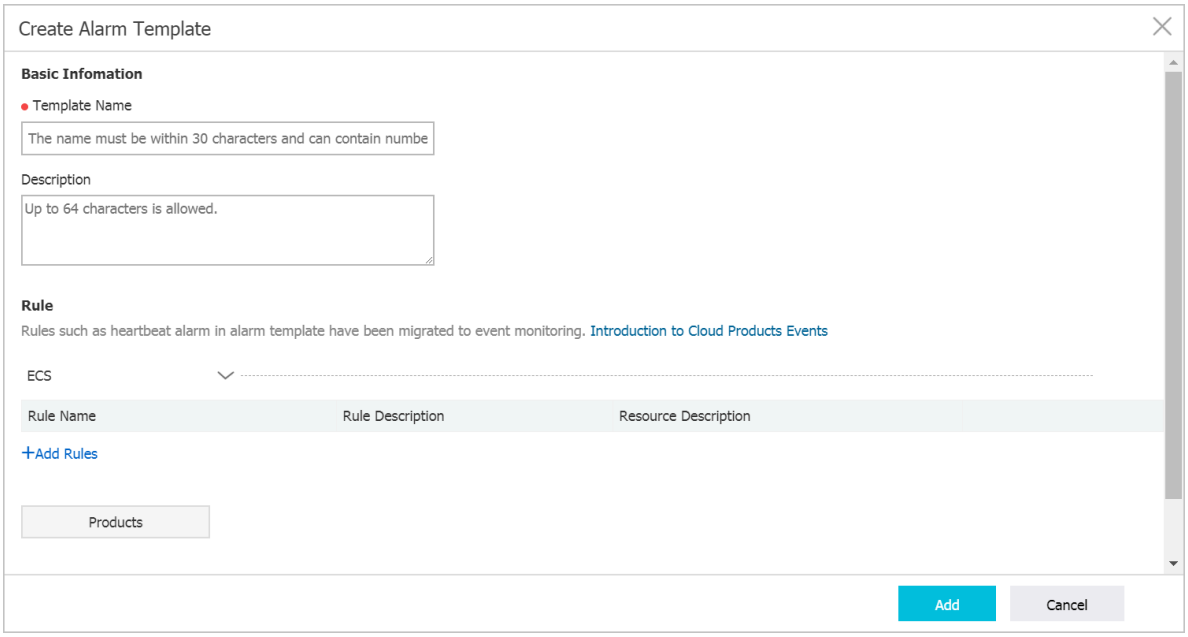


4. Enter a Template Name and Description in the Basic Information area.
5. Set an alarm rule. To add more alarm rules, click Add Rules.
6. Click Add.

### 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 alarm rule

This topic describes how to create a threshold alarm rule, so you can receive an alarm when a metric value reaches the specified threshold, and perform timely troubleshooting.

### Background

You can create threshold alarm rules to manage and monitor the usage and operation of cloud service resources. When a metric value reaches the specified threshold, you can receive an alarm. In this way, you can be informed of exceptions and handle them efficiently.

### Prerequisites

We recommend that you create an alarm contact and an alarm contact group before creating a threshold alarm rule. When you create an alarm rule, you can select the alarm contact group as the alarm receiver. For more information about how to create an alarm contact and an alarm contact group, see #unique_34.

If you want to use alarm callbacks in alarm rules, you must prepare a callback URL that is accessible on the Internet. In addition, you must enable the URL callback as a notification method in the existing operations and maintenance (O&M) or message notification system.

### Procedure

Note

CloudMonitor can notify you of alarms by means of phone calls, SMS messages, emails, TradeManager messages, or DingTalk ChatBot. If you want to receive alarms based on multiple methods, enter correct information when you configure alarm 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 page is displayed by default.

3. Click Create Alarm Rule to go to the Create Alarm Rule page.

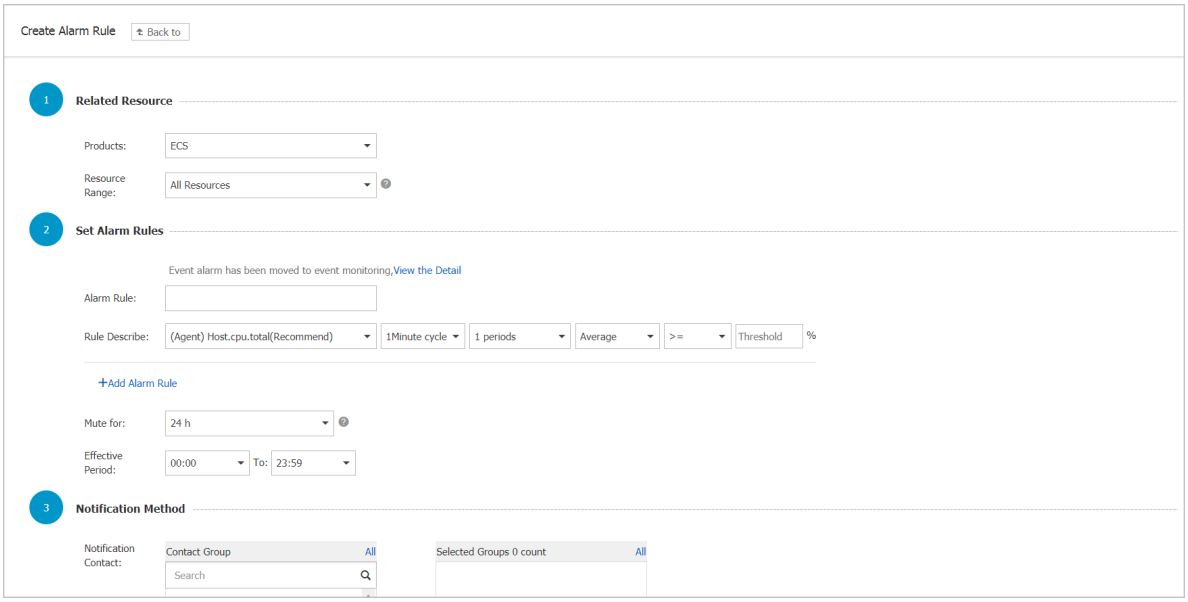


4. Select a resource range, set alarm rule parameters, select a notification method, and then click Confirm. For more information about alarm rule parameters, see #unique_45.

### References

- Create an alarm callback
- Use alarm templates
- Use the Initiative Alarm feature

## 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 #unique_34.

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.

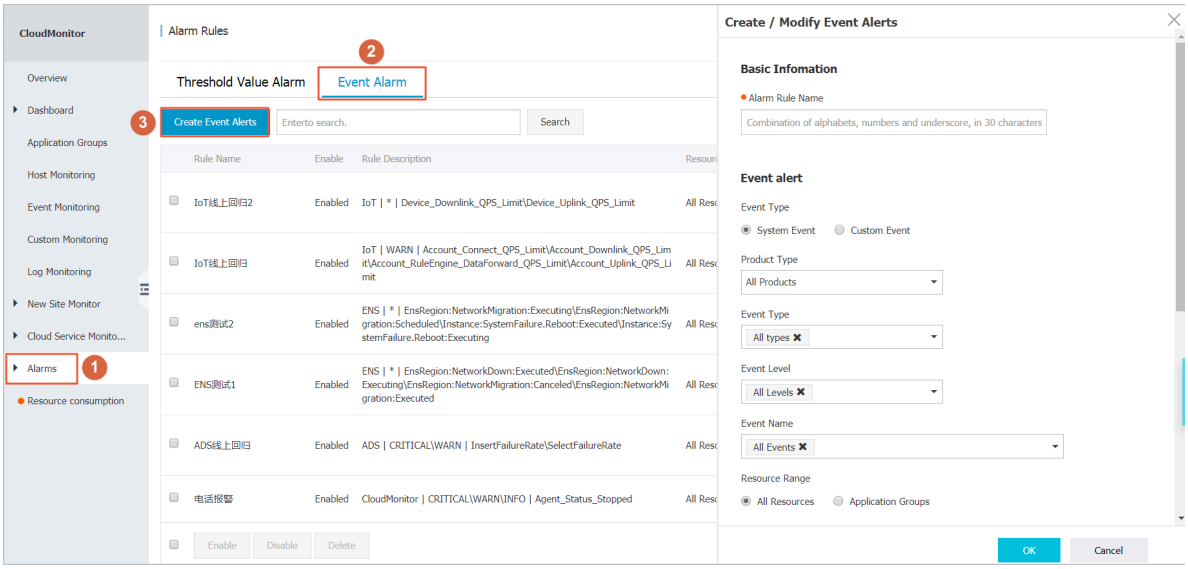


4. In the Basic Information section, enter the alert rule name.
5. Set Event Alarm Rule:
   a. If you set event type to System Event:
      - Product Type, Event Level, and Event Name: Set these parameters as needed.
      - Resource Range: If you select All Resources, notification is sent based on the 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.

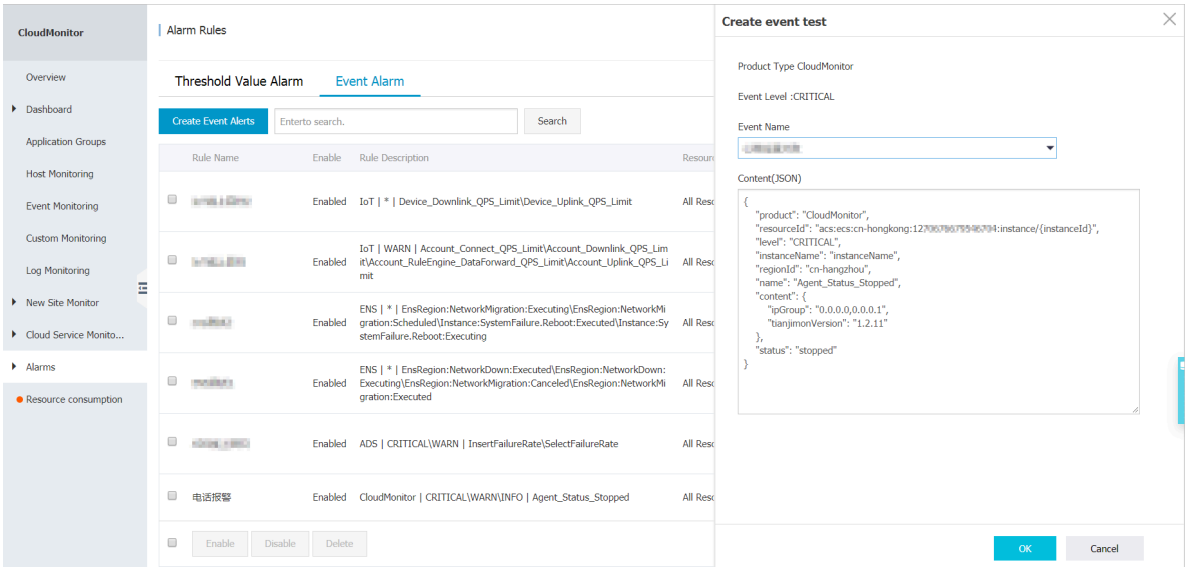


5. Select an event that you want to test. The event content is displayed. You can modify the fields such as instance ID in the content as needed.
6. Click OK. The system will send an event based on the content, triggering alert notification, MNS queue, Function Compute, or URL callback that you configure in the alert rule.

## 4.3.3 Alarm rule parameters

This topic describes parameters of a threshold alarm rule.

### Parameters

- Product: the monitored service, such as ECS, ApsaraDB for RDS, and Object Storage Service (OSS).
- Resource Range: the scope of the alarm rule. Valid values: All Resources and Instances.

> **Note:**
> If you set Resource Range to All Resources, the alarm rule is applicable to 1,000 instances or fewer. If the number of monitored resources is more than 1,000,

you may not receive alarms when the specified metric reaches the threshold. We recommend that you add resources to service-specific application groups before creating the alarm rule. To create a threshold alarm rule for a group, go to the Group Instances page, and click Threshold alarm.

  - All Resources: specifies that the alarm rule is applicable to all your instances of the specified service. The system sends alarm notifications if any metric of these instances reaches the specified threshold.
  - Instances: specifies that the alarm rule is applicable to a specified instance. The system sends alarm notifications if any metric of the instance reaches the specified threshold.

- Alarm Rule: the name of the alarm rule.
- Rule Description: the content of the alarm rule. This parameter defines the metric conditions that cause alarms.

  Alarm rule example: in host monitoring, a data point on the metric of a single host is reported at a 15-second interval. Therefore, 20 data points are reported in 5 minutes.

  - Average CPU usage in a 5-minute cycle greater than 90% in three consecutive cycles: specifies that the average value of the 20 data points on CPU usage reported in a 5-minute cycle is greater than 90% in three consecutive cycles. The system sends alarm notifications if the specified metric reaches the threshold.
  - CPU usage in a 5-minute cycle always greater than 90% in three consecutive cycles: specifies that the values of the 20 data points on CPU usage reported in a 5-minute cycle are greater than 90% in three consecutive cycles. The system sends alarm notifications if the specified metric reaches the threshold.
  - CPU usage in a 5-minute cycle greater than 90% for once in three consecutive cycles: specifies that the value of at least one of the 20 data points on CPU usage reported in a 5-minute cycle is greater than 90% in three consecutive cycles. The system sends alarm notifications if the specified metric reaches the threshold.
  - Total public network outbound traffic in a 5-minute cycle greater than 50 MB/s in three consecutive cycles: specifies that the sum of the values of the 20 data points on public network outbound traffic reported in a 5-minute cycle is greater than 50 MB/s in three consecutive cycles. The system sends alarm notifications if the specified metric reaches the threshold.

- Mute For: CloudMonitor sends an alarm notification only after detecting the specified exceptions consecutively for specified times. The minimum value is 5 minutes and the maximum value is 24 hours.
- Effective Period: the period when an alarm rule is effective. The system only sends alarm notifications within the effective period according to the alarm rule. The system only records alarms if the alarms occur during a non-effective period.
- Notification Contact: the contact group that CloudMonitor sends alarm notifications to.
- Alarm Levels: specifies the alarm severity level that corresponds to a specified notification method. Valid values: CRITICAL, WARN, and INFO.
  - INFO: sends alarm notifications by means of emails and DingTalk ChatBot.
- Auto Scaling: an alarm triggers the corresponding scaling rule after you select Auto Scaling and configure the rule.
- Email Remark: custom supplementary information of an alarm email. CloudMonitor sends the remarks along with the alarm email.
- HTTP CallBack: CloudMonitor uses a POST request to push an alarm to the public URL address you provided. This callback supports HTTP-based requests.

## 4.3.4 Manage alarm rules

CloudMonitor provides monitoring and alarms for your cloud services, and helps you timely locate exceptional metrics and efficiently perform troubleshooting.

You can start to manage alarm rules in the CloudMonitor console in three ways: in the left-side navigation pane, choose Application Groups to go to the Application Groups page, or choose a required monitoring type to go to the corresponding monitoring metrics page, or choose Alarms > Alarm Rules to go to the Alarm Rules page.

- Go to the Application Groups page to manage alarm rules.
- Go to the Host Monitoring page to manage alarm rules.
- Set alarm rules in Cloud Service Monitoring.
- Go to the Custom Monitoring page to set alarm rules.

# 4.3.5 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.

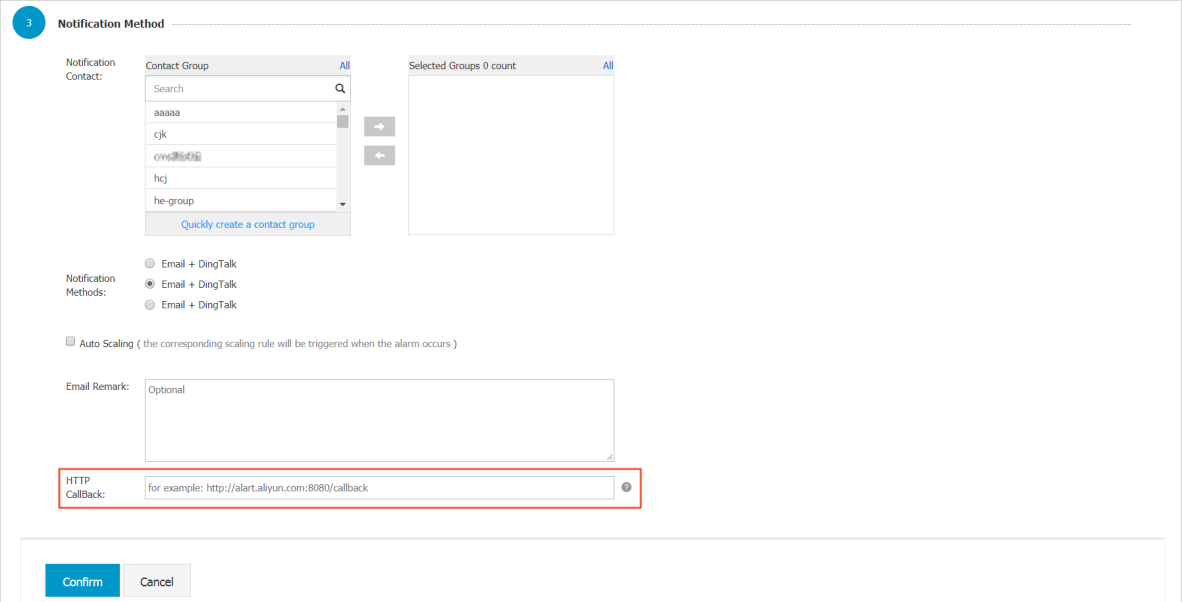


3. In the notification method section, enter the URL address for alert callback and click OK. When an alert rule is triggered, CloudMonitor sends an alert to your specified URL.

## 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 ":" putNewAlar  m_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 "
}
```

## 4.3.6 Write alarms to MNS

This topic describes how to write a threshold alarm to Message Service (MNS).

### Procedure

1. To authorize CloudMonitor to write an alarm to MNS, click here.

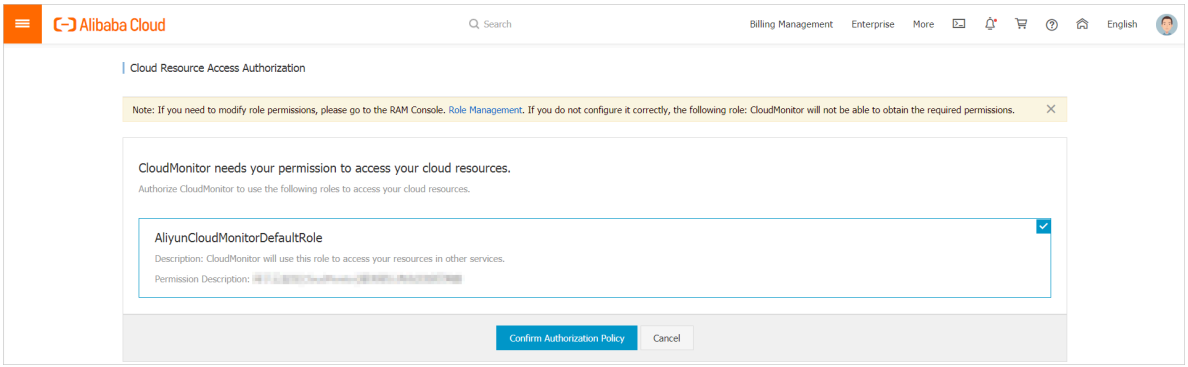


2. Start the OpenAPI Explorer service, and call the PutResourceMetricRule operation to create an alarm rule.
3. Call the PutMetricRuleTargets operation to create an alarm for the specified alarm rule, and set corresponding MNS parameters.

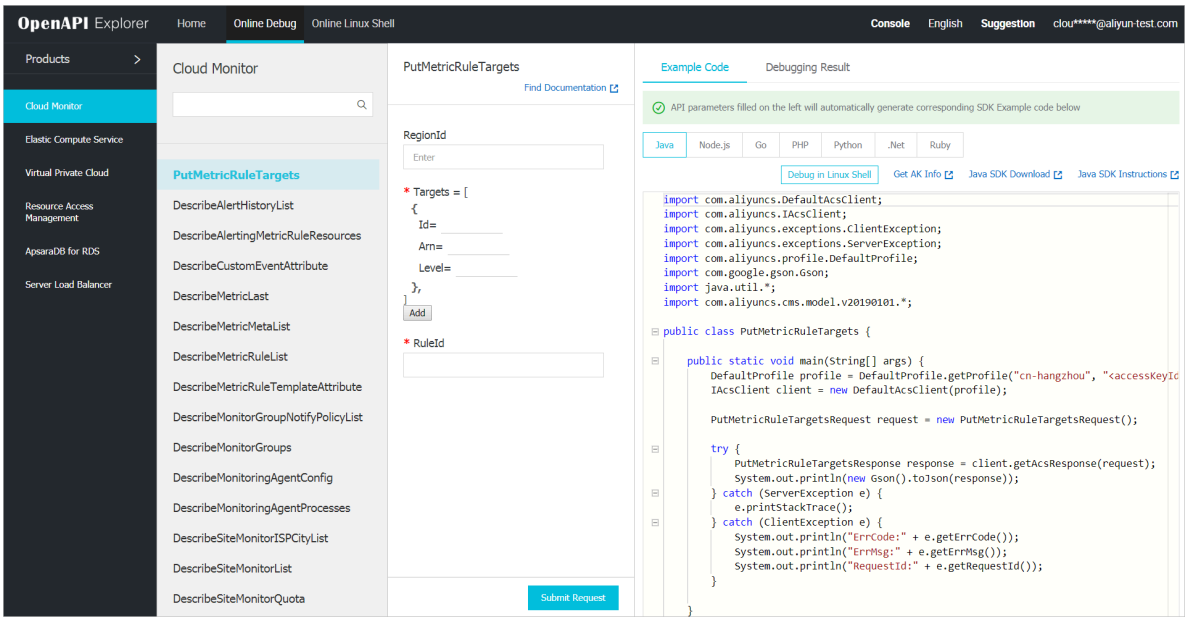


ARN: specifies the target MNS queue in the format of `" acs : mns :{$ RegionId }:{$ UserId }:/ queues /{$ queueName }/ messages "`, or specifies the target MNS topic in the format of `" acs : mns :{$ RegionId }:{$ UserId }:/ topics /{$ queueName }/ messages "`.

The following example shows parameters of PutMetricRuleTargets:

```
 RuleId :" db17 - 4afc - b11a - 568512d5a1  f9 ",
 Targets :[{
     Id :  1 ,
     Arn :" acs : mns :{$ RegionId }:{$ UserId }:/ queues /{$
 queueName }/ messages ",
     Level : [" INFO ", " WARN ", " CRITICAL "],
}]
```

### Message body written to MNS

CloudMonitor writes a message body to MNS in the JSON string format. When MNS consumes the message body, your client parses the message structure as a JSON string as follows:

```
{
  " ruleId ": " putNewAlar  m_group_77  8af9ba - a291 - 46ab - ac53 -
 3983bcee ****",
  " ruleName ": " test ",
  // Current   level .
  " curLevel ": " WARN ",
  // Previous   level .
  " preLevel ": " OK ",
  // The   instance   that   triggers   the   alarm .
  " resources ": "{\" instanceId \": \" i - uf61rfofjd  2iku7e ****
\"}",
  // The   condition   that   triggers   the   alarm .
  " escalation ": {
    " comparison  Operator ": " GreaterTha  nYesterday ",
    " level ":  3 ,
    " statistics ": " Average ",
    " tag ": " WARN ",
    " threshold ": " 0 ",
    " times ":  1
  },
  " metricData ": {
    " timestamp ":  1534736160  000 ,
    " userId ": " 1270676679  54 ****",
    " instanceId ": " i - uf61rfofjd  2iku7e ****",
    " Average ":  470687744 ,
    " Maximum ":  470794240 ,
    " Minimum ":  470556672 ,
    // Compare   some   metrics   with   those   in   the   previous
 month   and   those   in   the   same   period   of   the   previous
   year .-- Start .
    " AliyunCmsP  revValues ": { // Compared   values .
      " timestamp ":  1534649760  000 ,
      " userId ": " 1270676679  54 ****",
      " instanceId ": " i - uf61rfofjd  2iku7e ****",
      " Average ":  468463616 ,
      " Maximum ":  468549632 ,
      " Minimum ":  468258816
    },
    // Comparison   formula .
    " AliyunCmsC  omplexExpr  ession ": " 100 . 0  * ($ Average -$$
 prevAverag  e )/$$ prevAverag  e ",
    // Conversion   formula .
    " AliyunCmsC  omplexMath ": " 100 . 0  * ( 470687744 - 468463616
 )/ 468463616 ",
```

```
    // Calculatio  n   result .
    " AliyunCmsC  omplexValu  e ":  0 . 4747707023  6336133
    // Compare   some   metrics   with   those   in   the   previous
 month   and   those   in   the   same   period   last   year .-- End
 .
  },
  // Metric   parameters .
  " metricName ": " memory_act  ualusedspa  ce # 60 ",
  " namespace ": " acs_ecs_da  shboard ",
  " period ": " 60 ",

  // Applicatio  n   group   parameters .
  " groupBy ": " group ",
  " productGro  upName ": " RDS   instance   group ",
  " groupId ":" 44958 ",

  // Alarm   time
  " lastTime ":  327362743 , // The   duration   of   the   alarm .
  " time ":  1534736160  000 , // The   time   when   the   data
 occurred .

  " userId ": " 1736511134  38 ****",
  " eventName ": " AlertOk ",
  " eventType ": " Alert ",
  // Use   the   following   parameters   to   trace   the   alarm .
  " batchId ": " 4272653 - 152082 ****- 0 ",
  " version ": " 1 . 0 "
}
```

# 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 notificati on 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.

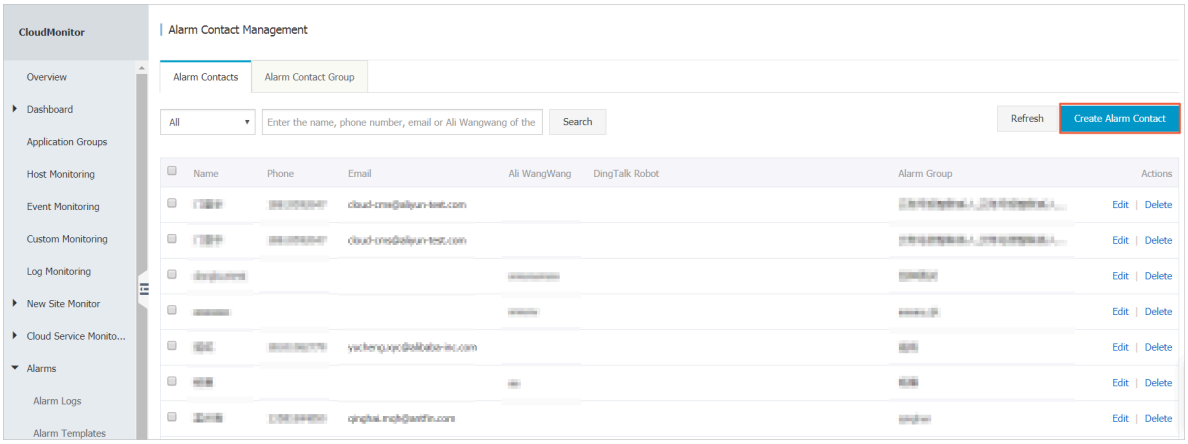

3. Click Create Alarm Contact in the upper-right corner. In the dialog box that appears, enter your name and email address.

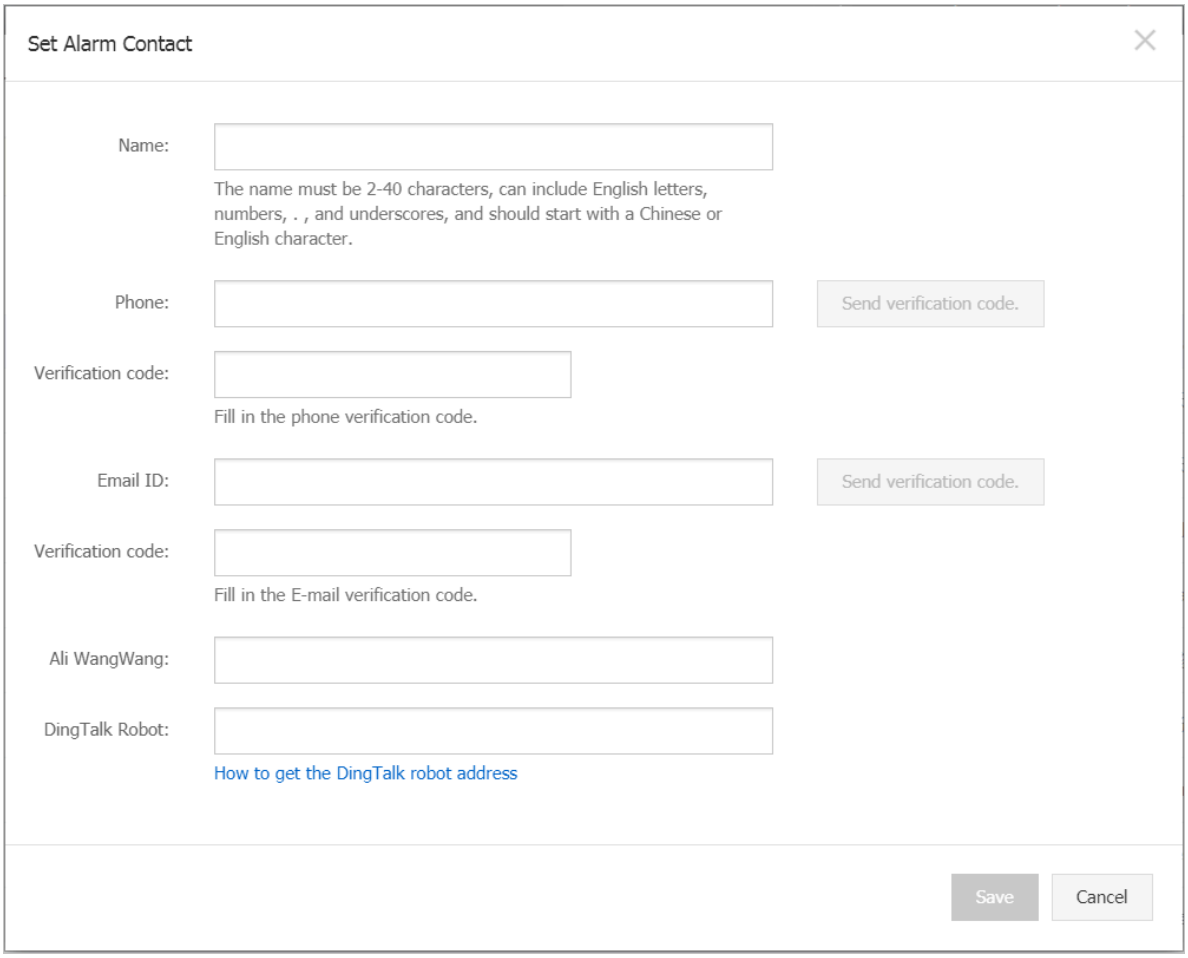


4. After the information is verified, click Save.

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.

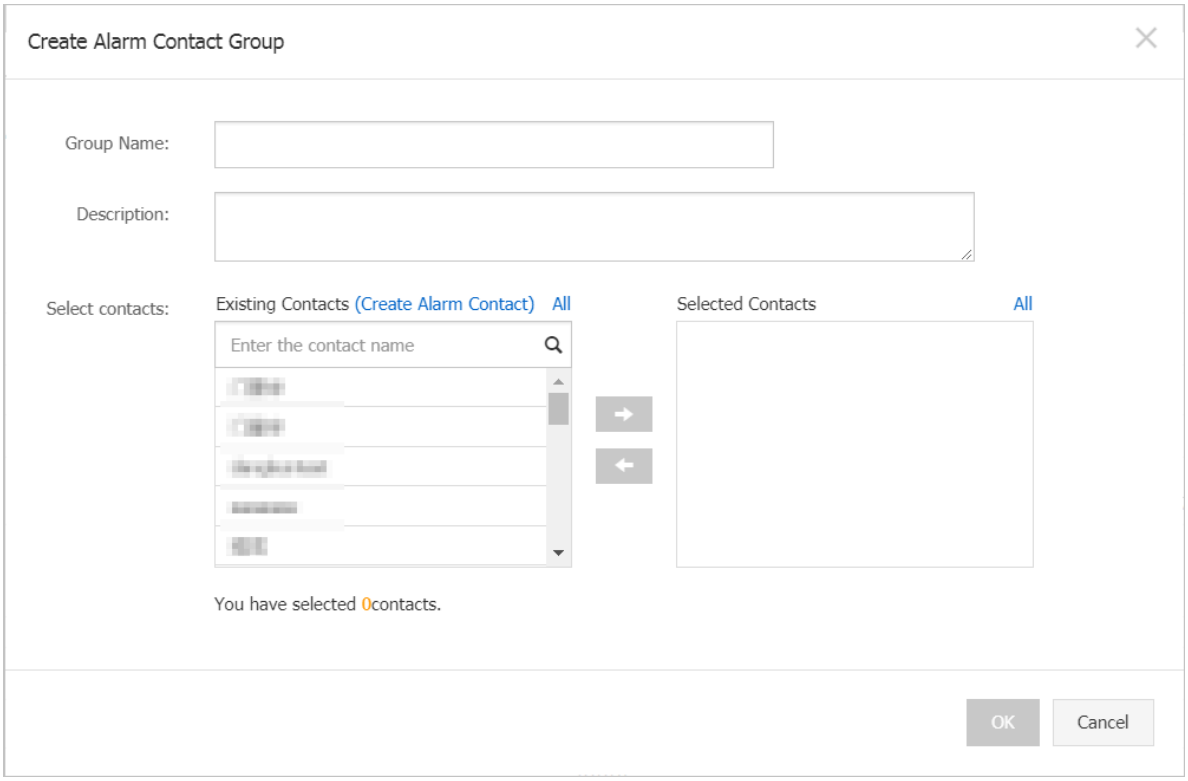


5. Enter the group name, select the contacts to be added to the group, and click OK.

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 View alarm logs

This topic describes how to view alarm logs.

You can search for alarm logs by rule name or group name in the CloudMonitor console.

## Procedure

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Alarms > Alarm Logs to open the Alarm Logs page.

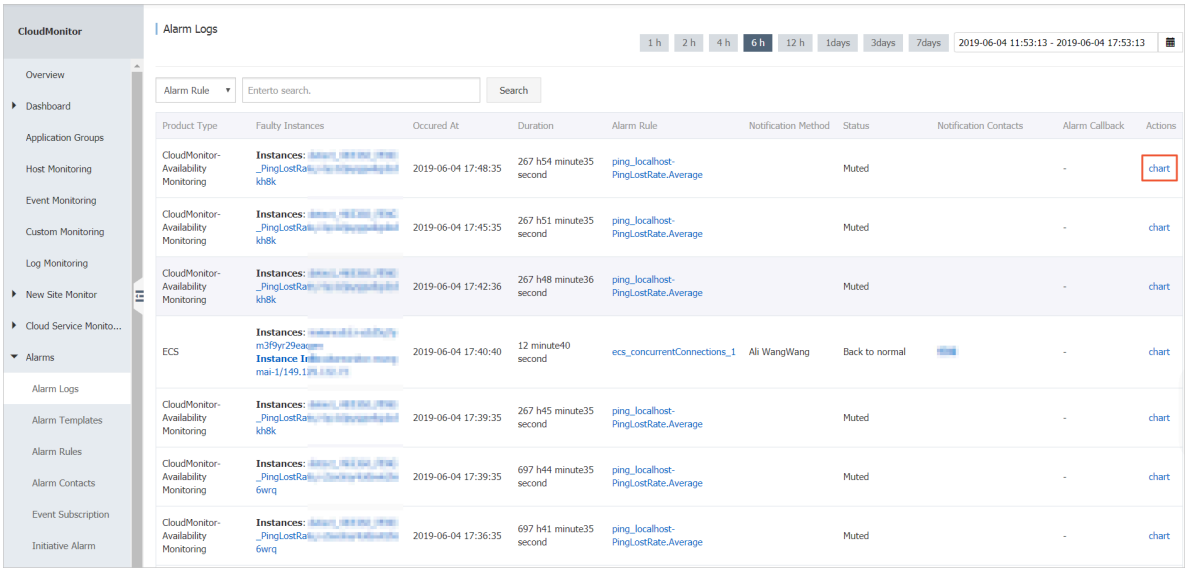


3. Select a search criterion (Alarm Rule or Group Name) from the drop-down list, enter a keyword in the search bar, and click Search.
4. Find the record that you want to view, and click chart in the Actions column.

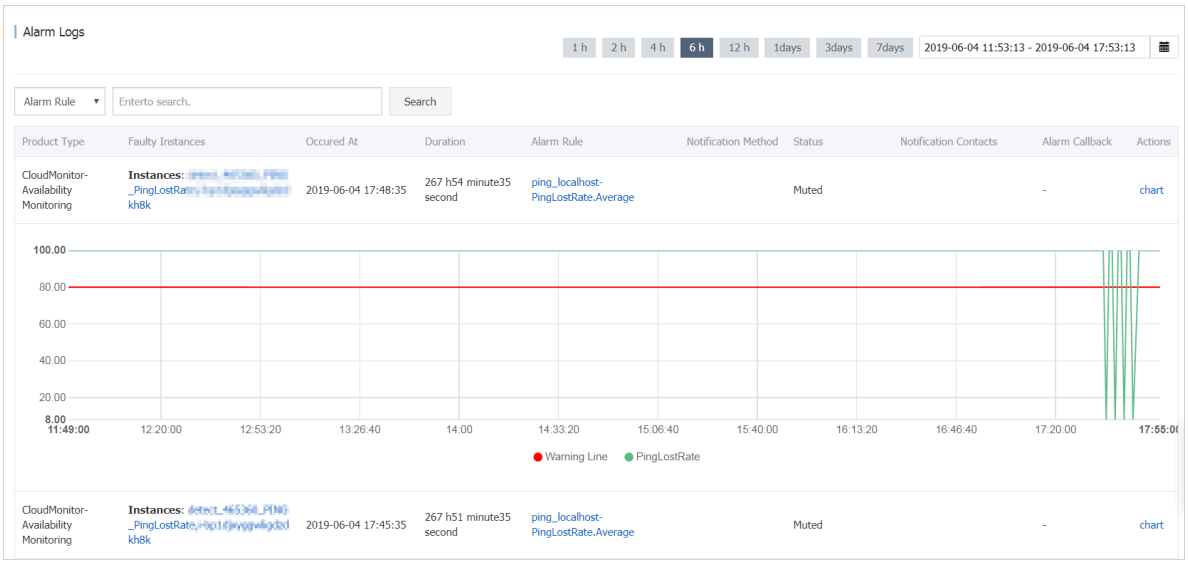

5. Select a time range within which you want to view alarm logs. You can only view the alarm logs that were generated within the last 31 days.

## 4.6 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 |

| Service name | Metric name | Rule description |
| --- | --- | --- |
| | InternetOutRate_Percent | Maximum value in 1 minute greater than 90%, five consecutive times, 1-hour mute duration, email notification |
| 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 |

| Service name | Metric name | Rule description |
| --- | --- | --- |
| | DropTrafficRX | Maximum value in 1 minute greater than 0, five consecutive times, 1-hour mute duration, email notification |
| | 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 |

| Service name | Metric name | Rule description |
| --- | --- | --- |
| | MemoryUtilization | Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification |
| | 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 |

| Service name | Metric name | Rule description |
| --- | --- | --- |
| | ShardingConnectionUtilization | Maximum value in 5 minutes greater than 80%, five consecutive times, 1-hour mute duration, email notification |
| 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 |

| Service name | Metric name | Rule description |
| --- | --- | --- |
| | NodeDiskUtilization | Maximum value in 1 minute greater than 75%, ten consecutive times, 1-hour mute duration, email notification |
| | 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.

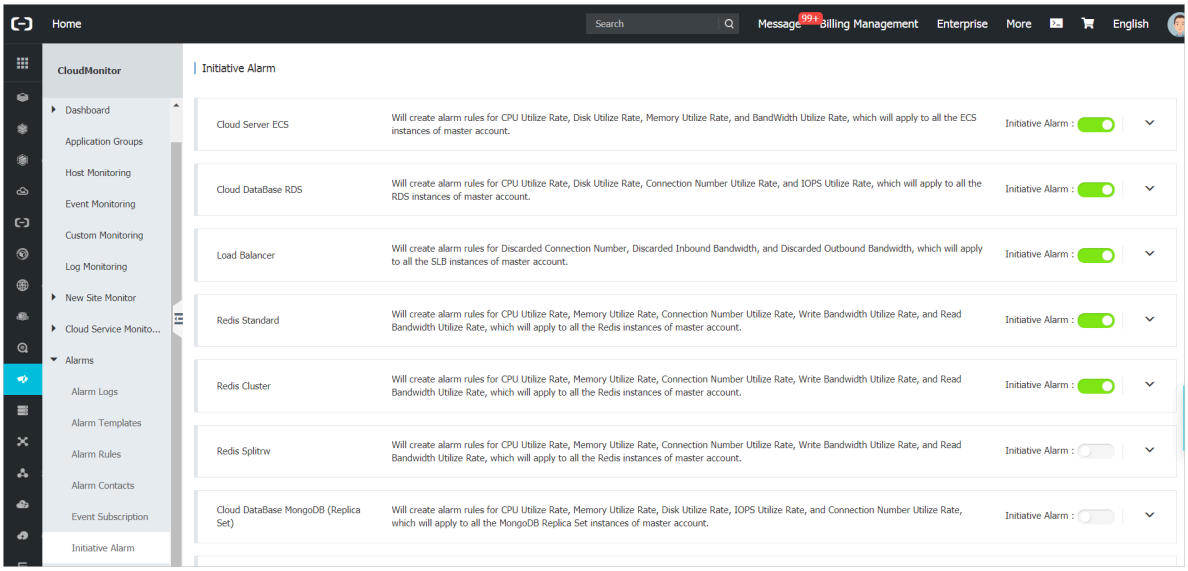


4. Click the drop-down arrow to the right of the One-click Alarm switch to view the alert rules that are automatically generated by CloudMonitor.

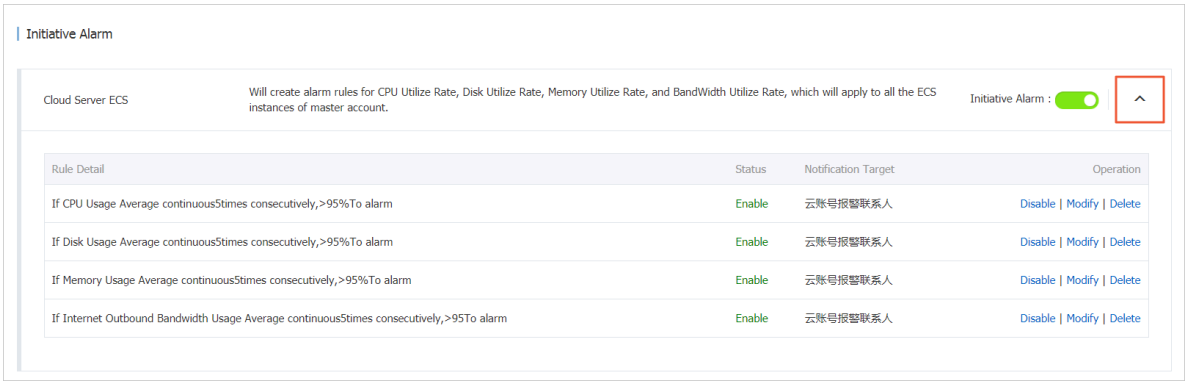


5. (Optional) You can click Disable, Modify, or Delete in the Actions column corresponding to an alert rule to disable, modify, or delete the rule.

# 5 Availability monitoring

## 5.1 Create an availability monitoring task

This topic describes how to configure availability monitoring so that you can receive alarms if a local service or a dependent remote service does not respond within a specified timeout period or returns an error status code.

### Background

Based on availability monitoring, CloudMonitor helps you quickly locate issues when a local service or a remote service has no response. CloudMonitor can send an alarm to you if the local service or the remote service fails to respond within a specified timeout period or returns an error status code. In this way, you can efficiently check response conditions of local or remote paths and ports.

### Prerequisites

- You have created a resource group for availability monitoring. For more information, see #unique_42.
- You have installed the CloudMonitor agent on the monitored host. For more information, see#unique_30.

### Procedure

#### Restrictions

Note:

- Availability monitoring depends on the CloudMonitor agent. Make sure that you have installed the CloudMonitor agent on the monitored host.
- CloudMonitor performs the availability detection once a minute.

#### Procedure

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, click Application Groups to go to the Application Groups page.
3. Click the name of the application group where you want to create an availability monitoring task to go to the Basic Information page of the application group.

4. In the left-side navigation pane, click Availability Monitoring to go to the Availability Monitoring page.

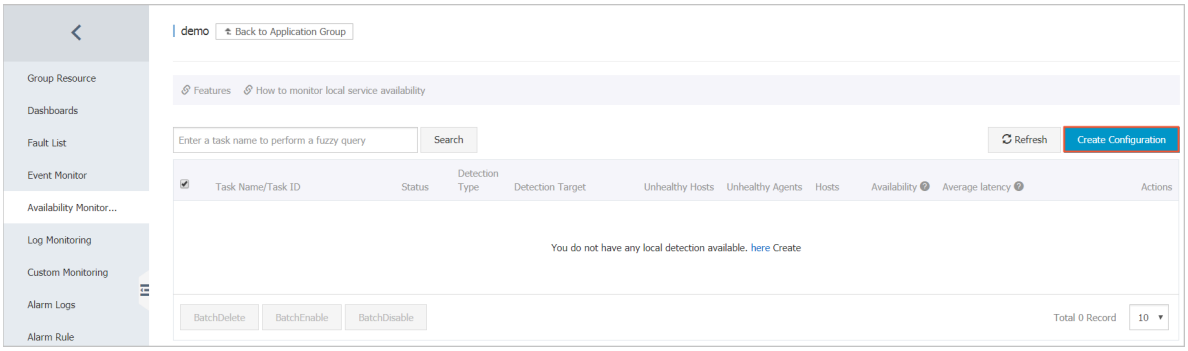

5. Click Create Configuration in the upper-right corner to go to the Create Availability Monitoring page.

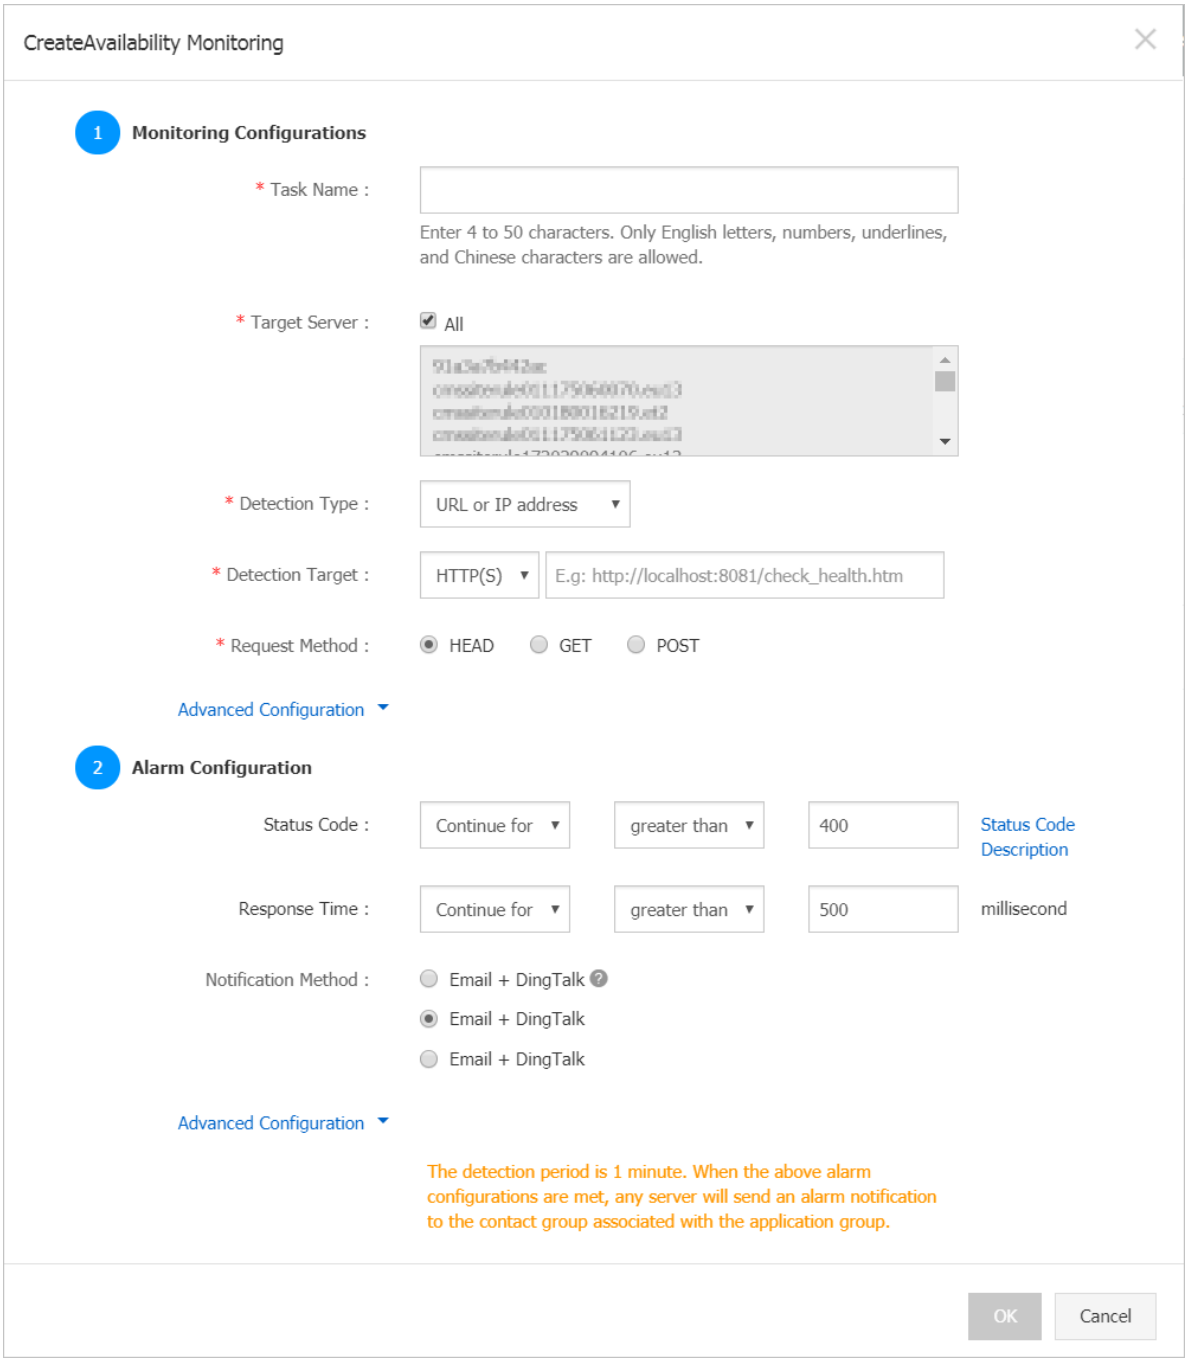


6. Set Task Name and Target Server. You can configure the same detection rule for all hosts in the group or some hosts in the group.
7. Set Detection Type to URL or IP Address, ApsaraDB for RDS, or ApsaraDB for Redis. Afterward, set Detection Target.
   - If you set Detection Type to URL or IP Address, 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 for Redis, you can view the instances in your group and their connection addresses.

8. In the Alarm Configuration field, set the Status Code and Response Time metrics. CloudMonitor generates an alarm if either of these metrics reaches the specified threshold. The system sends alarms to the contact group of the corresponding application group.
   - Status Code: the system generates an alarm if the local or remote service returns a status code as specified.
   - Response Time: the system generates an alarm if the local or remote service failed to respond within the specified timeout period.
   - Notification Method: the method that the system uses to send alarms.
   - Advanced Configuration: you can configure Mute For and Effective Period.
     - Mute For is a period when your alarm rules are muted so that the system does not send any alarms even when the local or remote service runs in the specified alarm conditions.
     - Effective Period: the time when an alarm rule takes effect. The system only sends alarms within the effective period according to the alarm rule. The system only records alarms if the alarms occur during a non-effective 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 applicatio n 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 #unique_30.
- Before you use local service availability monitoring, you must #unique_42.

### 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.

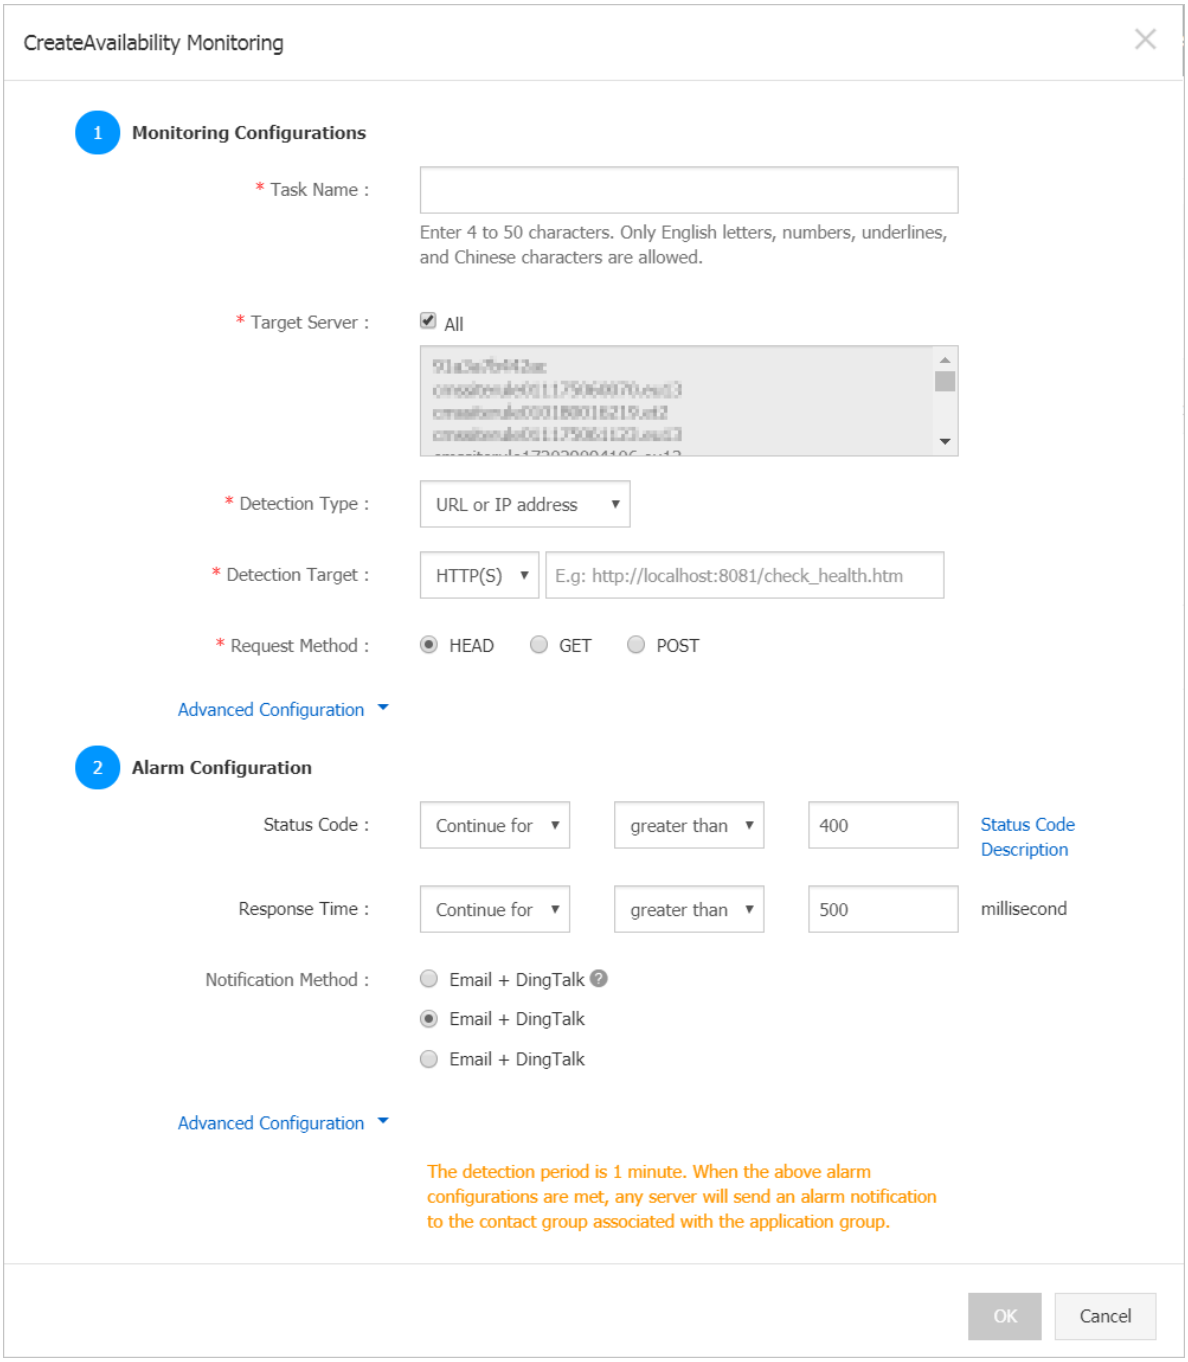


6. Set Task Name and Target Server. You can configure the same detection rule to be applied to all hosts in the group or just a portion of hosts in the group.
7. Set Detection Type to URL or IP Address, ApsaraDB for RDS, or ApsaraDB RDS for Redis. Set Detection Target.
8. In the Alert Configuration section, 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

CloudMonitor provides multiple metrics, such as the disk usage, input/output operations per second (IOPS) usage, connection usage, and CPU usage, to help you monitor the status of ApsaraDB for Relational Database Service (RDS). After you purchase RDS, CloudMonitor automatically collects data based on these metrics.

Note:

- Only primary and read-only instances in RDS support the monitoring and alerting services.
- By default, CloudMonitor creates alert rules for each primary instance and read-only instance. The alert threshold is 80% for the CPU usage, connection usage, IOPS usage, and disk usage. When the usage of a resource exceeds 80%, an SMS message and an email are sent to the specified contacts.

Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Disk usage | The percentage of the disk capacity used by the instance. | Instance | Percentage | 5 minutes |
| IOPS usage | The percentage of the IOPS used by the instance. | Instance | Percentage | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Connection usage | The percentage of instances that the current application connects to. The number of instances that an application can connect to is limited. This metric indicates the percentage of connected instances. | Instance | Percentage | 5 minutes |
| CPU usage | The percentage of the CPU capacity used by the instance. The CPU usage is determined by the database memory size. | Instance | Percentage | 5 minutes |
| Memory usage | The percentage of the memory used by the instance. Currently, only MySQL databases support this metric. | Instance | Percentage | 5 minutes |
| Read-only instance latency | The latency of the MySQL read-only instance. | Instance | Seconds | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Inbound traffic | The inbound traffic per second to the instance. | Instance | bit/s | 5 minutes |
| Outbound traffic | The outbound traffic per second from the instance. | Instance | bit/s | 5 minutes |
| Instance failure | An event-type metric. You can set alert rules for this metric. | - | - | - |
| Instance failover | An event-type metric. You can set alert rules for this metric. | - | - | - |

The metrics of inbound and outbound traffic support only 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. The ApsaraDB for RDS page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

### Alerting service

- Set alert rules
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for RDS. The ApsaraDB for RDS page appears.
  3. Click Alarm Rules in the Actions column for an instance to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.2 SLB

CloudMonitor provides multiple metrics, such as the inbound traffic and outbound traffic, to help you to monitor the status of Server Load Balancer (SLB). CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs. After you create an SLB instance, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics
  - Layer-4 protocol metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Inbound traffic on a port | The traffic consumed for accessing the port from the Internet. | Port | bit/s | 1 minute |
| Outbound traffic on a port | The traffic consumed for accessing the Internet from the port. | Port | bit/s | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Received packets on a port | The number of packets received on the port per second. | Port | Count/second | 1 minute |
| Transmitted packets on a port | The number of packets transmitted on the port per second. | Port | Count/second | 1 minute |
| New connections on a port | The average number of times that the status is SYN_SENT at first for a TCP three-way handshake in the monitoring period. | Port | Count | 1 minute |
| Active connections on a port | The number of connections in the ESTABLISHED status on the port in the monitoring period. | Port | Count | 1 minute |
| Inactive connections on a port | The number of connections in statuses other than ESTABLISHED on the port in the monitoring period. | Port | Count | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Concurrent connections on a port | The total number of connections on the port (including both active and inactive connections) in the monitoring period. | Port | Count | 1 minute |
| Healthy backend Elastic Compute Service (ECS) instances | The number of instances that pass the health test. | Port | Count | 1 minute |
| Faulty backend ECS instances | The number of instances that fail the health test. | Port | Count | 1 minute |
| Connections discarded on a port | The average number of connections discarded on the port per second. | Port | Count/second | 1 minute |
| Received packets discarded on a port | The average number of received packets discarded on the port per second. | Port | Count/second | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Transmitted packets discarded on a port | The average number of transmitted packets discarded on the port per second. | Port | Count/second | 1 minute |
| Inbound traffic discarded on a port | The average inbound traffic discarded on the port per second. | Port | bit/s | 1 minute |
| Outbound traffic discarded on a port | The average outbound traffic discarded on the port per second. | Port | bit/s | 1 minute |
| Active connections on an instance | The number of connections in the ESTABLISHED status on the instance in the monitoring period. | Instance | Count/second | 1 minute |
| Inactive connections on an instance | The number of connections in statuses other than ESTABLISHED on the instance in the monitoring period. | Instance | Count/second | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Connections discarded on an instance | The number of connections discarded on the instance per second. | Instance | Count/second | 1 minute |
| Received packets discarded on an instance | The number of received packets discarded on the instance per second. | Instance | Count/second | 1 minute |
| Transmitted packets discarded on an instance | The number of transmitted packets discarded on the instance per second. | Instance | Count/second | 1 minute |
| Inbound traffic discarded on an instance | The amount of inbound traffic discarded on the instance per second. | Instance | bit/s | 1 minute |
| Outbound traffic discarded on an instance | The amount of outbound traffic discarded on the instance per second. | Instance | bit/s | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Concurrent connections on an instance | The total number of connections on the instance (including both active and inactive connections) in the monitoring period. | Instance | Count/second | 1 minute |
| New connections on an instance | The average number of times that the status is SYN_SENT at first for a TCP three-way handshake in the monitoring period. | Instance | Count/second | 1 minute |
| Received packets on an instance | The number of packets received on the instance per second. | Instance | Count/second | 1 minute |
| Transmitted packets on an instance | The number of packets transmitted on the instance per second. | Instance | Count/second | 1 minute |
| Inbound traffic on an instance | The traffic consumed for accessing the instance from the Internet. | Instance | bit/s | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Outbound traffic on an instance | The traffic consumed for accessing the Internet from the instance. | Instance | bit/s | 1 minute |

- Layer-7 protocol metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| QPS on a port | The QPS on the port. | Port | Count/second | 1 minute |
| Response time (RT) on a port | The average response time to requests on the port. | Port | Milliseconds | 1 minute |
| Status codes 2xx on a port | The number of status codes 2xx returned by SLB to the client on the port. | Port | Count/second | 1 minute |
| Status codes 3xx on a port | The number of status codes 3xx returned by SLB to the client on the port. | Port | Count/second | 1 minute |
| Status codes 4xx on a port | The number of status codes 4xx returned by SLB to the client on the port. | Port | Count/second | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Status codes 5xx on a port | The number of status codes 5xx returned by SLB to the client on the port. | Port | Count/second | 1 minute |
| Other status codes on a port | The number of other status codes returned by SLB to the client on the port. | Port | Count/second | 1 minute |
| Upstream status codes 4xx on a port | The number of status codes 4xx returned by RS to SLB on the port. | Port | Count/second | 1 minute |
| Upstream status codes 5xx on a port | The number of status codes 5xx returned by RS to the client on the port. | Port | Count/second | 1 minute |
| Upstream RT on a port | The average response time to requests from RS to the proxy on the port. | Port | Milliseconds | 1 minute |
| QPS on an instance | The QPS on the instance. | Instance | Count/second | 1 minute |
| RT on an instance | The average response time to requests on an instance. | Instance | Count/second | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Status codes 2xx on an instance | The number of status codes 2xx returned by SLB to the client on the instance. | Instance | Count/second | 1 minute |
| Status codes 3xx on an instance | The number of status codes 3xx returned by SLB to the client on the instance. | Instance | Count/second | 1 minute |
| Status codes 4xx on an instance | The number of status codes 4xx returned by SLB to the client on the instance. | Instance | Count/second | 1 minute |
| Status codes 5xx on an instance | The number of status codes 5xx returned by SLB to the client on the instance. | Instance | Count/second | 1 minute |
| Other status codes on an instance | The number of other status codes returned by SLB to the client on the instance. | Instance | Count/second | 1 minute |
| Upstream status codes 4xx on the instance | The number of status codes 4xx returned by RS to SLB on the instance. | Instance | Count/second | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Upstream status codes 5xx on an instance | The number of status codes 5xx returned by RS to SLB on the instance. | Instance | Count/second | 1 minute |
| Upstream RT on an instance | The average response time to requests from RS to the proxy on the instance. | Instance | Milliseconds | 1 minute |

 Note:

In the preceding table, new connections, active connections, and inactive connections refer to TCP connection requests sent from clients 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. The Server Load Balancer page appears.
3. Click a region. All instances in the region appear in the instance list.
4. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Server Load Balancer. The Server Load Balancer page appears.
  3. Click a region. All instances in the region appear in the instance list.
  4. Click Alarm Rules in the Actions column for an instance to view the alert rules.
  5. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the alert method, and then click Confirm.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 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 #unique_71.
- For more information about the alarm service for OSS monitoring, see #unique_72.

## 6.4 CDN

CloudMonitor provides multiple metrics, such as the queries per second (QPS), bandwidth, and byte hit ratio, to help you monitor the status of Content Delivery Network (CDN). After you add a CDN domain name, CloudMonitor automatically

collects data based on these metrics. You can log on to the CloudMonitor console and view the monitoring details on the CDN monitoring page. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

### Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Visits per second | The number of visits in the monitoring period divided by the monitoring period. | Domain name | QPS | 1 minute |
| Bandwidth | The maximum traffic per unit time. | Domain name | bit/s | 1 minute |
| Hit ratio | The probability that request bytes are found in the cache in the monitoring period. The number of request bytes is the number of requests multiplied by traffic. The byte hit ratio reflects the back-to-origin traffic. | Domain name | Percentage | 1 minute |
| Outbound traffic to the Internet | The traffic from CDN to the Internet. | Domain name | Bytes | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Percentage of status codes 4xx | The percentage of HTTP status codes 4xx to all the returned HTTP status codes in the monitoring period. | Domain name | Percentage | 1 minute |
| Percentage of status codes 5xx | The percentage of HTTP status codes 5xx to all the returned HTTP status codes in the monitoring period. | Domain name | Percentage | 1 minute |

- View monitoring data

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Alibaba Cloud CDN. The CDN page appears.
  3. Click the Domain Name List tab.
  4. Click a domain name or click Monitoring Charts in the Actions column for a domain name to view the monitoring charts.
  5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Alibaba Cloud CDN. The CDN page appears.
  3. Click the Domain Name List tab.
  4. Click Alarm Rules in the Actions column for a domain name to view the alert rules.
  5. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.5 EIP

CloudMonitor provides multiple metrics, such as the inbound traffic, outbound traffic, received packets, and transmitted packets, to help you monitor the status of Elastic IP Address (EIP). CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs. After you purchase EIP, CloudMonitor automatically collects data based on these metrics.

Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Inbound bandwidth | The traffic per second that passes through EIP to Elastic Compute Service (ECS). | Instance | bit/s | 1 minute |
| Outbound bandwidth | The traffic per second that passes through EIP from ECS. | Instance | bit/s | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Received packets | The number of packets per second that pass through EIP to ECS. | Instance | PPS | 1 minute |
| Transmitted packets | The number of packets per second that pass through EIP from ECS. | Instance | PPS | 1 minute |
| Packet loss rate due to throttling | The packet loss rate when the actually used bandwidth exceeds the configured upper limit. | Instance | PPS | 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. The Elastic IP Address page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

### Alerting service

- Set alert rules
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Elastic IP Address. The Elastic IP Address page appears.
  3. Click Alarm Rules in the Actions column for an instance to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.6 ApsaraDB for Memcache

CloudMonitor provides multiple metrics, such as the used cache and read hit ratio, to help you monitor the status of ApsaraDB for Memcache. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs. After you purchase Memcache, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Used cache | The amount of used cache. | Instance | Bytes | 1 minute |
| Read hit ratio | The success rate of reading key-values ( KVs). | Instance | Percentage | 1 minute |
| QPS | The number of KV read requests per second. | Instance | QPS | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
|---|---|---|---|---|
| Records | The total number of KVs read in the monitoring period. | Instance | Count | 1 minute |
| Cache inbound bandwidth | The bandwidth used for writing data to the cache. | Instance | bit/s | 1 minute |
| Cache outbound bandwidth | The bandwidth used for reading data from the cache. | Instance | bit/s | 1 minute |
| Evictions | The number of KVs evicted per second. | Instance | Count/second | 1 minute |

Note:

- Monitoring data is retained 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 > ApsaraDB for Memcache. The ApsaraDB for Memcache List page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

### Alerting service

CloudMonitor allows you to set alert rules for the metrics of ApsaraDB for Memcache so that you can receive alerts when any exception occurs.

- Set alert rules for a single instance
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Memcache. The ApsaraDB for Memcache List page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart to set alert rules for the corresponding metric of this instance.
- Set alert rules for multiple instances at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Memcache. The ApsaraDB for Memcache List page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.7 ApsaraDB for Redis

CloudMonitor provides multiple metrics, such as the used capacity and used connections, to help you monitor the status of ApsaraDB for Redis. After you create an ApsaraDB for Redis instance, CloudMonitor automatically collects data based on these metrics. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

### Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Used capacity | The used capacity of the instance. | Instance | Bytes | 1 minute |
| Used connections | The number of client connections. | Instance | Count | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Write bandwidth | The write traffic per second. | Instance | bit/s | 1 minute |
| Read bandwidth | The read traffic per second. | Instance | bit/s | 1 minute |
| Failed operations | The number of failed KVStore operations. | Instance | Count | 1 minute |
| Used capacity percentage | The percentage of the used capacity to the total capacity. | Instance | Percentage | 1 minute |
| Used connection percentage | The percentage of used connections to total connections. | Instance | Percentage | 1 minute |
| Write bandwidth usage | The percentage of the write bandwidth to the total bandwidth. | Instance | Percentage | 1 minute |
| Read bandwidth usage | The percentage of the read bandwidth to the total bandwidth. | Instance | Percentage | 1 minute |
| Instance failure | An event-type metric. You can set alert rules for this metric. | - | - | - |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Instance failover | An event-type metric. You can set alert rules for this metric. | - | - | - |

- View monitoring data

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Redis. The Redis Monitoring List page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for Redis. The Redis Monitoring List page appears.
  3. Click Alarm Rules in the Actions column for an instance to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.8 ApsaraDB for MongoDB

CloudMonitor provides multiple metrics, such as the CPU usage and memory usage, to help you monitor the status of ApsaraDB for MongoDB. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs. After you purchase ApsaraDB for MongoDB, CloudMonitor automatically collects data based on these metrics.

Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| CPU usage | The CPU usage of the instance. | User, instance, and primary/secondary node | Percentage | 5 minutes |
| Memory usage | The memory usage of the instance. | User, instance, and primary/secondary node | Percentage | 5 minutes |
| Disk usage | The disk usage of the instance. | User, instance, and primary/secondary node | Percentage | 5 minutes |
| Input/Output operations per second (IOPS) usage | The IOPS usage of the instance. | User, instance, and primary/secondary node | Percentage | 5 minutes |
| Connection usage | The percentage of instances to which the current application connects. The number of instances to which an application can connect is limited. This metric indicates the percentage of connected instances. | User, instance, and primary/secondary node | Percentage | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Average number of SQL queries per second | The average number of SQL queries per second for the instance. | User, instance, and primary/secondary node | Count | 5 minutes |
| Connections in use | The number of instances to which the current application connects. | User, instance, and primary/secondary node | Count | 5 minutes |
| Disk space occupied by an instance | The total disk space occupied by the instance. | User, instance, and primary/secondary node | Bytes | 5 minutes |
| Disk space occupied by data | The disk space occupied by data. | User, instance, and primary/secondary node | Bytes | 5 minutes |
| Disk space occupied by logs | The disk space occupied by logs. | User, instance, and primary/secondary node | Bytes | 5 minutes |
| Inbound internal network traffic | The inbound traffic of the instance. | User, instance, and primary/secondary node | Bytes | 5 minutes |
| Outbound internal network traffic | The outbound traffic of the instance. | User, instance, and primary/secondary node | Bytes | 5 minutes |
| Request count | The total number of requests sent to the server. | User, instance, and primary/secondary node | Count | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Insert operations | The total number of insert commands received since the last time the instance was started. | User, instance, and primary/secondary node | Count | 5 minutes |
| Query operations | The total number of query commands received since the last time the instance was started. | User, instance, and primary/secondary node | Count | 5 minutes |
| Update operations | The total number of update commands received since the last time the instance was started. | User, instance, and primary/secondary node | Count | 5 minutes |
| Delete operations | The total number of delete operations performed since the last time the instance was started. | User, instance, and primary/secondary node | Count | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Getmore operations | The total number of getmore operations performed since the last time the instance was started. | User, instance, and primary/secondary node | Count | 5 minutes |
| Command operations | The total number of commands sent to the database since the last time the instance was started. | User, instance, and primary/secondary node | Count | 5 minutes |
| Instance failure | An event-type metric. You can set alert rules for this metric. | - | - | - |

Note:

- Monitoring data is retained 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 > ApsaraDB for MongoDB. The MongoDB page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
  5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single instance

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for MongoDB. The MongoDB page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart to set alert rules for the corresponding metric of this instance.

- Set alert rules for multiple instances at a time

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > ApsaraDB for MongoDB. The MongoDB page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.9 MNS

CloudMonitor provides multiple metrics, such as the numbers of delayed messages, invalid messages, and active messages, to help you monitor the status of Message Service (MNS). After you create an MNS queue, CloudMonitor automatically collects

data based on these metrics. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

### Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| ActiveMessages | The total number of active messages in the queue. | User, region, bucket, and queue | Count | 5 minutes |
| InactiveMessages | The total number of inactive messages in the queue. | User, region, bucket, and queue | Count | 5 minutes |
| DelayMessage | The total number of delayed messages in the queue. | User, region, bucket, and queue | Count | 5 minutes |
| SendMessageCount | The number of requests for sending a message. | User, region, and queue | Count | 60 minutes |
| BatchSendMessageCount | The number of requests for sending multiple messages at a time. | User, region, and queue | Count | 60 minutes |
| ReceiveMessageCount | The number of requests for receiving a message. | User, region, and queue | Count | 60 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| BatchReceiveMessageCount | The number of requests for receiving multiple messages at a time. | User, region, and queue | Count | 60 minutes |
| BatchDeleteMessageCount | The number of requests for deleting multiple messages at a time. | User, region, and queue | Count | 60 minutes |
| ChangeMessageVisibilityCount | Change the number of visible messages. | User, region, and queue | Count | 60 minutes |

- View monitoring data

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Message Service. The MNS List page appears.
  3. Click the name of a queue or click Monitoring Charts in the Actions column for a queue to view the monitoring charts.
  4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Message Service. The MNS List page appears.
  3. Click Alarm Rules in the Actions column for a queue to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.10 AnalyticDB

CloudMonitor provides multiple metrics, such as the rated disk capacity, occupied disk capacity, and disk usage, to help you monitor the status of AnalyticDB. After you purchase AnalyticDB, CloudMonitor automatically collects data based on these metrics. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

### Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| diskSize | The rated disk capacity. | Instance, table schema, and worker | MB | 1 minute |
| diskUsed | The occupied disk capacity. | Instance, table schema, and worker | MB | 1 minute |
| diskUsedPercent | The disk usage. | Instance, table schema, and worker | Percentage | 1 minute |

- View monitoring data

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > AnalyticDB. The ADS page appears.
  3. Click the name of a database or click Monitoring Charts in the Actions column for a database to view the monitoring charts.
  4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > AnalyticDB. The ADS page appears.
  3. Click Alarm Rules in the Actions column for a database to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.11 Log Service

CloudMonitor provides multiple metrics, such as the inbound traffic and outbound traffic, total queries per second (QPS), and log statistics, to help you monitor the status of Log Service. After you create a Log Service project, CloudMonitor automatically collects data based on these metrics. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Inflow | The inbound traffic of the Logstore per minute. | User, project, and Logstore | Bytes | 1 minute |
| Outflow | The outbound traffic of the Logstore per minute. | User, project, and Logstore | Bytes | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| SumQPS | The number of writes per minute in the Logstore. | User, project, and Logstore | Count | 1 minute |
| LogMethodQPS | The number of writes per minute for each method in the Logstore. | User, project, Logstore, and method | Count | 1 minute |
| LogCodeQPS | The number of writes per minute for each status code in the Logstore. | User, project, Logstore, and status | Count | 1 minute |
| SuccessdByte | The number of resolved bytes in the Logstore. | User, project, and Logstore | Bytes | 10 minutes |
| SuccessdLines | The number of lines in resolved logs in the Logstore. | User, project, and Logstore | Count | 10 minutes |
| FailedLines | The number of lines in logs that failed to be resolved in the Logstore. | User, project, and Logstore | Count | 10 minutes |
| AlarmPV | The total number of Elastic Compute Service (ECS) configuration errors in the Logstore. | User, project, and Logstore | Count | 5 minutes |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| AlarmUv | The total number of ECS instances with incorrect configurations in the Logstore. | User, project, and Logstore | Count | 5 minutes |
| AlarmIPCount | The number of errors incurred by each IP address in the Logstore. | User, project, Logstore, alert type, and source IP address | 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. The Log Service page appears.
  3. Click Monitoring Charts in the Actions column for a project to view the monitoring charts.
  4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Log Service. The Log Service page appears.
  3. Click Alarm Rules in the Actions column for a project to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Parameters

  Note:

- When setting alert rules, you can specify the status for a status-related metric. Valid values of the status field include 200, 400, 401, 403, 405, 500, and 502.
- You can specify the method for a metric related to the operation count. Valid values of the method field include PostLogStoreLogs, GetLogtailConfig, PutData, GetCursorOrData, GetData, GetLogStoreHistogram, GetLogStoreLogs, ListLogStores, and ListLogStoreTopics.

For more information about alert rule parameters, see #unique_45.

## 6.12 Container Service

CloudMonitor provides multiple metrics, such as the CPU usage and memory usage, to help you monitor the status of Container Service. After you create a Container Service cluster, CloudMonitor automatically collects data based on these metrics. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

### Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
|---|---|---|---|---|
| containerCpuUtilization | The CPU usage of the container. | User and container | Percentage | 30 seconds |
| containerMemoryUtilization | The memory usage of the container. | User and container | Percentage | 30 seconds |
| containerMemoryAmount | The amount of memory used by the container. | User and container | Bytes | 30 seconds |
| containerInternetIn | The inbound traffic of the container. | User and container | Bytes | 30 seconds |
| containerInternetOut | The outbound traffic of the container. | User and container | Bytes | 30 seconds |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| containerIOPS | The I/O read traffic of the container. | User and container | Bytes | 30 seconds |
| containerIOWrite | The I/O write traffic of the container. | User and container | Bytes | 30 seconds |

Note:

- Monitoring data is retained 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 > Container Service. The Clusters page appears.
3. Click Monitoring Charts in the Actions column for a cluster to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single cluster

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Cloud Service Monitoring > Container Service. The Clusters page appears.
3. Click Monitoring Charts in the Actions column for a cluster to view the monitoring charts.
4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple clusters at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Container Service. The Clusters page appears.
  3. Select target clusters and click Set Alarm Rules under the list to set alert rules for the selected clusters.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.13 Shared Bandwidth

CloudMonitor provides multiple metrics, such as the inbound bandwidth and outbound bandwidth, to help you monitor the status of Shared Bandwidth. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Shared Bandwidth, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Inbound bandwidth of a bandwidth package | User and instance | bit/s | 1 minute |
| Outbound bandwidth of a bandwidth package | User and instance | bit/s | 1 minute |
| Received packets of a bandwidth package | User and instance | PPS | 1 minute |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Transmitted packets of a bandwidth package | User and instance | PPS | 1 minute |
| Outbound bandwidth usage of a bandwidth package | User and instance | Percentage | 1 minute |

Note:

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

- View monitoring data

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Shared Bandwidth. The Shared Bandwidth page appears.
  3. Click the ID of a bandwidth package or click Monitoring Charts in the Actions column for a bandwidth package to view the monitoring charts.
  4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to seven consecutive days.
  5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single bandwidth package

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Shared Bandwidth. The Shared Bandwidth page appears.
  3. Click the ID of a bandwidth package or click Monitoring Charts in the Actions column for a bandwidth package to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple bandwidth packages at a time

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Shared Bandwidth. The Shared Bandwidth page appears.
  3. Select target bandwidth packages and click Set Alarm Rules under the list to set alert rules for the bandwidth packages.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.14 Global Acceleration

CloudMonitor provides multiple metrics, such as the inbound bandwidth and outbound bandwidth, to help you monitor the status of Global Acceleration. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Global Acceleration, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Inbound bandwidth | User and instance | bit/s | 1 minute |
| Outbound bandwidth | User and instance | bit/s | 1 minute |
| Received packets | User and instance | PPS | 1 minute |
| Transmitted packets | User and instance | PPS | 1 minute |

Note:

- Monitoring data is retained for up to 31 days.
- You can view the monitoring data for up to seven 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. The Global Acceleration page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to seven consecutive days.
  5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single instance

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Global Acceleration. The Global Acceleration page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple instances at a time

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Global Acceleration. The Global Acceleration page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.15 TSDB

CloudMonitor provides multiple metrics, such as the disk usage, timeline quantity, and time point increment, to help you monitor the status of Time Series and Spatial-

Temporal Database (TSDB). CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase TSDB, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Disk usage | User and instance | Percentage | 20 seconds |
| Timeline quantity | User and instance | Count | 20 seconds |
| Time point increment | User and instance | Percentage | 20 seconds |

Note:

- Monitoring data is retained 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. The HiTSDB page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single instance
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > HiTSDB. The HiTSDB page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Set alert rules for multiple instances at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > HiTSDB. The HiTSDB page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.16 VPN Gateway

CloudMonitor provides multiple metrics, such as the inbound bandwidth and outbound bandwidth, to help you monitor the status of Virtual Private Network (VPN) Gateway. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase VPN Gateway, CloudMonitor automatically collects data based on these metrics.

Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Inbound bandwidth of a bandwidth package | User and instance | bit/s | 1 minute |
| Outbound bandwidth of a bandwidth package | User and instance | bit/s | 1 minute |
| Received packets of a bandwidth package | User and instance | PPS | 1 minute |
| Transmitted packets of a bandwidth package | User and instance | PPS | 1 minute |

Note:

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

- View monitoring data

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Cloud Service Monitoring > VPN. The VPN page appears.
3. Click the ID of a VPN or click Monitoring Charts in the Actions column for a VPN to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to seven consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single VPN
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > VPN. The VPN page appears.
  3. Click the ID of a VPN or click Monitoring Charts in the Actions column for a VPN to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Set alert rules for multiple VPNs at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > VPN. The VPN page appears.
  3. Select target VPNs and click Set Alarm Rules under the list to set alert rules for the selected VPNs.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.17 API Gateway

CloudMonitor provides multiple metrics, such as the inbound traffic, outbound traffic, and response time, to help you monitor the status of API Gateway.

After you purchase API Gateway, CloudMonitor automatically collects data based on these metrics. You can log on to the CloudMonitor console and view the monitoring details on the API Gateway monitoring page. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

Monitoring service

- Metrics

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Error distribution | The numbers of 2XX, 4XX, and 5XX status codes returned for an API in the monitoring period. | User and API | Count | 1 minute |
| Inbound traffic | The total traffic of requests received by an API in the monitoring period. | User and API | Bytes | 1 minute |
| Outbound traffic | The total traffic of responses sent by an API in the monitoring period. | User and API | Bytes | 1 minute |
| Response time | The latency between the time when API Gateway calls the backend service of an API and the time when the result is received from the backend service in the monitoring period. | User and API | Seconds | 1 minute |

| Metric | Description | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- | --- |
| Total requests | The total number of requests received by an API in the 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. The API Gateway Monitoring List page appears.
  3. Click the name of an API or click Monitoring Charts in the Actions column for an API to view the monitoring charts.
  4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > API Gateway. The API Gateway Monitoring List page appears.
  3. Click Alarm Rules in the Actions column for an API to view the alert rules.
  4. Click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.18 DirectMail

CloudMonitor provides multiple metrics, such as metrics about the Web or API messaging sending method, Simple Mail Transfer Protocol (SMTP) message sending method, and abnormal accounts, to help you monitor the status of DirectMail.

CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase DirectMail, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Unit | Minimum frequency |
| --- | --- | --- |
| Web or API over-length-error QPS | Count/minute | 1 minute |
| Web/API over-quota-error QPS | Count/minute | 1 minute |
| Web/API spam QPS | Count/minute | 1 minute |
| Web/API success QPS | Count/minute | 1 minute |
| SMTP authentication failure QPS | Count/minute | 1 minute |
| SMTP authentication success QPS | Count/minute | 1 minute |
| SMTP over-length-error QPS | Count/minute | 1 minute |
| SMTP over-quota-error QPS | Count/minute | 1 minute |
| SMTP spam QPS | Count/minute | 1 minute |

Note:

- Monitoring data is retained 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 > DirectMail. On the DirectMail page that appears, you can view the metrics of DirectMail.

### Alerting service

CloudMonitor allows you to set alert rules for the metrics of DirectMail so that you can receive alerts when any exception occurs.

- Set alert rules
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > DirectMail. The DirectMail page appears.
  3. Click the Alarm Rules tab and then click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.19 Elasticsearch

CloudMonitor provides multiple metrics, such as the cluster status, cluster queries per second (QPS), and cluster write QPS, to help you monitor the status of Elasticsearch. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Elasticsearch, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
|---|---|---|---|
| Cluster status | Cluster | | 1 minute |
| Cluster QPS | Cluster | Count/second | 1 minute |
| Cluster write QPS | Cluster | Count/second | 1 minute |
| CPU usage of a node | Node | Percentage | 1 minute |
| Disk usage of a node | Node | Percentage | 1 minute |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Heap memory usage of a node | Node | Percentage | 1 minute |
| Load of a node within 1 minute | Node | | 1 minute |
| FullGC times of a node | Node | Count | 1 minute |
| Exceptions of a node | Node | Count | 1 minute |
| Cluster snapshot status | Cluster | A value of -1 indicates that no snapshot exists. A value of 0 indicates that the snapshot is created. A value of 1 indicates that the snapshot is being created. A value of 2 indicates that the snapshot fails to be created. | 1 minute |

Note:

- Monitoring data is retained 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. The Elasticsearch page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Elasticsearch. The Elasticsearch page appears.
  3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.20 Auto Scaling

CloudMonitor provides multiple metrics, such as the minimum and maximum numbers of instances, to help you monitor the status of Auto Scaling. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Auto Scaling, CloudMonitor automatically collects data based on these metrics.

Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Minimum number of instances | User and scaling group | Count | 5 minutes |
| Maximum number of instances | User and scaling group | Count | 5 minutes |
| Total instances | User and scaling group | Count | 5 minutes |
| Running instances | User and scaling group | Count | 5 minutes |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Instances being added | User and scaling group | Count | 5 minutes |
| Instances being removed | User and scaling group | Count | 5 minutes |

Note:

- Monitoring data is retained 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 > Auto Scaling. The Alibaba Cloud Auto Scaling page appears.
3. Click the name of a scaling group or click Monitoring Charts in the Actions column for a scaling group to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single scaling group

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Cloud Service Monitoring > Auto Scaling. The Alibaba Cloud Auto Scaling page appears.
3. Click the name of a scaling group or click Monitoring Charts in the Actions column for a scaling group to view the monitoring charts.
4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple scaling groups at a time

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Auto Scaling. The Alibaba Cloud Auto Scaling page appears.
  3. Select target scaling groups and click Set Alarm Rules under the list to set alert rules for the selected scaling groups.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.21 E-MapReduce

CloudMonitor provides multiple metrics, such as the CPU idle rate, memory capacity, and disk capacity, to help you monitor the status of Elastic MapReduce (E-MapReduce). CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase E-MapReduce, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Inbound traffic rate | User, cluster, and role | bit/s | 30 seconds |
| Outbound traffic rate | User, cluster, and role | bit/s | 30 seconds |
| CPU idle rate | User, cluster, and role | Percentage | 1 minute |
| User-mode CPU usage | User, cluster, and role | Percentage | 30 seconds |
| System-mode CPU usage | User, cluster, and role | Percentage | 30 seconds |
| Idle disk capacity | User, cluster, and role | Bytes | 30 seconds |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Total disk capacity | User, cluster, and role | Bytes | 30 seconds |
| 15-minute load average | User, cluster, and role | - | 30 seconds |
| 5-minute load average | User, cluster, and role | - | 30 seconds |
| 1-minute load average | User, cluster, and role | - | 30 seconds |
| Idle memory capacity | User, cluster, and role | Bytes | 30 seconds |
| Total memory capacity | User, cluster, and role | Bytes | 30 seconds |
| Received packets per second | User, cluster, and role | PPS | 30 seconds |
| Transmitted packets per second | User, cluster, and role | PPS | 30 seconds |
| Running processes | User, cluster, and role | Count | 30 seconds |
| Total processes | User, cluster, and role | Count | 30 seconds |
| Blocked processes | User, cluster, and role | Count | 30 seconds |
| Created processes or threads | User, cluster, and role | Count | 30 seconds |
| MemNonHeap UsedM | User, cluster, and role | Bytes | 30 seconds |
| MemNonHeap CommittedM | User, cluster, and role | Bytes | 30 seconds |
| MemNonHeap MaxM | User, cluster, and role | Bytes | 30 seconds |
| MemHeapUsedM | User, cluster, and role | Bytes | 30 seconds |
| MemHeapCommittedM | User, cluster, and role | Bytes | 30 seconds |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| MemHeapMaxM | User, cluster, and role | Bytes | 30 seconds |
| MemMaxM | User, cluster, and role | Bytes | 30 seconds |
| ThreadsNew | User, cluster, and role | - | 30 seconds |
| ThreadsRunnable | User, cluster, and role | - | 30 seconds |
| ThreadsBlocked | User, cluster, and role | - | 30 seconds |
| ThreadsWaiting | User, cluster, and role | - | 30 seconds |
| ThreadsTimedWaiting | User, cluster, and role | - | 30 seconds |
| ThreadsTerminated | User, cluster, and role | - | 30 seconds |
| GcCount | User, cluster, and role | - | 30 seconds |
| GcTimeMillis | User, cluster, and role | - | 30 seconds |
| CallQueueLength | User, cluster, and role | - | 30 seconds |
| NumOpenConnections | User, cluster, and role | - | 30 seconds |
| ReceivedByte | User, cluster, and role | - | 30 seconds |
| SentByte | User, cluster, and role | - | 30 seconds |
| BlockCapacity | User, cluster, and role | - | 30 seconds |
| BlocksTotal | User, cluster, and role | - | 30 seconds |
| CapacityRemaining | User, cluster, and role | - | 30 seconds |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| CapacityTotal | User, cluster, and role | - | 30 seconds |
| CapacityUsed | User, cluster, and role | - | 30 seconds |
| CapacityUsedNonDFS | User, cluster, and role | - | 30 seconds |
| CorruptBlocks | User, cluster, and role | - | 30 seconds |
| ExcessBlocks | User, cluster, and role | - | 30 seconds |
| ExpiredHeartbeats | User, cluster, and role | - | 30 seconds |
| MissingBlocks | User, cluster, and role | - | 30 seconds |
| PendingDataNodeMessageCount | User, cluster, and role | - | 30 seconds |
| PendingDeletionBlocks | User, cluster, and role | - | 30 seconds |
| PendingReplicationBlocks | User, cluster, and role | - | 30 seconds |
| PostponedMisreplicatedBlocks | User, cluster, and role | - | 30 seconds |
| ScheduledReplicationBlocks | User, cluster, and role | - | 30 seconds |
| TotalFiles | User, cluster, and role | - | 30 seconds |
| TotalLoad | User, cluster, and role | - | 30 seconds |
| UnderReplicatedBlocks | User, cluster, and role | - | 30 seconds |
| BlocksRead | User, cluster, and role | - | 30 seconds |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| BlocksRemoved | User, cluster, and role | - | 30 seconds |
| BlocksReplicated | User, cluster, and role | - | 30 seconds |
| BlocksUncached | User, cluster, and role | - | 30 seconds |
| BlocksVerified | User, cluster, and role | - | 30 seconds |
| BlockVerificationFailures | User, cluster, and role | - | 30 seconds |
| BlocksWritten | User, cluster, and role | - | 30 seconds |
| ByteRead | User, cluster, and role | - | 30 seconds |
| ByteWritten | User, cluster, and role | - | 30 seconds |
| FlushNanosAvgTime | User, cluster, and role | - | 30 seconds |
| FlushNanosNumOps | User, cluster, and role | - | 30 seconds |
| FsyncCount | User, cluster, and role | - | 30 seconds |
| VolumeFailures | User, cluster, and role | - | 30 seconds |
| ReadBlockOpNumOps | User, cluster, and role | - | 30 seconds |
| ReadBlockOpAvgTime | User, cluster, and role | Milliseconds | 30 seconds |
| WriteBlockOpNumOps | User, cluster, and role | - | 30 seconds |
| WriteBlockOpAvgTime | User, cluster, and role | Milliseconds | 30 seconds |
| BlockChecksumOpNumOps | User, cluster, and role | - | 30 seconds |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| BlockChecksumOpAvgTime | User, cluster, and role | Milliseconds | 30 seconds |
| CopyBlockOpNumOps | User, cluster, and role | - | 30 seconds |
| CopyBlockOpAvgTime | User, cluster, and role | Milliseconds | 30 seconds |
| ReplaceBlockOpNumOps | User, cluster, and role | - | 30 seconds |
| ReplaceBlockOpAvgTime | User, cluster, and role | Milliseconds | 30 seconds |
| BlockReportsNumOps | User, cluster, and role | - | 30 seconds |
| BlockReportsAvgTime | User, cluster, and role | Milliseconds | 30 seconds |
| NodeManager_AllocatedContainers | User, cluster, and role | - | 30 seconds |
| Containers Completed | User, cluster, and role | - | 30 seconds |
| ContainersFailed | User, cluster, and role | - | 30 seconds |
| ContainersIniting | User, cluster, and role | - | 30 seconds |
| ContainersKilled | User, cluster, and role | - | 30 seconds |
| Containers Launched | User, cluster, and role | - | 30 seconds |
| Containers Running | User, cluster, and role | - | 30 seconds |
| ActiveApplications | User, cluster, and role | - | 30 seconds |
| ActiveUsers | User, cluster, and role | - | 30 seconds |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| AggregateContainersAllocated | User, cluster, and role | - | 30 seconds |
| AggregateContainersReleased | User, cluster, and role | - | 30 seconds |
| AllocatedContainers | User, cluster, and role | - | 30 seconds |
| AppsCompleted | User, cluster, and role | - | 30 seconds |
| AppsFailed | User, cluster, and role | - | 30 seconds |
| AppsKilled | User, cluster, and role | - | 30 seconds |
| AppsPending | User, cluster, and role | - | 30 seconds |
| AppsRunning | User, cluster, and role | - | 30 seconds |
| AppsSubmitted | User, cluster, and role | - | 30 seconds |
| AvailableMB | User, cluster, and role | - | 30 seconds |
| AvailableVCores | User, cluster, and role | - | 30 seconds |
| PendingContainers | User, cluster, and role | - | 30 seconds |
| ReservedContainers | User, cluster, and role | - | 30 seconds |

Note:

- Monitoring data is retained 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 > E-MapReduce. The E-MapReduce Monitoring List page appears.
  3. Click the ID of a cluster or click Monitoring Charts in the Actions column for a cluster to view the monitoring charts.
  4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
  5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > E-MapReduce. The E-MapReduce Monitoring List page appears.
  3. Click the ID of a cluster or click Monitoring Charts in the Actions column for a cluster to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.22 Express Connect

CloudMonitor provides multiple metrics, such as the inbound traffic and outbound traffic, to help you monitor the status of Express Connect. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Express Connect, CloudMonitor automatically collects data based on these metrics.

Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Inbound traffic | User and instance | Bytes | 1 minute |
| Outbound traffic | User and instance | Bytes | 1 minute |
| Inbound bandwidth | User and instance | bit/s | 1 minute |
| Outbound bandwidth | User and instance | bit/s | 1 minute |
| Latency | User and instance | Milliseconds | 1 minute |
| Packet loss rate | User and instance | Percentage | 1 minute |

Note:

- Monitoring data is retained 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. The Express Connect page appears.
3. Click the ID of a router interface or click Monitoring Charts in the Actions column for a router interface to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

### Alerting service

- Set alert rules for a single router interface
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Express Connect. The Express Connect page appears.
  3. Click the ID of a router interface or click Monitoring Charts in the Actions column for a router interface to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Set alert rules for multiple router interfaces at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Express Connect. The Express Connect page appears.
  3. Select target router interfaces and click Set Alarm Rules under the list to set alert rules for the selected router interfaces.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.23 Function Compute

CloudMonitor provides multiple service-level and function-level metrics, such as the total invocations, average duration, and request status distribution, to help you monitor the status of Function Compute. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Function Compute, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| BillableInvocations | User, service, and function | Count | 1 minute |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| BillableInvocationsRate | User, service, and function | Percentage | 1 minute |
| ClientErrors | User, service, and function | Count | 1 minute |
| ClientErrorsRate | User, service, and function | Percentage | 1 minute |
| ServerErrors | User, service, and function | Count | 1 minute |
| ServerErrorsRate | User, service, and function | Percentage | 1 minute |
| Throttles | User, service, and function | Count | 1 minute |
| ThrottlesRate | User, service, and function | Percentage | 1 minute |
| TotalInvocations | User, service, and function | Count | 1 minute |
| Average duration | User, service, and function | Milliseconds | 1 minute |

Note:

- Monitoring data is retained 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. On the Function Compute page that appears, you can view the overall status of Function Compute.
3. Click the Service List tab to view the service-level or function-level monitoring information.

### Alerting service

CloudMonitor allows you to set alert rules for the metrics of Function Compute so that you can receive alerts when any exception occurs.

- Set alert rules for a single service
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Function Compute. The Function Compute page appears.
  3. Click the Alarm Rules tab and then click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Set alert rules for multiple services at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Function Compute. The Function Compute page appears.
  3. Click the Service List tab.
  4. Select target services and click Set Alarm Rules under the list to set alert rules for the selected services.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.24 Realtime Compute

CloudMonitor provides multiple metrics, such as the service latency, to help you monitor the status of Realtime Compute. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Realtime Compute, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Description | Minimum frequency |
| --- | --- | --- | --- | --- |
| Service latency | Project and job | Seconds | The latency between data production and data processing. | 1 minute |

| Metric | Dimension | Unit | Description | Minimum frequency |
| --- | --- | --- | --- | --- |
| Read records per second ( RPS) | Project and job | RPS | The average number of data records read per second for the job. | 1 minute |
| Write RPS | Project and job | RPS | The average number of data records written per second for the job. | 1 minute |
| FailoverRate | Project and job | Percentage | The failover rate of the job. The lower the better. | 1 minute |

Note:

- Monitoring data is retained 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 > Stream Computing. The ProjectList page appears.
3. Click the name of a project or click Monitoring Charts in the Actions column for a project to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

### Alerting service

- Set alert rules for a single project
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Stream Computing. The ProjectList page appears.
  3. Click the name of a project or click Monitoring Charts in the Actions column for a project to view the monitoring charts.
  4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.
- Set alert rules for multiple projects at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Stream Computing. The ProjectList page appears.
  3. Select target projects and click Set Alarm Rules under the list to set alert rules for the selected projects.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.25 AnalyticDB for PostgreSQL

CloudMonitor provides multiple metrics, such as the CPU usage and memory usage, to help you monitor the status of AnalyticDB for PostgreSQL. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase AnalyticDB for PostgreSQL, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Disk usage | User and instance | Percentage | 5 minutes |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Connection usage | User and instance | Percentage | 5 minutes |
| CPU usage | User and instance | Percentage | 5 minutes |
| Memory usage | User and instance | Percentage | 5 minutes |
| I/O throughput usage | User and instance | Percentage | 5 minutes |

Note:

- Monitoring data is retained 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 > AnalyticDB for PostgreSQL. The AnalyticDB for PostgreSQL Monitoring List page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single instance

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Cloud Service Monitoring > AnalyticDB for PostgreSQL. The AnalyticDB for PostgreSQL Monitoring List page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple instances at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > AnalyticDB for PostgreSQL. The AnalyticDB for PostgreSQL Monitoring List page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.26 HybridDB for MySQL

CloudMonitor provides multiple metrics, such as the disk usage, inbound bandwidth, and outbound bandwidth, to help you monitor the status of HybridDB for MySQL. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase HybridDB for MySQL, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Disk usage | User and instance | GB | 60 minutes |
| Inbound bandwidth | User and instance | KByte/s | 5 minutes |
| Outbound bandwidth | User and instance | KByte/s | 5 minutes |
| Requests per second | User and instance | Count/second | 5 minutes |
| CPU usage of a child node | User and instance | Percentage | 8 minutes |
| Disk usage of a child node | User and instance | GB | 8 minutes |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Input/Output operations per second (IOPS) of a child node | User and instance | Count/second | 8 minutes |

Note:

- Monitoring data is retained 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 > HybridDB for MySQL. The HybridDB for MySQL page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single instance

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Cloud Service Monitoring > HybridDB for MySQL. The HybridDB for MySQL page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple instances at a time

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > HybridDB for MySQL. The HybridDB for MySQL page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.

- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.27 NAT Gateway

CloudMonitor provides multiple metrics, such as the number of source network address translation (SNAT) connections, to help you monitor the status of Network Address Translation (NAT) Gateway. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase NAT Gateway, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| SNAT connections | User and instance | Count/minute | 1 minute |
| Inbound bandwidth of a bandwidth package | User and instance | bit/s | 1 minute |
| Outbound bandwidth of a bandwidth package | User and instance | bit/s | 1 minute |
| Received packets of a bandwidth package | User and instance | PPS | 1 minute |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Transmitted packets of a bandwidth package | User and instance | PPS | 1 minute |
| Outbound bandwidth usage of a bandwidth package | User and instance | Percentage | 1 minute |

Note:

- Monitoring data is retained 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 > NAT Gateway. The NAT Gateway List page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
5. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

Alerting service

- Set alert rules for a single instance

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose Cloud Service Monitoring > NAT Gateway. The NAT Gateway List page appears.
3. Click the ID of an instance or click Monitoring Charts in the Actions column for an instance to view the monitoring charts.
4. Click the bell icon in the upper-right corner of a monitoring chart or click Create Alarm Rule in the upper-right corner of the page. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Set alert rules for multiple instances at a time
  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > NAT Gateway. The NAT Gateway List page appears.
  3. Select target instances and click Set Alarm Rules under the list to set alert rules for the selected instances.
- Parameters

  For more information about alert rule parameters, see #unique_45.

## 6.28 Open Ad

CloudMonitor provides multiple metrics, such as the real-time bidding (RTB) page view (PV), queries per second (QPS) of RTB, and ad click PV, to help you to monitor the status of Open Ad. CloudMonitor also allows you to set alert rules for these metrics so that you can receive alerts when any exception occurs.

After you purchase Open Ad, CloudMonitor automatically collects data based on these metrics.

### Monitoring service

- Metrics

| Metric | Dimension | Unit | Minimum frequency |
|---|---|---|---|
| RTB PV | User | Count | 1 minute |
| RTB QPS | User | Count/second | 1 minute |
| Ad click PV | User | Count | 1 minute |
| Ad click QPS | User | Count/second | 1 minute |
| Ad click latency | User | Milliseconds | 1 minute |
| Ad impression PV | User | Count | 1 minute |
| Ad impression QPS | User | Count/second | 1 minute |
| Ad impression latency | User | Milliseconds | 1 minute |

| Metric | Dimension | Unit | Minimum frequency |
| --- | --- | --- | --- |
| Active crowd number of Data Management Platform (DMP) | User | Count/day | 1 hour |
| Valid crowd requests of DMP | User | Count/day | 1 hour |
| Storage space occupied by DMP | User | Bytes/day | 1 hour |
| Active crowd number of Umeng Data Intelligence Platform (DIP) | User | Count/day | 1 hour |
| Valid crowd requests of Umeng DIP | User | Count/day | 1 hour |

Note:

- Monitoring data is retained 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 > Openad. On the Openad page that appears, you can view the overall status of Open Ad.
3. In Time Range, select a preset time period or customize a time period. You can view the monitoring data for up to 14 consecutive days.
4. Optional. Click the zoom-in icon in the upper-right corner of a monitoring chart to enlarge the chart.

### Alerting service

CloudMonitor allows you to set alert rules for the metrics of Open Ad so that you can receive alerts when any exception occurs.

- Method 1

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Openad. The Openad page appears.
  3. Click the bell icon in the upper-right corner of a monitoring chart. Specify the resource range, set alert rules, set the notification method, and then click Confirm.

- Method 2

  1. Log on to the CloudMonitor console.
  2. In the left-side navigation pane, choose Cloud Service Monitoring > Openad. The Openad page appears.
  3. Click the Alarm Rules tab.
  4. On the Alarm Rules tab, click Create Alarm Rule in the upper-right corner. Specify the related resource, set alert rules, set the notification method, and then click Confirm.

- Parameters

  For more information about alert rule parameters, see #unique_45.

# 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 automatica lly 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 applicatio n 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.

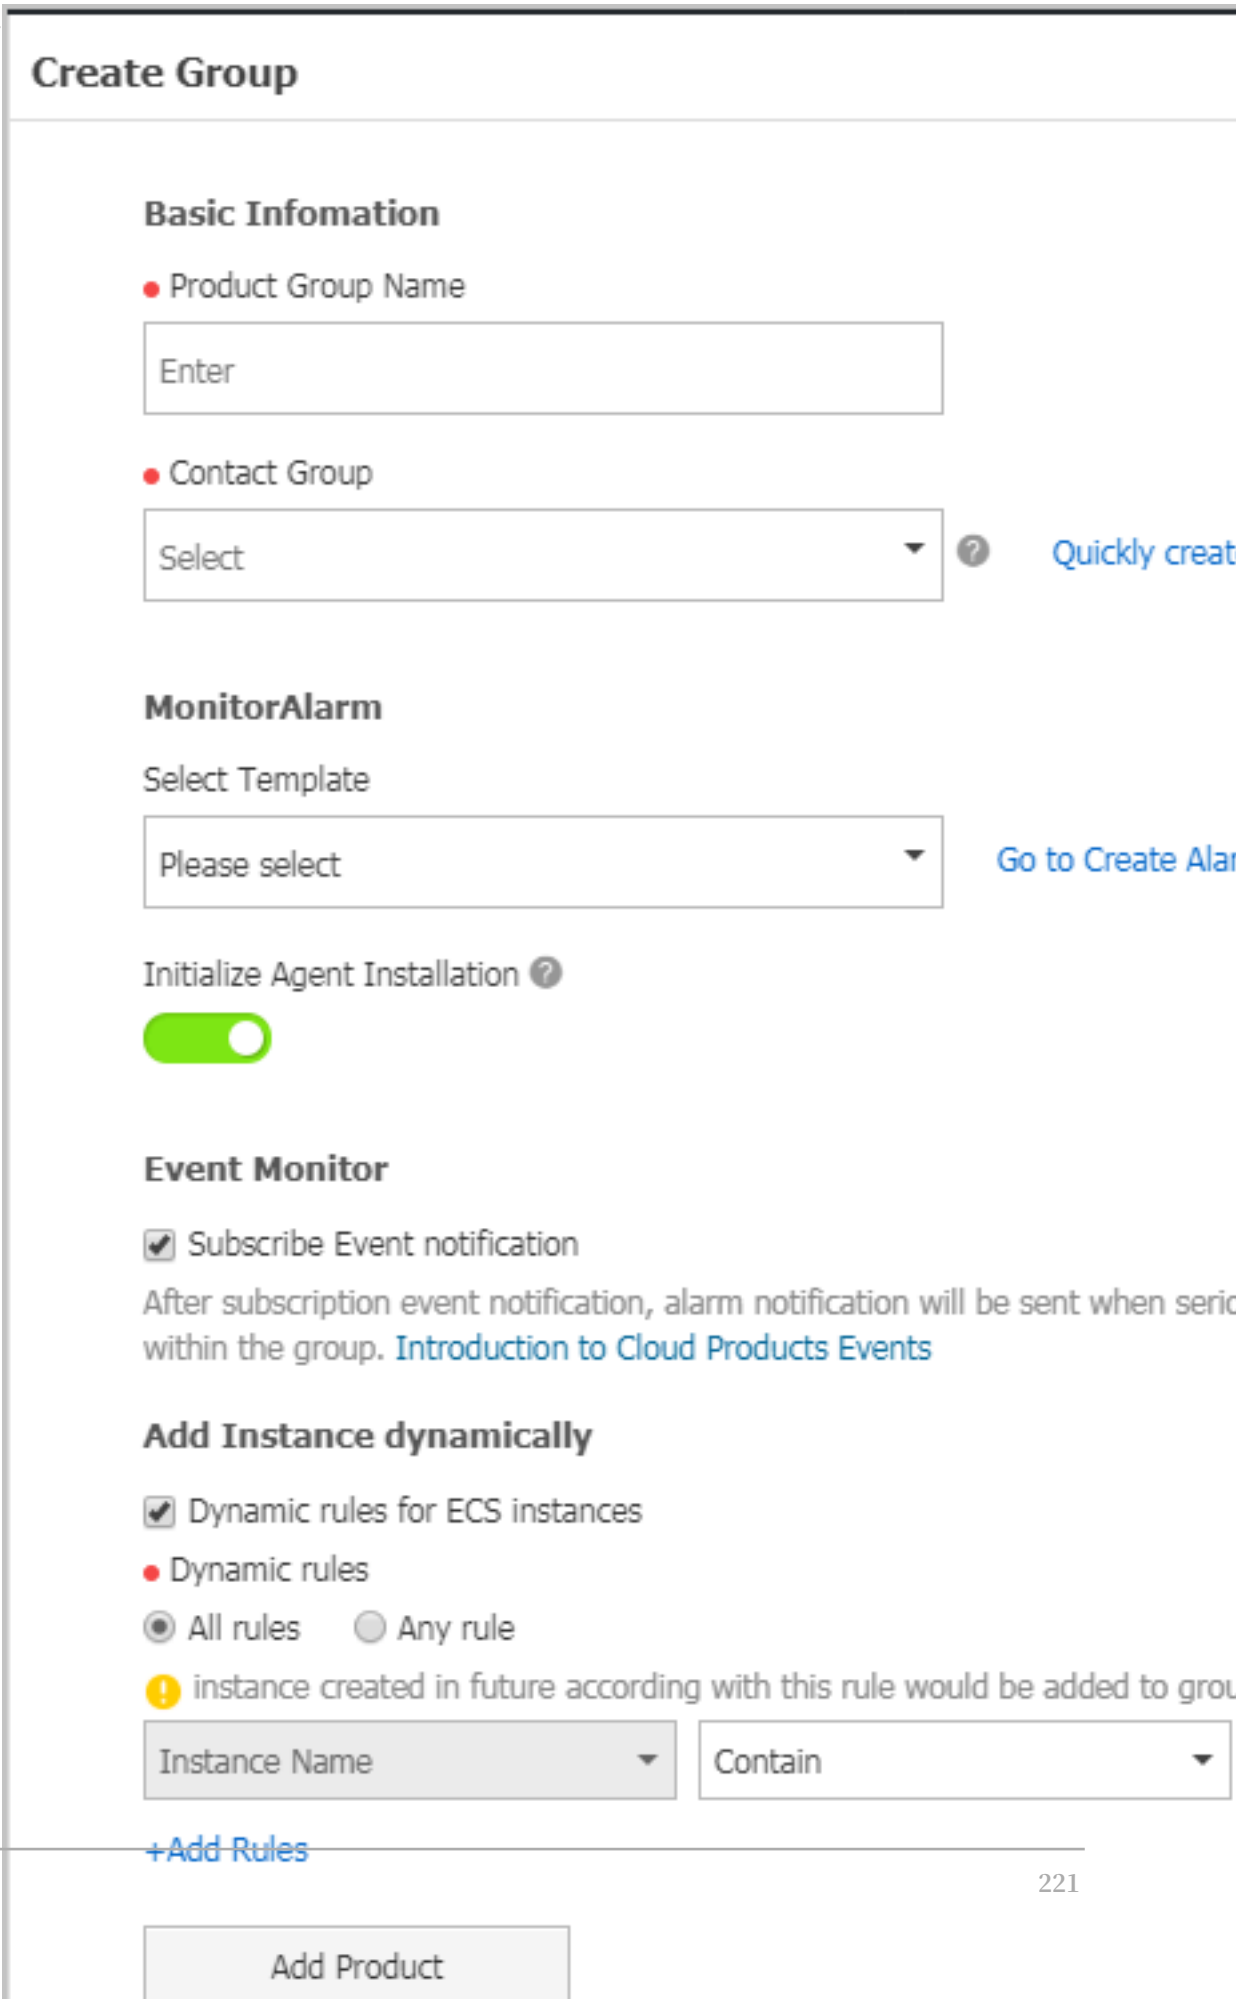

4. Enter Basic Information: Enter the group name and select one or more contact groups to receive alarm notifications.
5. Set MonitorAlarm: Select one or more templates to initialize alarm rules for the instances in the group (optional), and select the notification method. If you turn on the Initialize 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.
- Notificati on 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.
- Notificati on 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 descriptio n : 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.

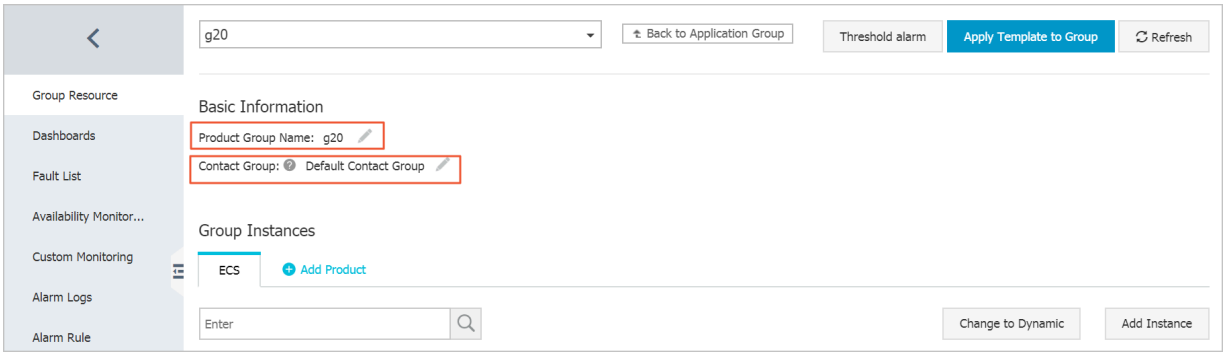


### Add or remove an instance

1. To delete an instance, click the tab of the target product, find the target instance, and click Delete in the Actions column.

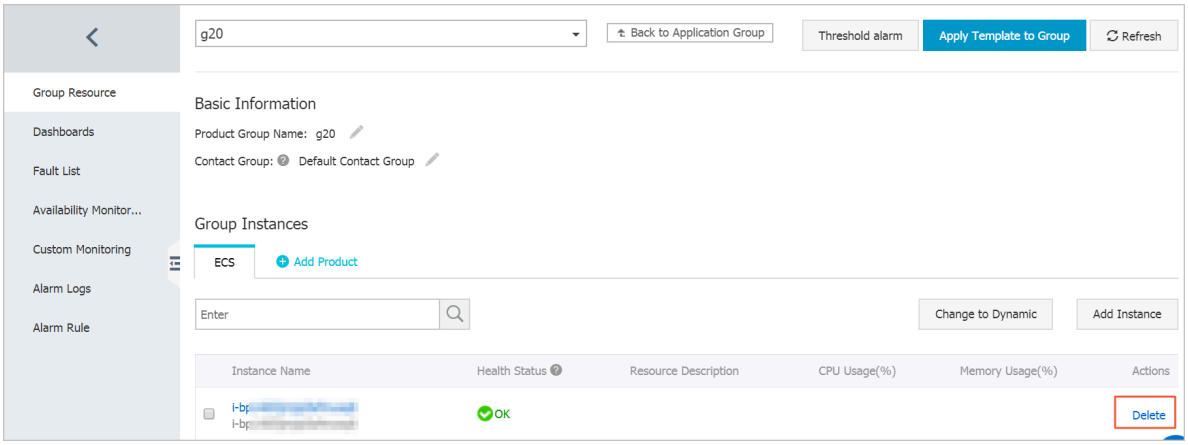

2. To add an instance, click the tab of the target product and in the upper-right corner of the tab page, click Add Instance. On the displayed AddResource page, select the target instance and click Confirm.

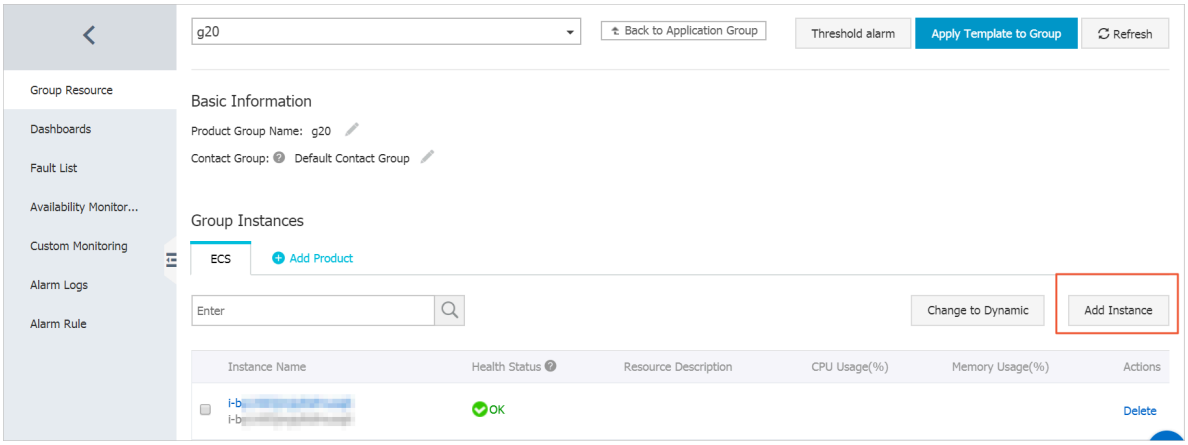


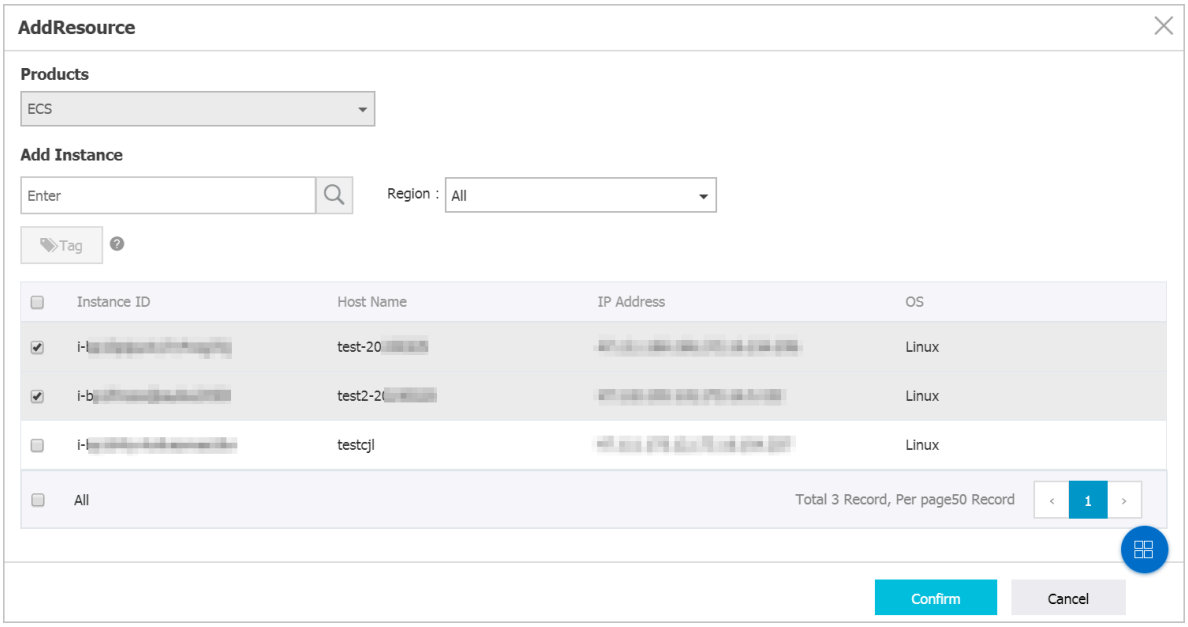


## Add a new product

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.

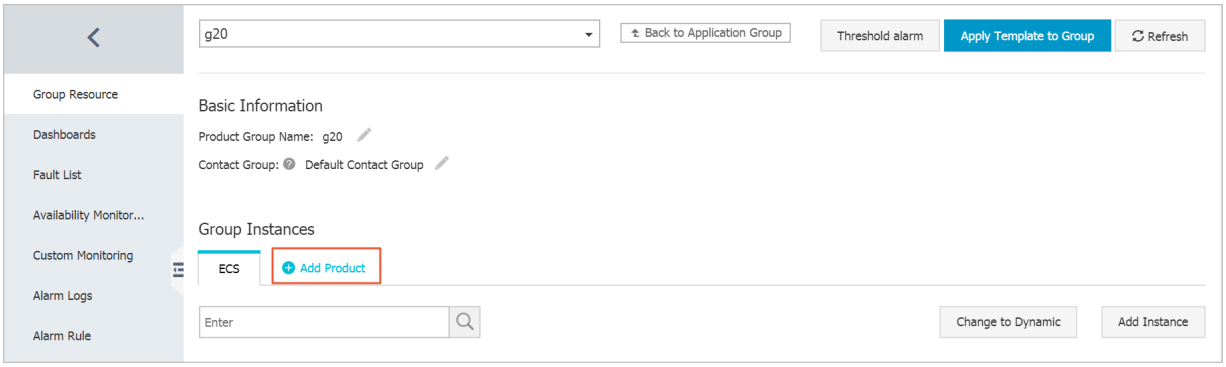


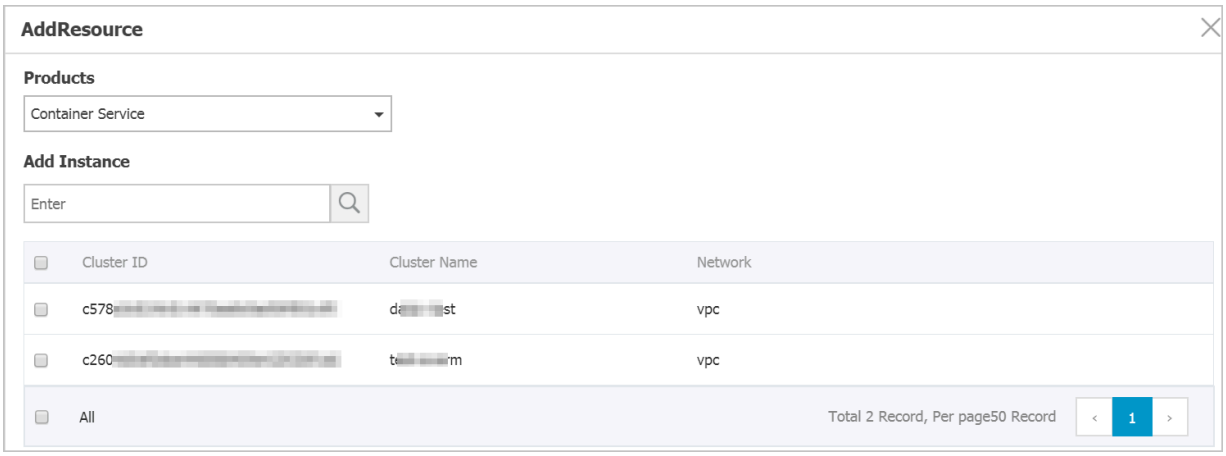


# 8.5 Add resources to an application group

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 #unique_42.

## 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.

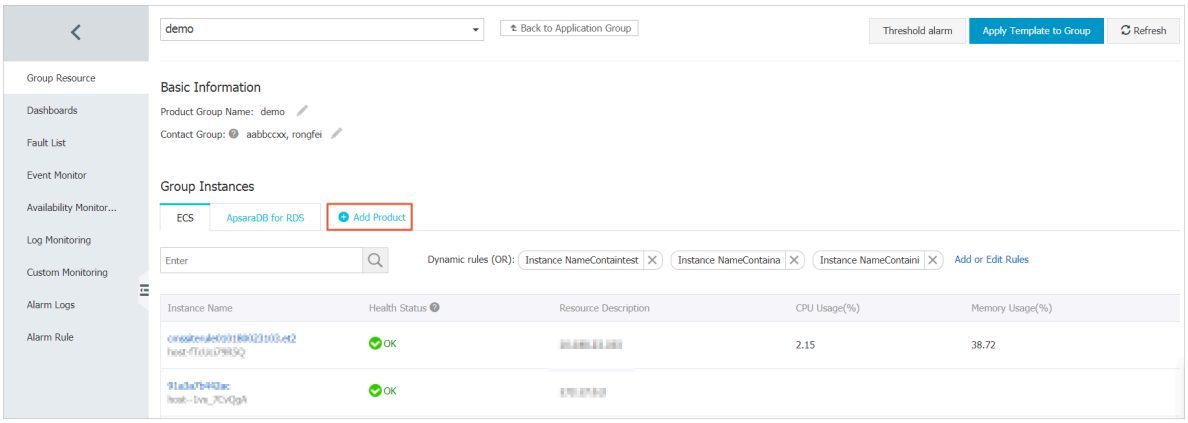


4. Click Add Product. The Add Resource page is displayed.

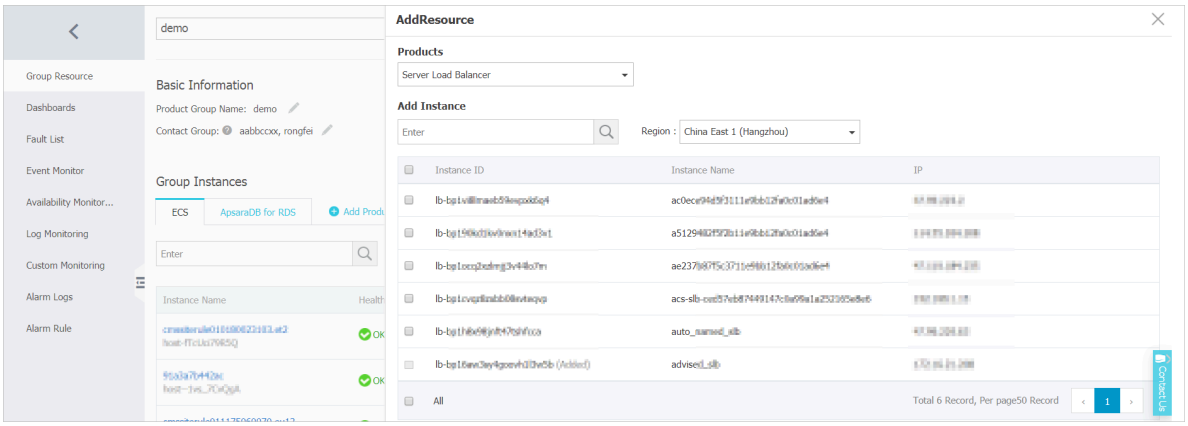


5. Select a service from the drop-down list, select the instances to be added from the instance list of the service, and click OK.

### 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.

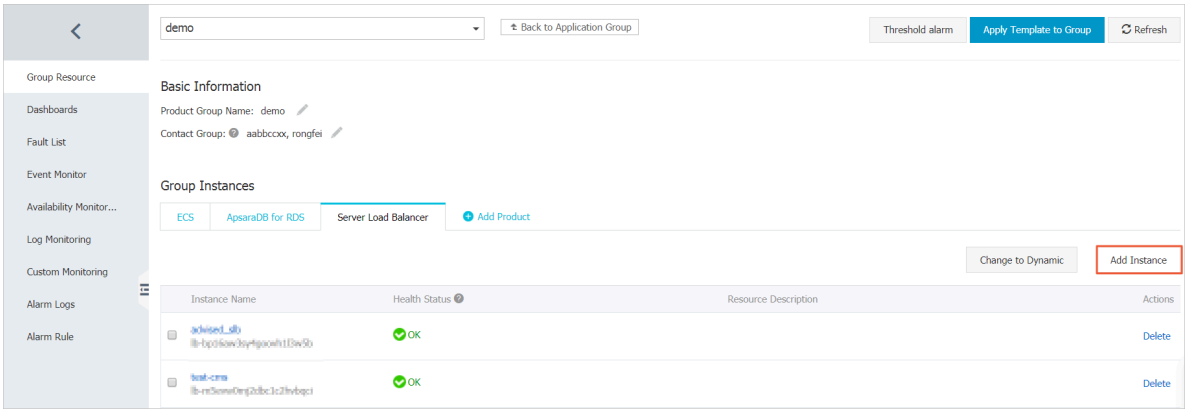


4. Click the tab of a service to be added such as OSS.
5. On the tab that appears, click Add Instance. The Add Resource page is displayed.

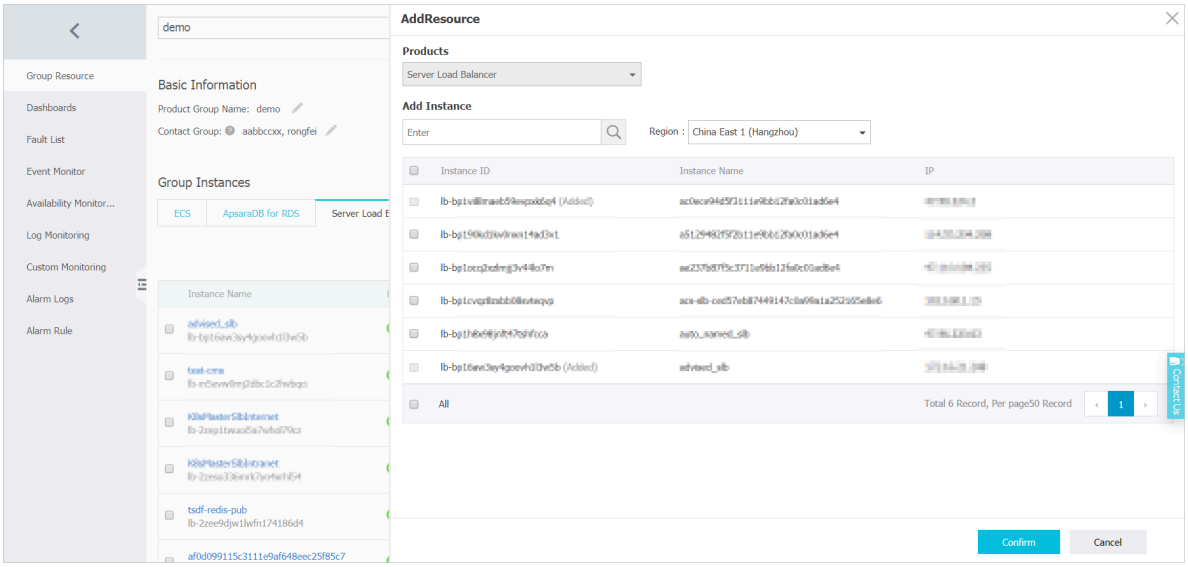


6. Select the instances to be added and click OK.

# 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 #unique_47.

## 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.
3. Click the name of the group to which you want to apply the alert template. The group details page is displayed.

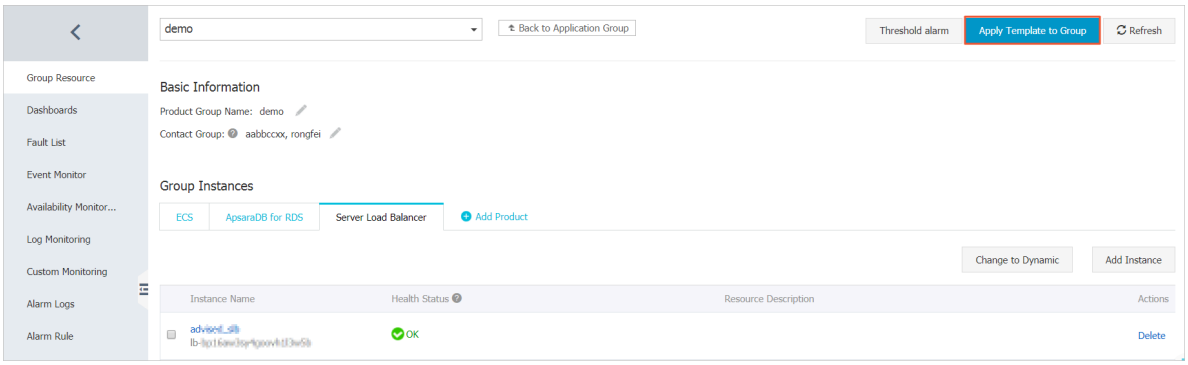

4. Click Apply Template to Group in the upper-right corner. The Apply Template to Group page is displayed.

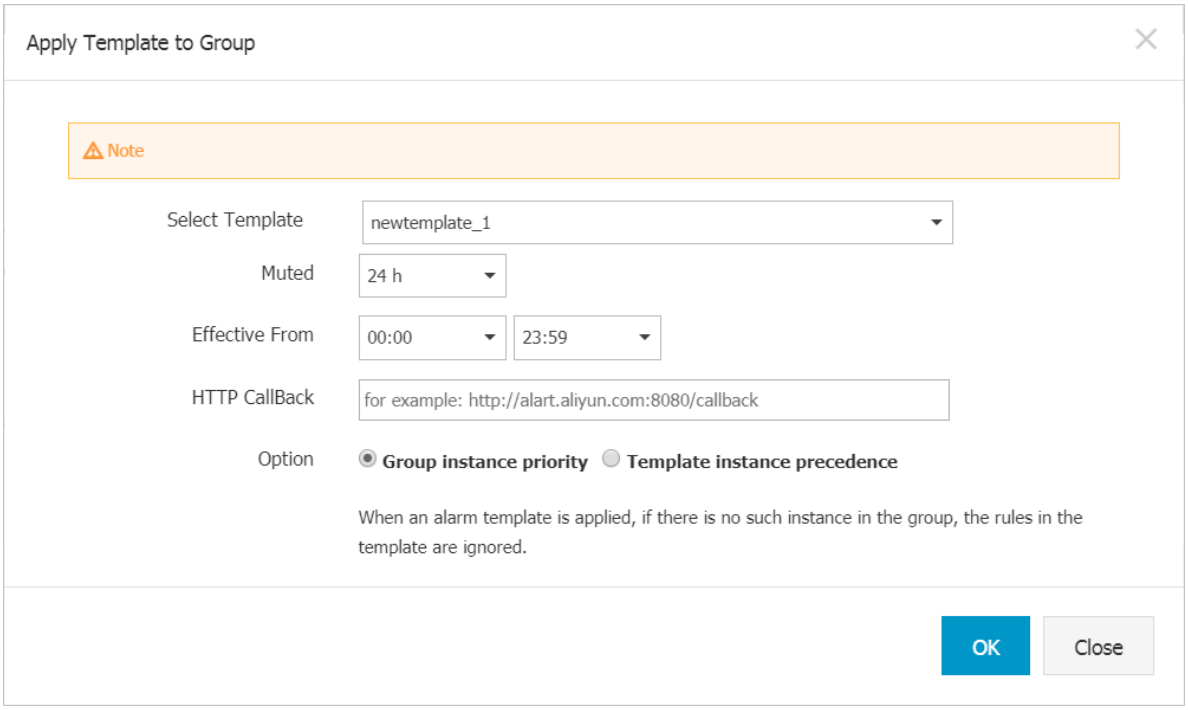


5. Select an alert template and click OK.

## 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 #unique_109.

### 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.

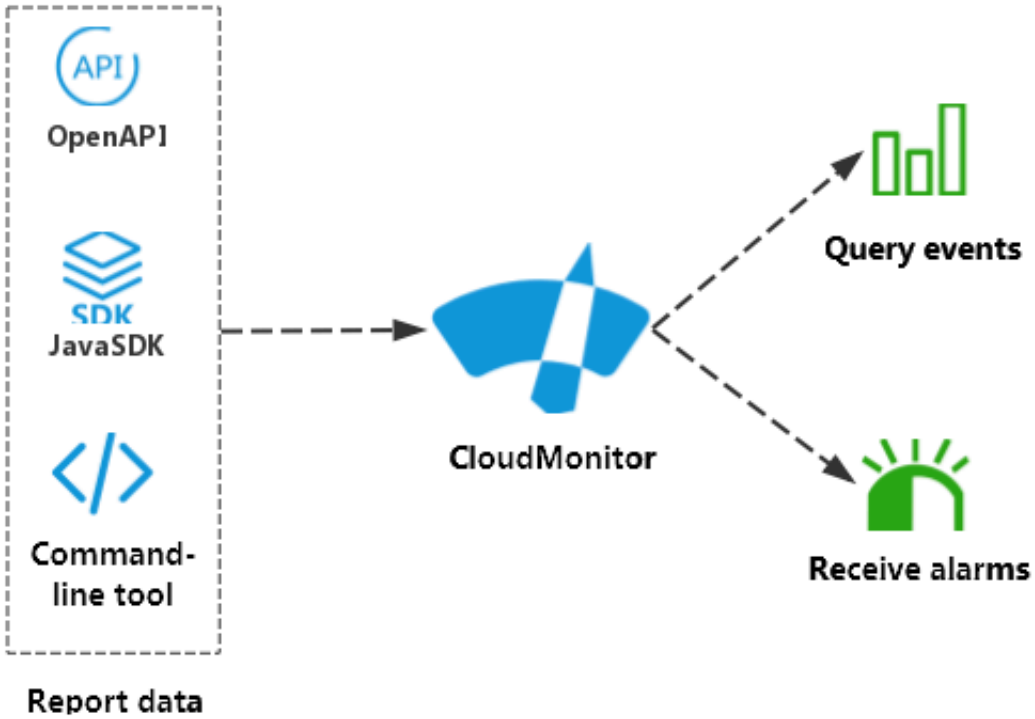


The difference between custom event monitoring and custom monitoring is as follows:

- 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 event monitoring supports.

#### Elastic Compute Service (ECS) system events

For more information about the description of ECS system events, see #unique_112.

| Name | Description | Type | Status | Level | Remarks |
| --- | --- | --- | --- | --- | --- |
| Instance:InstanceFailure.Reboot | Beginning of instance restart due to an instance error. | Exception | Executing | CRITICAL | |
| Instance:InstanceFailure.Reboot | End of instance restart due to an instance error. | Exception | Executed | CRITICAL | |
| Instance:SystemFailure.Reboot | Beginning of instance restart due to a system error. | Exception | Executing | CRITICAL | |
| Instance:SystemFailure.Reboot | End of instance restart due to a system error. | Exception | Executed | CRITICAL | |

| Name | Description | Type | Status | Level | Remarks |
| --- | --- | --- | --- | --- | --- |
| Instance:SystemMaintenance.Reboot | An instance restart is scheduled due to system maintenance. | Maintenance | Scheduled | CRITICAL | |
| Instance:SystemMaintenance.Reboot | An instance restart due to system maintenance is averted. | Maintenance | Avoided | CRITICAL | |
| Instance:SystemMaintenance.Reboot | An instance restart due to system maintenance is being executed. | Maintenance | Executing | CRITICAL | |
| Instance:SystemMaintenance.Reboot | An instance restart due to system maintenance is completed. | Maintenance | Executed | CRITICAL | |
| Instance:SystemMaintenance.Reboot | An instance restart due to system maintenance is canceled. | Maintenance | Canceled | CRITICAL | |
| Instance:SystemMaintenance.Reboot | An instance restart due to system maintenance has failed. | Maintenance | Failed | CRITICAL | |
| Disk:Stalled | Beginning of disk performance impact. | Exception | Executing | CRITICAL | |

| Name | Description | Type | Status | Level | Remarks |
| --- | --- | --- | --- | --- | --- |
| Disk:Stalled | End of disk performance impact. | Exception | Executed | CRITICAL | |
| Instance:StateChange | Notification of instance status change. | StatusNotification | Normal | INFO | Event details describe instance status, including Running, Stopped, Expired, Expires Soon, Locked, Stopping, To Be Released, and Deleted (the instance has been released). For more information about the instance lifecycle, see Instance lifecycle. |

| Name | Description | Type | Status | Level | Remarks |
| --- | --- | --- | --- | --- | --- |
| Instance:PreemptibleInstanceInterruption | Notification of preemptible instance interruption. | StatusNotification | Normal | WARN | Preemptible instances change to the To Be Released status due to various reasons. For example, the market price is higher than your bid, or the relationship between resource supplies and demand changes. For more information, see Preemptible instances. |
| Snapshot:CreateSnapshotCompleted | The disk snapshot is created. | StatusNotification | Normal | INFO | |

Server Load Balancer system events

| Name | Description | 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 |

| Name | Description | Level |
| --- | --- | --- |
| 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 | Description | 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 China (Hong Kong) and overseas regions by default. No bucket-level throttling threshold is configured by default. |

| Name | Description | 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 China (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 | Description | 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 China (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 | Description | 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 China (Hong Kong) and overseas regions by default. No bucket-level throttling threshold is configured by default. |

Auto Scaling system events

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| AUTOSCALING:SCALE_IN_ERROR | The scale-in activity of the scaling group failed. | 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. | Unnormal | CRITICAL |
| AUTOSCALING:SCALE_OUT_SUCCESS | The scale-out activity of the scaling group was successful. | Normal | INFO |

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| AUTOSCALING: SCALE_REJECT | The scaling activity of the scaling group was rejected. | Warn | WARN |
| AUTOSCALING :SCHEDULE_T ASK_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 Platform system events

| Name | Description | Type | Status | Level |
| --- | --- | --- | --- | --- |
| Account_Co nnect_QPS_ Limit | The number of connection requests per second for the current account has reached the upper limit. | Exception | Fail | WARN |
| Account_Do wnlink_QPS _Limit | The number of requests per second that the current account sent to devices has reached the upper limit. | Exception | Fail | WARN |

| Name | Description | Type | Status | Level |
| --- | --- | --- | --- | --- |
| Account_RuleEngine_DataForward_QPS_Limit | The number of requests per second that the current account sent to the rule engine has reached the upper limit. | Exception | Fail | WARN |
| Account_Uplink_QPS_Limit | The number of requests per second that the current account published has reached the upper limit. | Exception | Fail | WARN |
| Device_Downlink_QPS_Limit | The downstream QPS of any device has reached the upper limit. | Exception | Fail | WARN |
| Device_Uplink_QPS_Limit | The upstream QPS of any device has reached the upper limit. | Exception | Fail | WARN |

Smart Access Gateway system events

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| AccessGatewayFailover | The access point failed over. | Agwfailover | INFO |
| Connection Disconnect | The network is disconnected. | Disconnect | CRITICAL |
| DeviceHacked | The device is under attack. | Hacked | CRITICAL |
| DeviceOffline | The device is offline. | Offline | CRITICAL |

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| DeviceOnline | The device is online. | Online | INFO |

CloudMonitor system events

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| Group_AddResourcesFailed_QuotaReached | Failed to add instances to a group in real time because the resource usage has reached the upper limit. | Failed | CRITICAL |
| Agent_Status_Stopped | Heartbeat check failed. | Stopped | CRITICAL |
| Agent_Status_Running | Heartbeat check is resumed. | Running | CRITICAL |

Database Backup system events

| Name | Description | Status | 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 backup plan is suspended. | Failed | INFO |

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| InstanceStart | The backup plan has started. | Running | INFO |
| OpenContBackup | Incremental backup is enabled. | Running | INFO |

Relational Database Service system events

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| Instance_Failover | Instance failover occurred. | Executed | WARN |
| Instance_Failure_Start | Beginning of an instance failure. | Executing | CRITICAL |
| Instance_Failure_End | End of an instance failure. | Executed | CRITICAL |

ApsaraDB for Redis system events

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| Instance_Failover | Instance failover occurred. | Executed | WARN |
| Instance_Failure_Start | Beginning of an instance failure. | Executing | CRITICAL |
| Instance_Failure_End | End of an instance failure. | Executed | CRITICAL |

ApsaraDB for MongoDB system events

| Name | Description | Status | Level |
| --- | --- | --- | --- |
| Instance_Failure_Start | Beginning of an instance failure. | Executing | CRITICAL |
| Instance_Failure_End | End of an instance failure. | Executed | CRITICAL |

AnalyticDB events

| Name | Description | Level |
| --- | --- | --- |
| StorageUsage | The disk usage has exceeded 80%. | CRITICAL |

| Name | Description | Level |
| --- | --- | --- |
| InsertFailureRate | The insert failure rate is 10 %. | CRITICAL |
| SelectFailureRate | The query failure rate is 10 %. | CRITICAL |

Edge Node Service event

| Name | Description | Type | Status | Level |
| --- | --- | --- | --- | --- |
| EnsRegion: NetworkDown: Executing | The edge node failed. | Exception | Executing | CRITICAL |
| EnsRegion: NetworkDown: Executed | The edge node recovered. | Exception | Executed | CRITICAL |
| EnsRegion: NetworkMigration: Scheduled | Network cutover scheduled for the edge node. | Maintenance | Scheduled | WARN |
| EnsRegion: NetworkMigration: Executing | Network cutover is being executed for the edge node. | Maintenance | Executing | CRITICAL |
| EnsRegion: NetworkMigration: Executed | Network cutover is completed for the edge node. | Maintenance | Executed | INFO |
| EnsRegion: NetworkMigration: Canceled | Network cutover is canceled for the edge node. | Maintenance | Canceled | INFO |
| Instance: SystemFailure.Reboot: Executing | The instance is being restarted due to system errors. | Exception | Executing | CRITICAL |

| Name | Description | Type | Status | Level |
| --- | --- | --- | --- | --- |
| Instance:SystemFailure.Reboot:Executed | The instance has been restarted due to system errors. | Exception | 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.

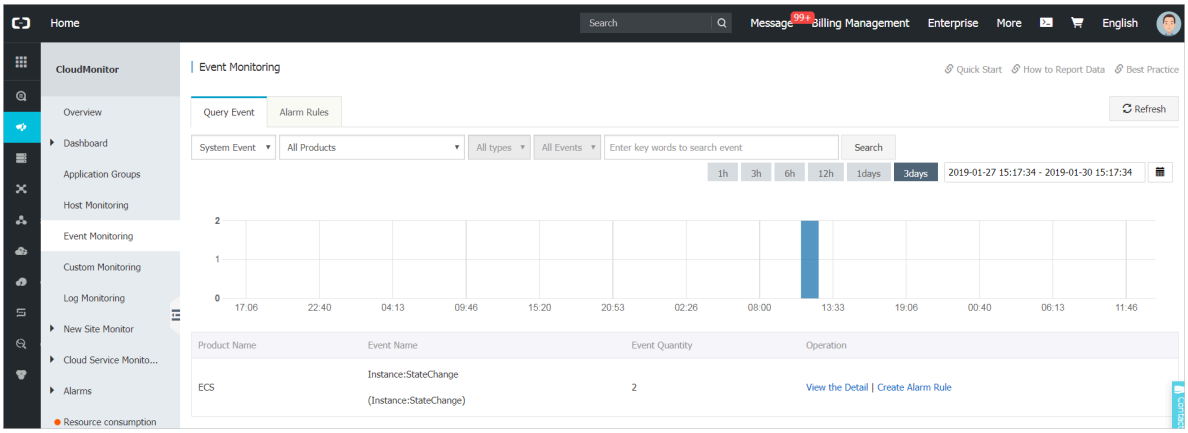

3. Click View Details in the Actions column corresponding to an event to view the details of the event.

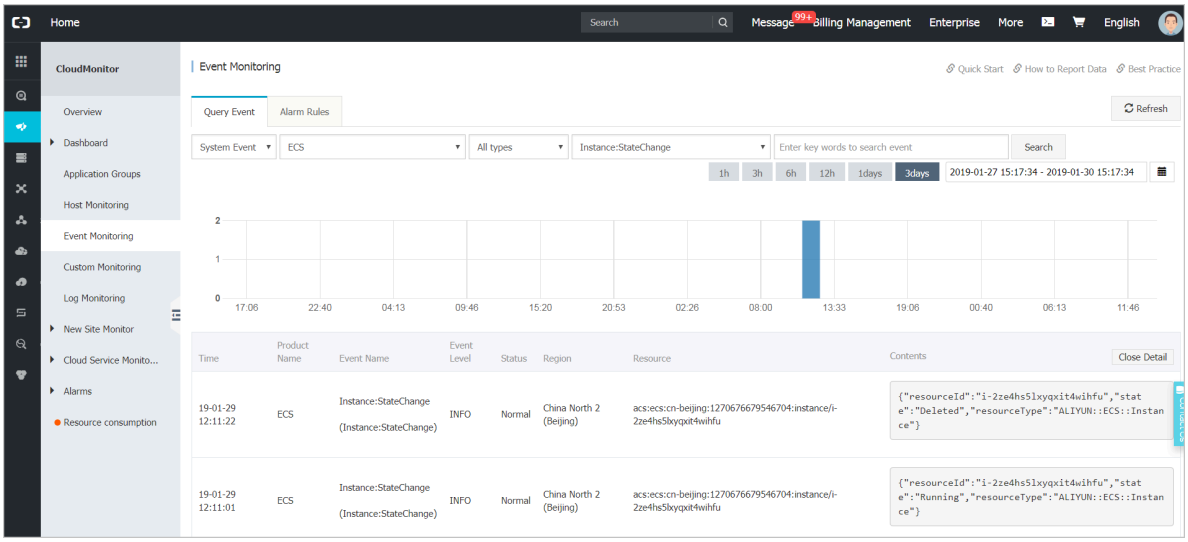


### View system events by group

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.

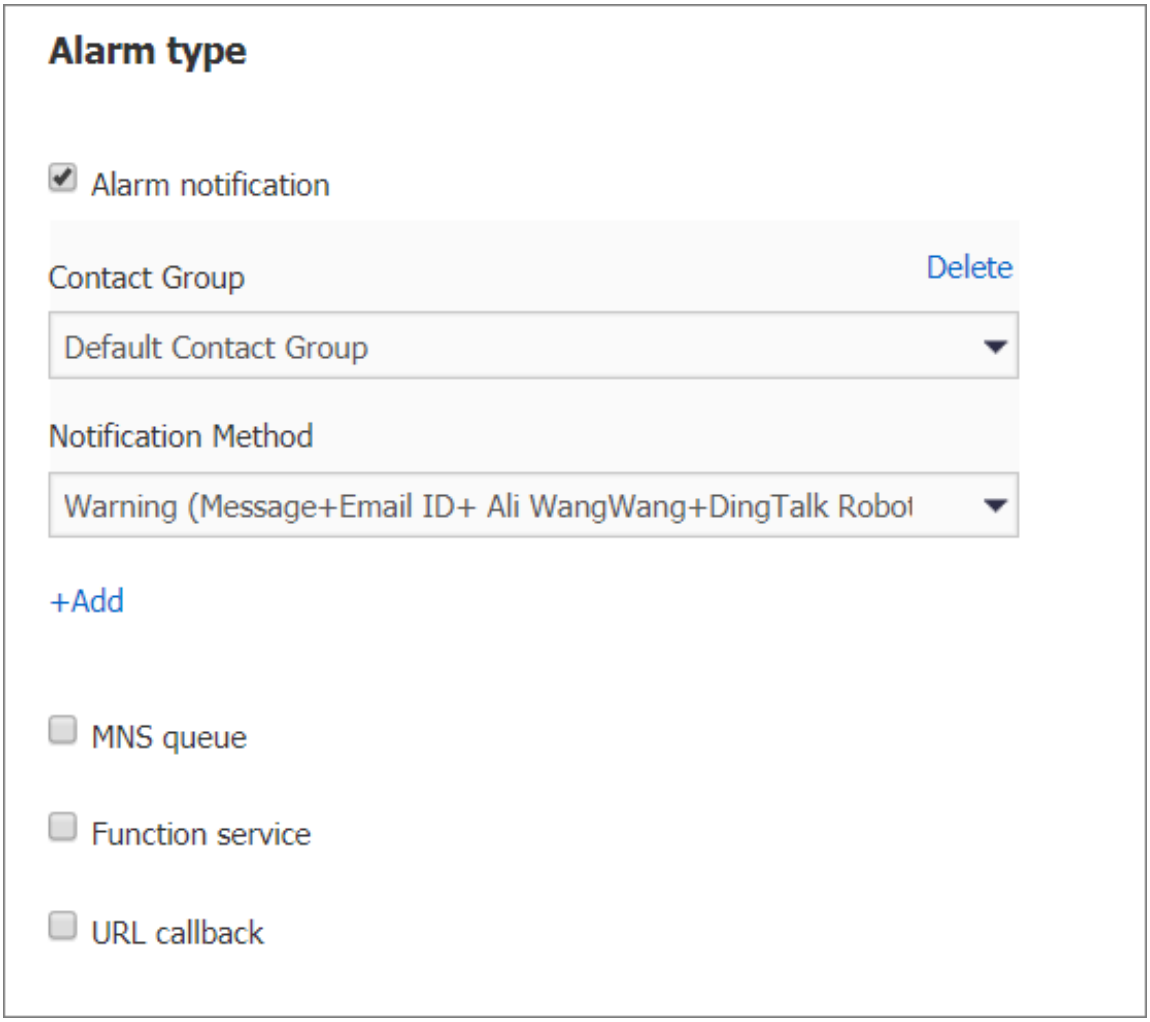

- Test an alert rule

1. Access the Alarm Rules tab for event monitoring.

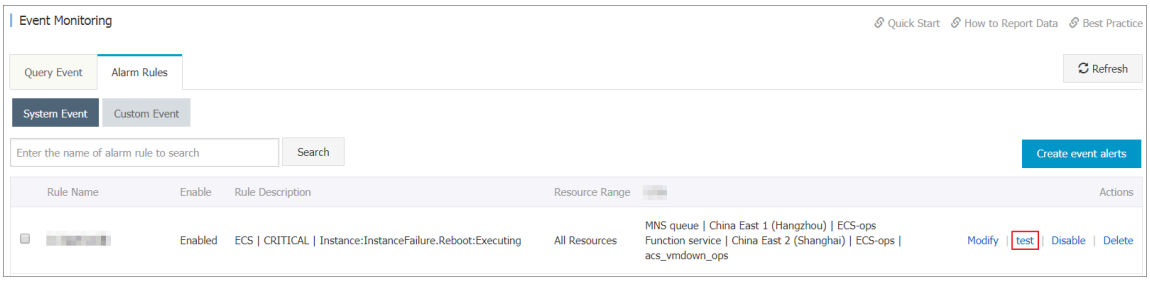


2. Click Test in the Actions column corresponding to the RAM user.
3. On the test page that appears, select the event to be tested. The correspond ing event content will be displayed. You can change content fields such as the instance ID as needed.
4. Click OK. The system will send an event based on the content, triggering alert notification, MNS queue, Function Compute, and URL callback that are configured in the alert rule.

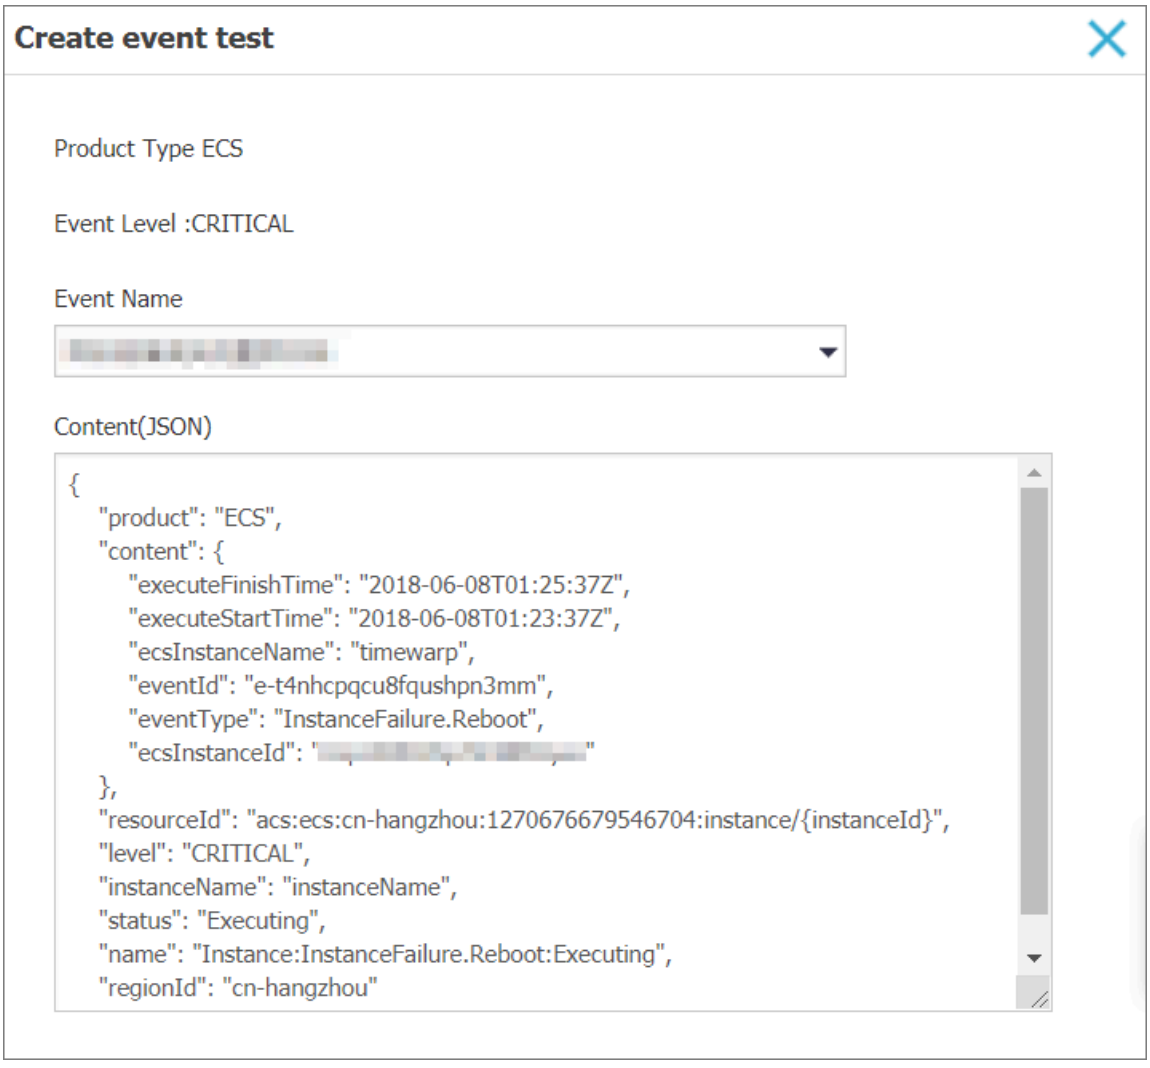

# 9.3 Custom events

## 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
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
[{" content ":" EventConte  nt "," groupId ": GroupId ," name ":"
 EventName "," time ":" 20171023T1  44439 . 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 "
            +  Canonicali  zedHeaders  + "\ n "
            +  Canonicali  zedResourc  e
```

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_STRI  NG  = " 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 ),  AccessKeyS  ecret ))
```

· 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
Authorizat  ion : YourAccKey : YourAccSec  ret
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
[{" 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 . openservic  es </ groupId >
    < artifactId > aliyun - cms </ artifactId >
    < version > 0 . 1 . 2 </ version >
```

```
</ dependency >
```

- **Sample code**

```
public   void   uploadEven  t ()  throws   CMSExcepti  on ,
Interrupte  dException  {
       //  Initialize   the   client .
        CMSClient   cmsClient  =  new   CMSClient ( endpoint ,
accKey ,  secret );
      // Construct   two   events   to   be   reported .
         CustomEven  tUploadReq  uest   request  =  CustomEven
tUploadReq  uest . builder ()
                  . append ( CustomEven  t . builder ()
                         . setContent (" abc , 123 ")
                         . setGroupId ( 101l )
                         . setName (" Event001 "). build ())
                  . append ( CustomEven  t . builder ()
                         . setContent (" abc , 123 ")
                         . setGroupId ( 101l )
                         . setName (" Event002 "). build ())
                  . build ();
            CustomEven  tUploadRes  ponse   response  =  cmsClient
. putCustomE  vent ( request );
            List < CustomEven  t >  eventList  =  new   ArrayList <
CustomEven  t >();
            eventList . add ( CustomEven  t . builder ()
                  . setContent (" abcd , 1234 ")
                  . setGroupId ( 101l )
                  . setName (" Event001 "). build ());
            eventList . add ( CustomEven  t . builder ()
                  . setContent (" abcd , 1234 ")
                  . setGroupId ( 101l )
                  . setName (" Event002 "). build ());
            request  =  CustomEven  tUploadReq  uest . builder ()
                  . setEventLi  st ( eventList ). build ();
            response  =  cmsClient . putCustomE  vent ( 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.**

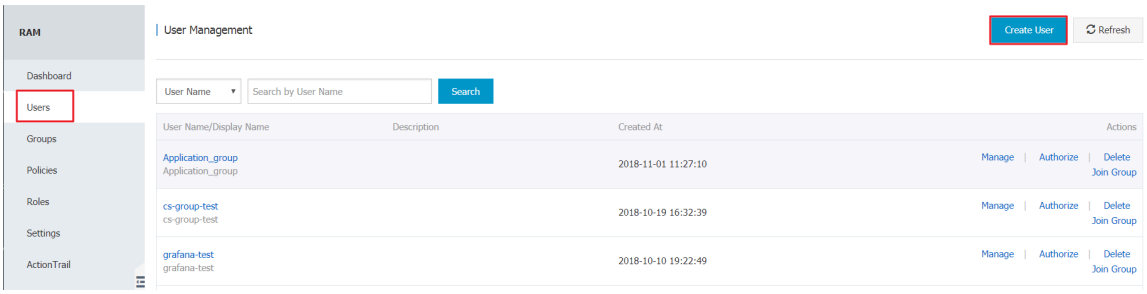


   b. **Generate an AccessKey ID and an AccessKey Secret for the RAM user.**

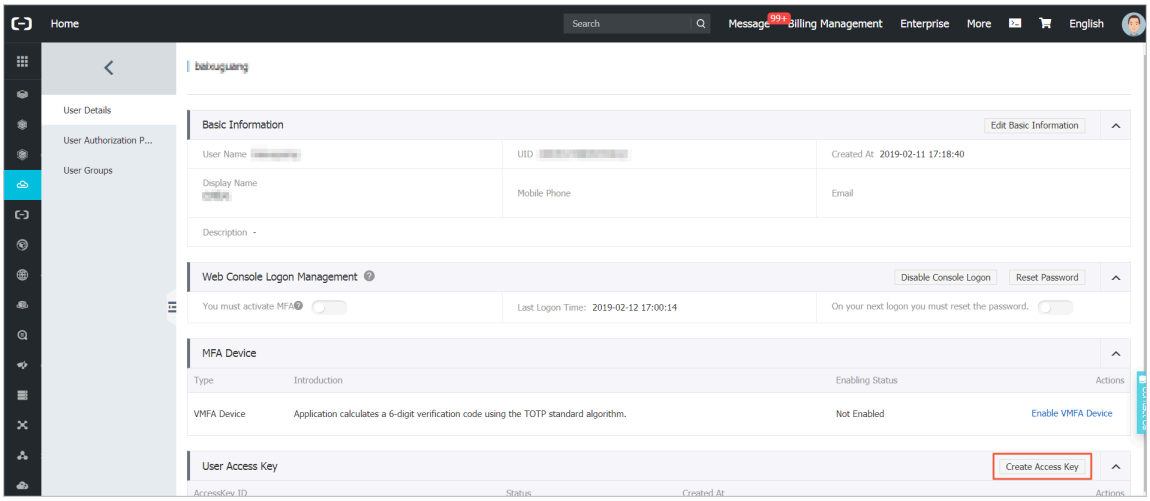


   c. **Grant CloudMonitor permissions to the RAM user.**

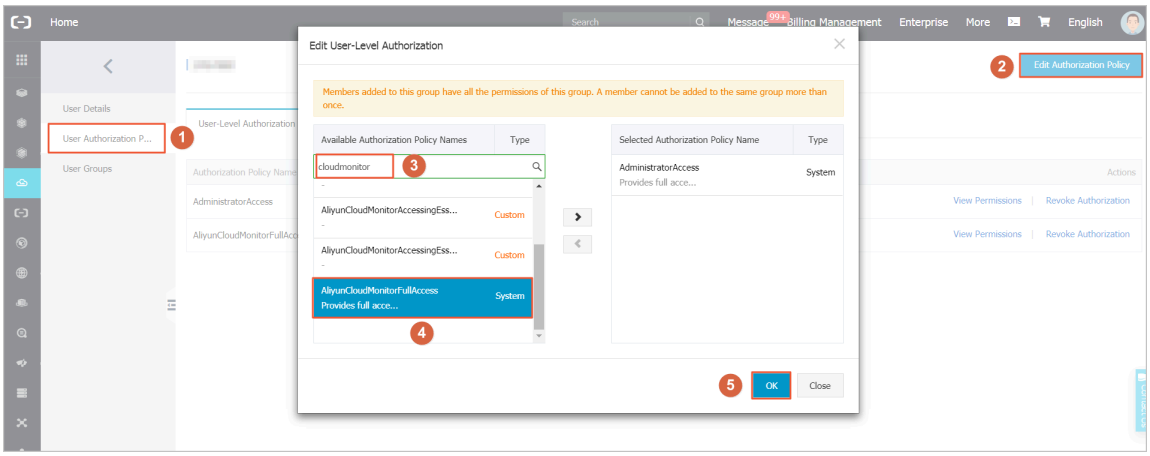

2. Install CMS SDK

   · 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 permission s, 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.

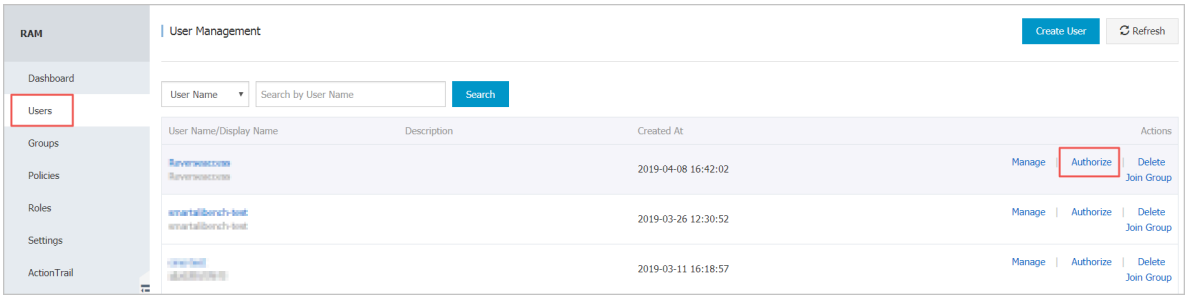


4. On the authorization page, select AliyunCloudMonitorFullAccess and click OK.

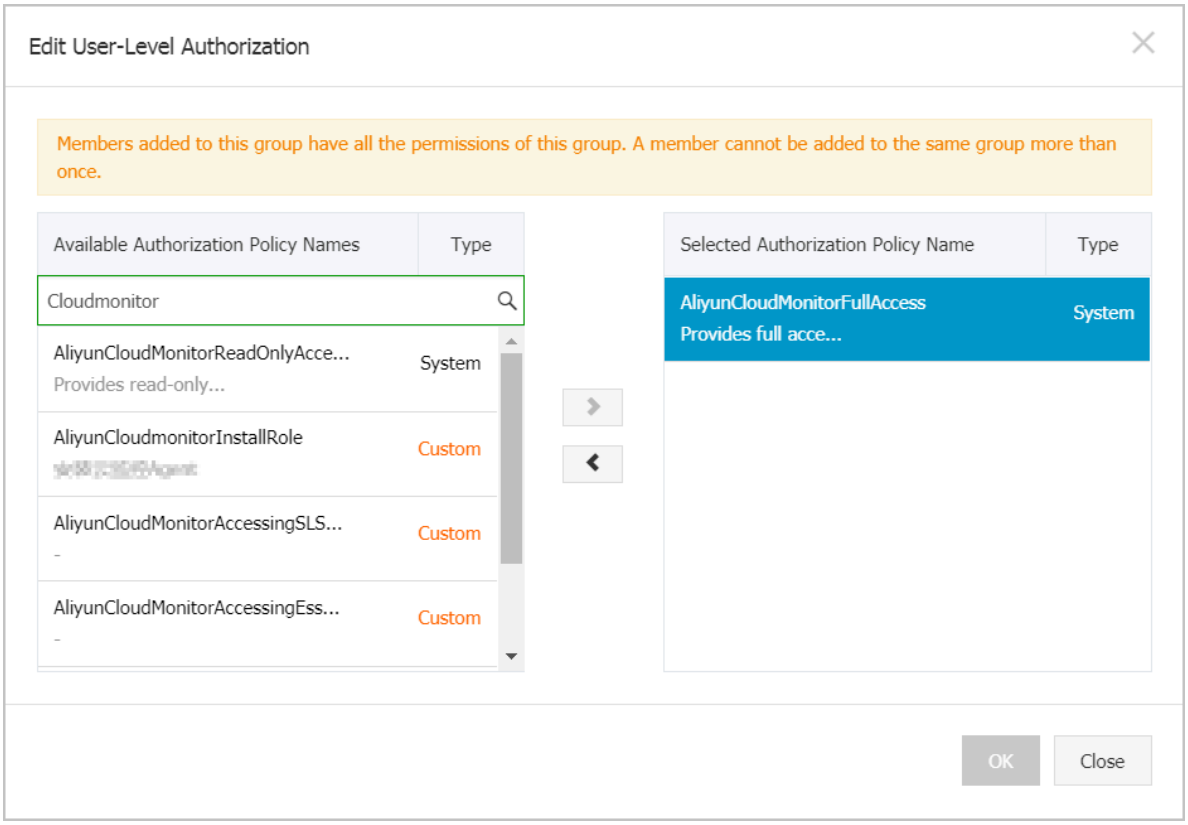

## 9.3.2 View custom events

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 troublesho oting

### 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.

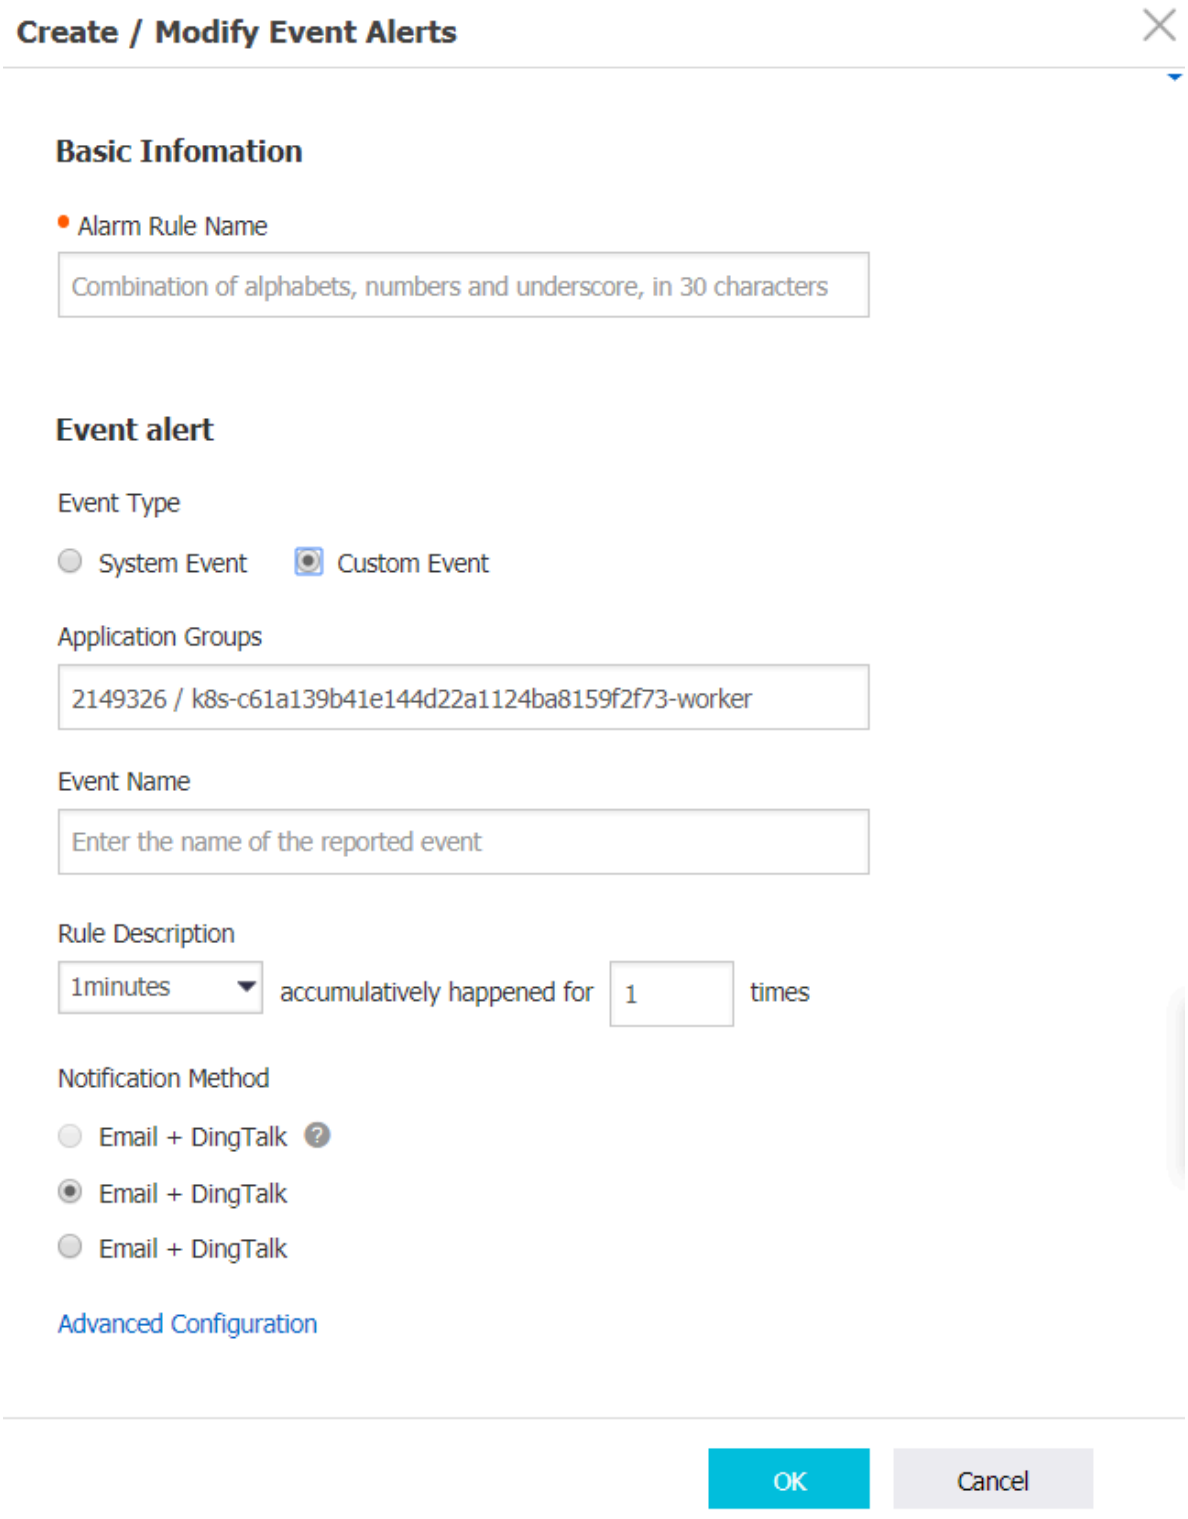


4. In the Basic Information area, enter a name for the alarm rule.

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 :// metrichub - 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   identifica  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   asynchrono  usly .  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 < CustomEven  t >  events  =  new   ArrayList <
 CustomEven  t >();
     for  ( int   i  =  0 ;  i  <  99 ;  i ++) {
         EventEntry   e  =  this . eventQueue . poll ();
         if  ( e  ==  null ) {
             break ;

         events . add ( CustomEven  t . builder (). setContent ( e
 . getContent ()). setName ( e . getName ()). build ());

     if  ( events . isEmpty ()) {
         return ;

     //  Report   events   in   batch   to   CloudMonit  or .  No
   retry   or   retry   in   SDK   is   added   here .  If   you
 have   high   requiremen  t   for   event   reliabilit  y ,  add
   retry   policies .
     try  {
         CustomEven  tUploadReq  uestBuilde  r   builder  =
 CustomEven  tUploadReq  uest . builder ();
         builder . setEventLi  st ( events );
         CustomEven  tUploadRes  ponse   response  =  cmsClient .
 putCustomE  vent ( builder . build ());
         if  (!"  200 ". equals ( response . getErrorCo  de ())) {
             logger . warn (" event   reporting   error :  msg
 : {},  rid : {}",  response . getErrorMs  g (),  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 informativ e for
searching and locating . Here , map is used for
organizing events and converted to Json format
as event content .
Map < String , String > eventConte nt = new HashMap <
String , String >();
eventConte nt . put (" method ", " GET "); // http request
  method
eventConte nt . put (" path ", "/ users "); // http path
eventConte nt . put (" exception ", e . getClass (). getName
()); // Exception class name for searching
eventConte nt . put (" error ", e . getMessage ()); //
Error message of exception
eventConte nt . put (" stack_trac e ", ExceptionU tils
. getStackTr ace ( e )); // Exception stack for
locating
// Finally submit the events in the preceding
asynchrono us reporting method . Since no retry
is performed in asynchrono us reporting , event
loss of small probabilit y may happen . However
, it is sufficient for alarms of unknown http
exceptions .
put (" http_error ", JsonUtils . toJson ( eventConte nt ));
image . png ]( http :// ata2 - img . cn - hangzhou . img - pub
. aliyun - inc . com / 864cf09597 7cf61bd340 dd1461a024 7c
. 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 organizati on of the message queue :
```

```
 Map < String ,  String >  eventConte  nt  =  new   HashMap <
 String ,  String >();
 eventConte  nt . put (" cid ",  consumerId ); //  Consumer
 ID
 eventConte  nt . put (" mid ",  msg . getMsgId ()); //
 Message   ID
 eventConte  nt . put (" topic ",  msg . getTopic ()); //
 Message   topic
 eventConte  nt . put (" body ",  body ); //  Message   body
 eventConte  nt . put (" reconsume_  times ",  String . valueOf
 ( msg . getReconsu  meTimes ())); //  The   number   of
 retries   after   message   failure
 eventConte  nt . put (" exception ",  e . getClass (). getName
 ()); // Exception   class   name   in   case   of   exception
 eventConte  nt . put (" error ",  e . getMessage ()); //
 Exception   message
 eventConte  nt . put (" stack_trac  e ",  ExceptionU  tils .
 getStackTr  ace ( e )); //  Exception   stack
//  Finally ,  report   the   event
```

```
put (" metaq_erro  r ",  JsonUtils . toJson ( eventConte  nt
));
```

Check the event after reporting:

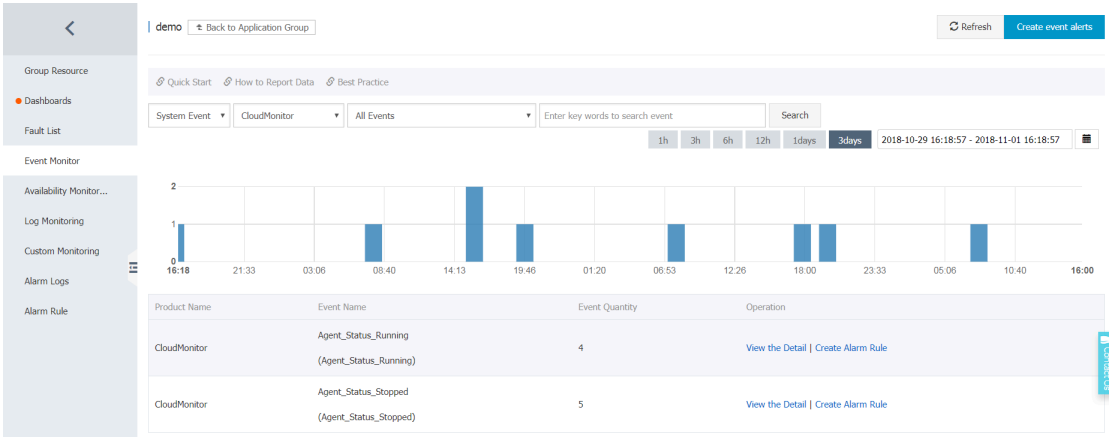


- Set alarms for queue message consumption exceptions:

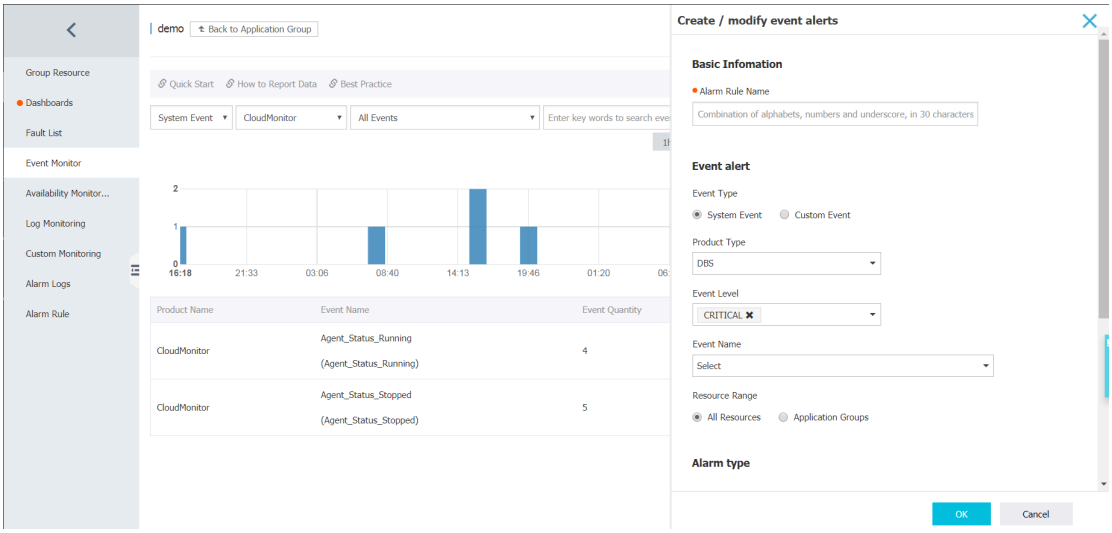


- Demo 3: Record important events

  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.

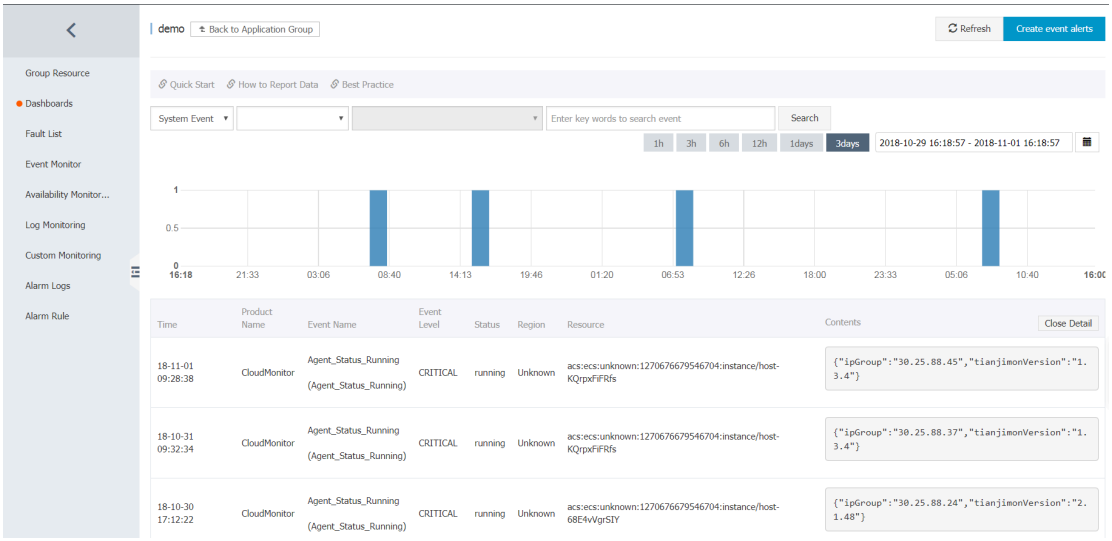


## 9.4 Request header definitions

Request headers for the event monitoring interface are defined as follows.

| Header | Type | Description |
| --- | --- | --- |
| Authorization | String | Content: AccessKeyId: SignString |
| User-Agent | String | Client descriptions |
| Content-MD5 | String | The string produced as an MD5 hash for a request body, which appears in all uppercase letters. If a request has no body, this request header is not required. |
| Content-Length | Value | The length of an HTTP request body as defined in RFC 2616. If a request has no body, this request header is not required. |
| Content-Type | String | Only application/json is supported |

| Header | Type | Description |
| --- | --- | --- |
| Date | String | The standard timestamp header in an HTTP request (following the RFC 1123 format and using GMT standard time): Mon, 3 Jan 2010 08:33:47 GMT. GMT |
| Host | String | The full Host name of an HTTP request (not including protocol headers such as https ://): metrichub-cms-cn-hangzhou.aliyuncs.com |
| x-cms-api-version | String | API v1.0 |
| x-cms-signature | String | Signature algorithm: HMAC-SHA1. |
| x-cms-ip | String | The IP for reporting events : 10.1.1.1 |

# 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.

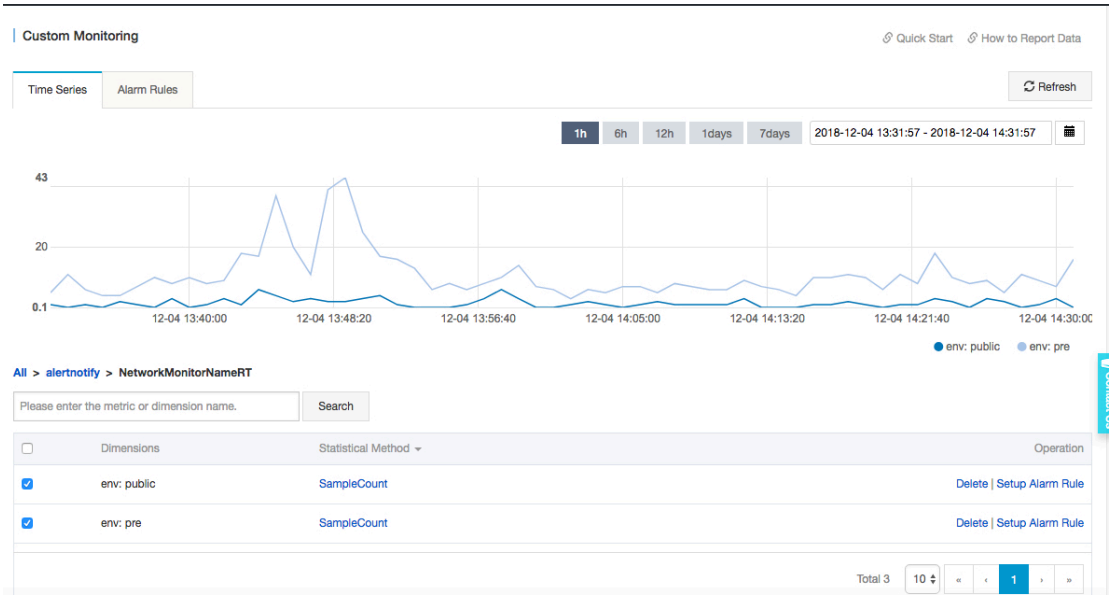


  - To view the custom monitoring data in an application group, complete the following steps:

    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.

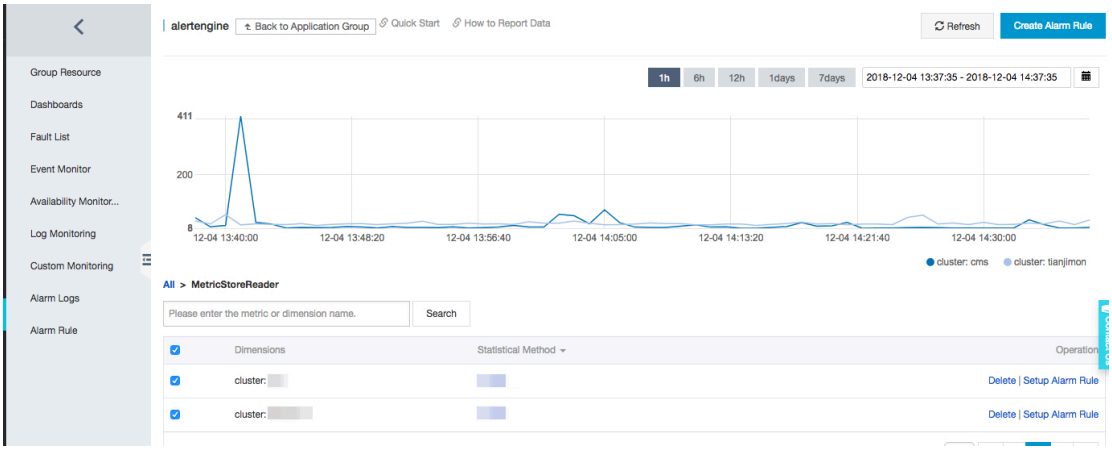


- Set an alarm rule.

  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

Custom monitoring allows you to customize metrics and alarm rules to meet your business requirements. Custom monitoring provides API operations for reporting monitoring data. You can use the API operations to report collected time series data to CloudMonitor (CMS). You can also configure alarm rules to receive notifications of corresponding exceptions.

CloudMonitor provides API operations, Java SDKs, and Alibaba Cloud command-line interface (CLI) for reporting data.

### Limits

- The upper limit of queries per second (QPS) is 200 QPS in China (Beijing), China (Shanghai), and China (Hangzhou), 100 QPS in China (Zhangjiakou) and China ( Shenzhen), and 50 QPS in all other regions.
- The system reports a maximum of 100 data entries at one time. The body size is 256 KB or less.
- The `metricName` field can only contain letters, digits, and underscores (_). This field must start with a letter. If the starting character is not a letter, this character 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` is 64 bytes or fewer. Otherwise, the key or value string is truncated.
- You have to pay for reporting raw data. You can obtain aggregated data free of charge. To obtain the free data, set the Type field to 1 in the request parameters.

### Report data by using API operations

After you report the raw data by using API operations, 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.
- P10: the value of the 10th percentile. This value is greater than 10% of all data in the aggregation period.
- P20: the value of the 20th percentile. This value is greater than 20% of all data in the aggregation period.
- P30: the value of the 30th percentile. This value is greater than 30% of all data in the aggregation period.
- P40: the value of the 40th percentile. This value is greater than 40% of all data in the aggregation period.
- P50: the value of the 50th percentile. This value is a median value and greater than 50% of all data in the aggregation period.
- P60: the value of the 60th percentile. This value is greater than 60% of all data in the aggregation period.
- P70: the value of the 70th percentile. This value is greater than 70% of all data in the aggregation period.
- P75: the value of the 75th percentile. This value is greater than 75% of all data in the aggregation period.
- P80: the value of the 80th percentile. This value is greater than 80% of all data in the aggregation period.
- P90: the value of the 90th percentile. This value is greater than 90% of all data in the aggregation period.
- P95: the value of the 95th percentile. This value is greater than 95% of all data in the aggregation period.
- P98: the value of the 98th percentile. This value is greater than 98% of all data in the aggregation period.
- P99: the value of the 99th percentile. This value 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-b | 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 |
| China (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 ": 0 ," metricName ":" diskUtiliz ation "," dimensions
":{" instanceId ":" xxxxxx1 "," disk ":"/"}," time ":" 20190701T1
2345 . 888 + 0800 "," type ": 0 ," period ": 60 ," values ":{"
value ": 60 }}]
```

- Signature algorithm

  CloudMonitor only supports the HMAC-SHA1 signature algorithm.

  1. Prepare an Alibaba Cloud AccessKey pair.

     To generate a digital signature for an API request, you must use an AccessKey pair that consists of the AccessKey ID and 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 "
           + Canonicali zedHeaders + "\ n "
```

```
                + Canonicali zedResourc e
```

In this formula, `\ n` indicates the newline escape character and the plus sign (`+`) indicates the string concatenation operator. The other parts are defined as follows:

| Parameter | Description | 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. This value 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. This timestamp header follows the RFC 1123 specification and uses GMT standard time. | Mon, 3 Jan 2010 08:33:47 GMT |
| CanonicalizedHeaders | The string constructed by the custom headers that are prefixed with x-cms and x-acs of the HTTP request. | x-cms-api-version:0.1.0\nx-cms-signature |

| Parameter | Description | Examples |
| --- | --- | --- |
| CanonicalizedResource | The string constructed by the HTTP request resources, as described in the following section. | /event/custom/upload |

The CanonicalizedHeaders string in the preceding table is constructed as follows:

a. Convert the names of all HTTP request headers that are prefixed with `x - cms` and `x - acs` to lowercase letters.
b. Sort the CMS custom request headers generated in the preceding step in lexicographic order.
c. Delete any space on either side of a delimiter between a request header and the corresponding content.
d. Separate all headers and content with separators (`\ n`) to form the final CanonicalizedHeaders string.

The CanonicalizedResource string 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 in `QUERY_STRI NG`, add a question mark (`?`) and the query string to the end of the CanonicalizedResource string.

   In the request, `QUERY_STRI NG` is the lexicographically sorted 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-

parameter value pairs are then sorted in lexicographic order and connected with ampersands (&) to form a string. The formula is as follows:

```
QUERY_STRI  NG  = " KEY1 = VALUE1 " + "&" + " KEY2 = VALUE2 "
```

3. Generate a digital signature for the request.

   The default signature algorithm is HMAC-SHA1. You can use the following formula to generate a signature:

```
Signature  =  base16 ( hmac - sha1 ( UTF8 - Encoding - Of (
SignString ),  AccessKeyS  ecret ))
```

4. Signature example

```
 SignString =" POST " + \ n
+" 0B9BE351E5  6C90FED853  B32524253E  8B " + \ n
+" applicatio  n / json " + \ n
+" Tue ,  11   Dec   2018   21 : 05 : 51  + 0800 " + \ n
+" x - cms - api - version : 1 . 0 " + \ n
+" x - cms - ip : 127 . 0 . 0 . 1 " + \ n
+" x - cms - signature : hmac - sha1 " + \ n
+"/ metric / custom / upload "

 accesskey =" testkey "
 accessSecr  et =" testsecret "  Signature   key

 Signature   result :  1DC19ED63F  755ACDE203  614C8A1157
 EB1097E922
```

- Request parameters

| Parameter | Type | Required | Description |
| --- | --- | --- | --- |
| groupId | long | Yes | The ID of your application group. |

| Parameter | Type | Required | Description |
| --- | --- | --- | --- |
| metricName | string | Yes | The name of a metric that you want to monitor. 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 are truncated from the string. |
| dimensions | object | Yes | The dimension map. All key-value pairs are strings. A string can contain letters, digits, and connectors such as underscores (_), hyphens (-), periods (.), forward slashes (/), and backslashes (\). 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 are truncated from the string. |

| Parameter | Type | Required | Description |
| --- | --- | --- | --- |
| time | string | Yes | The time when the metric value was generated. The time supports timestamps in the yyyyMMdd ’T’ HHmmss.SSSZ format or long format, such as 20171012T132456.888+0800 or 1508136760000. |
| type | int | Yes | The data type of the reported value. A value of 0 specifies raw data and a value of 1 specifies aggregate data.<br>When you report aggregate data, we recommend that you report data for both 60s and 300s aggregation periods. Otherwise, you cannot query monitoring data in a time span that is more than seven days. |

| Parameter | Type | Required | Description |
| --- | --- | --- | --- |
| period | string | No | The aggregation period. Unit: seconds.<br>If the type parameter is set to 1, this field is required. Valid values: 60 and 300. |
| values | object | Yes | The collection of metric values. If the type parameter is set to 0, the key of this parameter must be set to value, so raw data is reported. CloudMonitor aggregates raw data over the aggregation period into several data types, such as maximum, count, and sum. |

Report data by using the Java SDK (Recommended)

- Install the Java SDK

  When you install the Java SDK based on Maven, add the following dependency:

```
< 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 code (Java)**

```
import   com . aliyuncs . DefaultAcs  Client ;
import   com . aliyuncs . IAcsClient ;
import   com . aliyuncs . exceptions . ClientExce  ption ;
import   com . aliyuncs . exceptions . ServerExce  ption ;
import   com . aliyuncs . profile . DefaultPro  file ;
import   com . google . gson . Gson ;
import   java . util .*;
import   com . aliyuncs . cms . model . v20190101 .*;

public   class   PutCustomM  etric  {

    public   static   void   main ( String []  args ) {
        DefaultPro  file   profile  =  DefaultPro  file
. getProfile (" cn - hangzhou ", "< accessKeyI  d >", "<
accessSecr  et >");
        IAcsClient   client  =  new   DefaultAcs  Client (
profile );

        PutCustomM  etricReque  st   request  =  new
PutCustomM  etricReque  st ();

        List < PutCustomM  etricReque  st . MetricList >
metricList  List  =  new   ArrayList < PutCustomM  etricReque  st
. MetricList >();

        PutCustomM  etricReque  st . MetricList   metricList  1
=  new   PutCustomM  etricReque  st . MetricList ();
        metricList  1 . setGroupId (" 0 ");
        metricList  1 . setMetricN  ame (" diskUtiliz  ation ");
        metricList  1 . setDimensi  ons ("{\" hostName \":\"
xxxxx \",\" disk \":\"/\"}");
        metricList  1 . setTime (" 20190612T1  32456 . 888
0800Z ");
        metricList  1 . setType (" 0 ");
        metricList  1 . setPeriod (" 60 ");
        metricList  1 . setValues ("{\" value \": 20 }");
        metricList  List . add ( metricList  1 );
        request . setMetricL  ists ( metricList  List );

        try  {
            PutCustomM  etricRespo  nse   response  =  client .
getAcsResp  onse ( request );
            System . out . println ( new   Gson (). toJson (
response ));
       }  catch  ( ServerExce  ption   e ) {
            e . printStack  Trace ();
       }  catch  ( ClientExce  ption   e ) {
            System . out . println (" ErrCode :" +  e .
getErrCode ());
            System . out . println (" ErrMsg :" +  e . getErrMsg
());
            System . out . println (" RequestId :" +  e .
getRequest  Id ());
       }

   }
```

```
}
```

- **Sample code (Golang)**

```
 package   main

 import  (
    " fmt "
      " github . com / aliyun / alibaba - cloud - sdk - go /
 services / cms "

)

 func   main () {
     client ,  err  :=  cms . NewClientW  ithAccessK  ey (" cn -
 hangzhou ", "< accessKeyI  d >", "< accessSecr  et >")

     request  :=  cms . CreatePutC  ustomMetri  cRequest ()
     request . Scheme  = " https "

   request . MetricList  = &[] cms . PutCustomM  etricMetri  cList
 {
    {
       GroupId : " 0 ",
       MetricName : " diskUtiliz  ation ",
       Dimensions : "{" hostName ":" xxxxx "," disk ":"/"}",
       Time : " 20190612T1  32456 . 888   0800Z ",
       Type : " 0 ",
       Period : " 60 ",
       Values : "{" value ": 20 }",
    },
  }

     response ,  err  :=  client . PutCustomM  etric ( request )
     if   err  ! =  nil  {
         fmt . Print ( err . Error ())
    }
     fmt . Printf (" response   is  %# v \ n ",  response )
}
```

- **Sample response**

```
{
    " Message ": " success ",
    " RequestId ": " E25EE651 - 9C97 - 4EFD - AF22 - A753B674E8
 D4 ",
    " Code ": " 200 "
}
```

- Sample code for other programming languages

- **Automatically report aggregate data for multiple aggregation periods**

  **The Java SDK supports data reporting after local aggregation. Aggregation periods can be 60 seconds and 300 seconds.**

| Data type | Description | Aggregate value | Memory consumption ( excluding the name, dimension , individual time series, and individual aggregation periods) |
|---|---|---|---|
| value | Typical value type | All attributes except LastValue | Approximately 4 KB |
| gauge | Sample value | LastValue | 4 bytes |
| meter | Sum and rate | Sum, SumPerSecond | 50 bytes |
| counter | Count | SampleCount | 10 bytes |
| timer | Computing time | SampleCount, CountPerSecond, Average, Maximum , Minimum, and PXX(P10-P99) | Approximately 4 KB |
| histogram | Distribution | SampleCount, Average, Maximum , Minimum, and PXX(P10-P99) | Approximately 4 KB |

```
// Initializa  tion
        CMSClientI  nit . groupId  =  0L ;
        CMSClient   cmsClient  =  new   CMSClient ( accKey ,
 secret ,  endpoint );// Create   a   client .
        CMSMetricR  egistryBui  lder   builder  =  new
 CMSMetricR  egistryBui  lder ();
        builder . setCmsClie  nt ( cmsClient );
        final   MetricRegi  stry   registry  =  builder . build
 ();// Create   a   registry   that   includes   two   aggregatio  n
   periods .
       // Or   final   MetricRegi  stry   registry  =  builder .
 build ( RecordLeve  l .  _60S );// Create   a   registry   that
   only   includes   aggregate   data   in   a   1 - minute
 aggregatio  n   period .
```

```
// Use   value .
 ValueWrapp  er   value  =  registry . value ( MetricName . build ("
 value "));
 value . update ( 6 . 5 );
// Use   meter .
 MeterWrapp  er   meter  =  registry . meter ( MetricName . build ("
 meter "));
 meter . update ( 7 . 2 );
// Use   counter .
 CounterWra  pper   counter  =  registry . counter ( MetricName .
 build (" counter "));
 counter . inc ( 20 );
 counter . dec ( 5 );
// Use   timer .
 TimerWrapp  er   timer  =  registry . timer ( MetricName . build ("
 timer "));
 timer . update ( 30 ,  TimeUnit . MILLISECON  DS );
// Use   histogram .
 HistogramW  rapper   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 by using Alibaba Cloud CLI

#### Prepare an Alibaba Cloud account

Make sure that you have created an Alibaba Cloud account, created a RAM user for your account, and generated an AccessKey pair for the RAM user to grant CloudMonit or permissions.

- Create a RAM user.

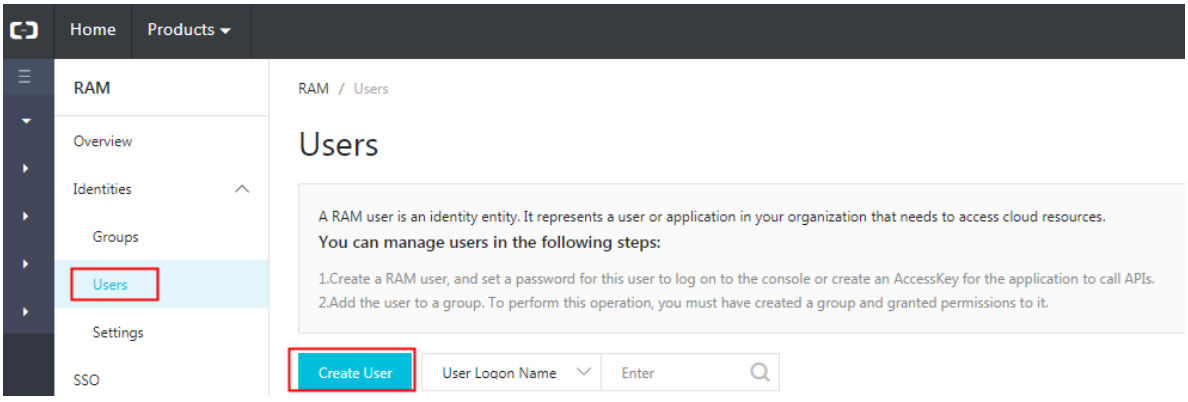

- Generate an AccessKey ID and an AccessKey Secret for the RAM user.

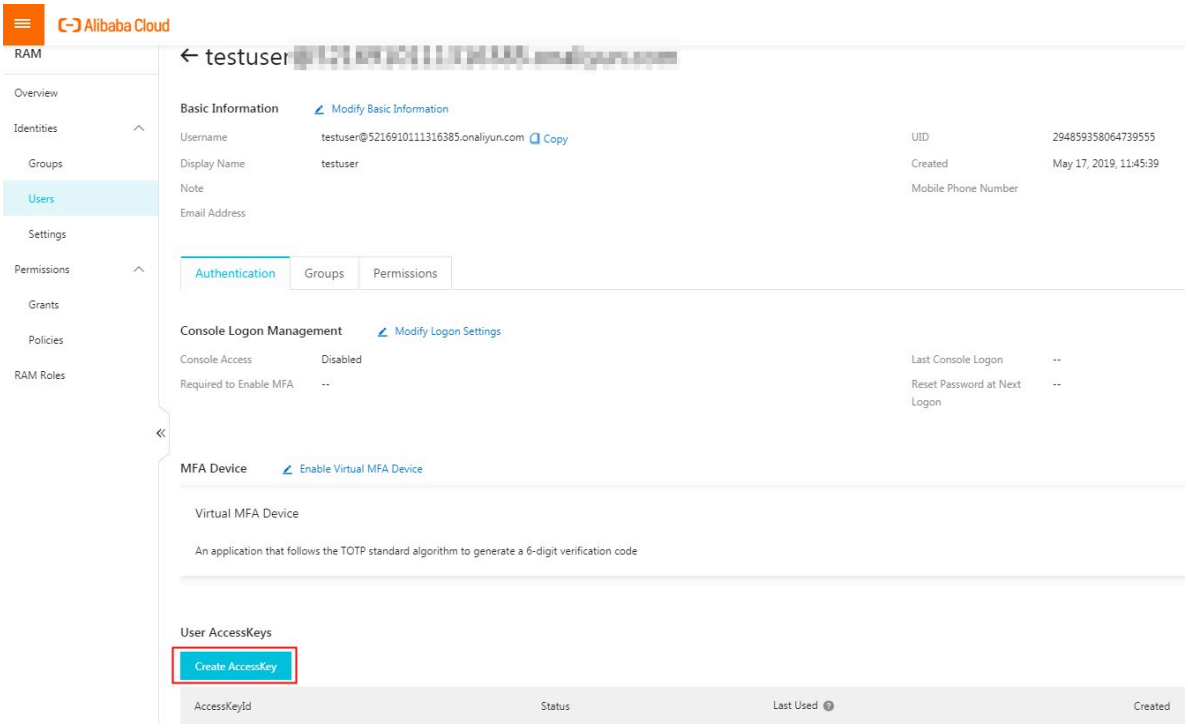


- Grant CloudMonitor permissions to the RAM user.

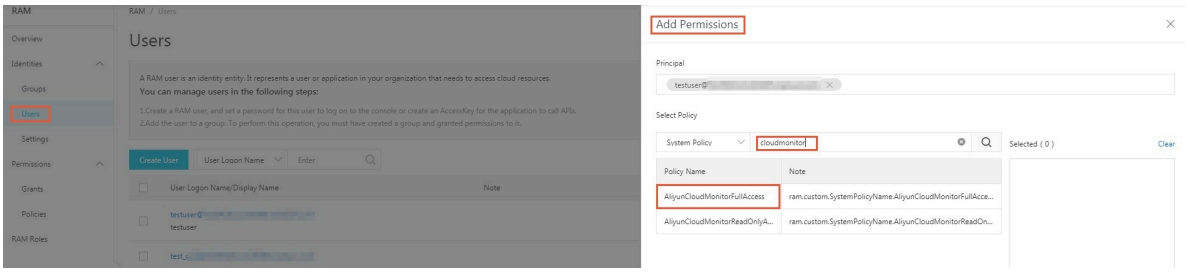


## Install Alibaba Cloud CLI

Required operating systems: Linux, Unix, or macOS.

- Method 1: download the installation package

  You can download the installation package of the latest CLI tool from Alibaba Cloud CLI GitHub, and decompress the package to use the CLI. The CLI supports Mac, Linux, and Windows (x64) terminals. After decompression, you can move the aliyun file to the /usr/local/bin directory or add the file to the $PATH environment variable.

- Method 2: compile source code

  Install and configure the Golang environment, and follow these steps to download and compile the source code:

  ```
  ```

```
$  mkdir  - p  $ GOPATH / src / github . com / aliyun
$  cd  $ GOPATH / src / github . com / aliyun
$  git   clone   http :// github . com / aliyun / aliyun - cli .
 git
$  git   clone   http :// github . com / aliyun / aliyun - openapi -
 meta . git
$  cd   aliyun - cli
$  make   install
```

**Configure the CLI**

**Before using Alibaba Cloud CLI, you must run the aliyun configure command to configure the AccessKey pair, region, and language. You use the configured data to call Alibaba Cloud resources. You can create and view your AccessKey pair on the AccessKey page in the Alibaba Cloud console, or contact your system administrator to obtain the AccessKey pair.**

```
$  aliyun   configure
 Configurin  g   profile  ' default ' ...
 Aliyun   Access   Key   ID  [ None ]: < Your   AccessKey   ID >
 Aliyun   Access   Key   Secret  [ None ]: < Your   AccessKey   Secret
 >
 Default   Region   Id  [ None ]:  cn - hangzhou
 Default   output   format  [ json ]:  json
 Default   Language  [ zh ]:  zh
```

**Multi-user configuration: Alibaba Cloud CLI supports multi-user configuration. You can run the $ aliyun configure --profile user1 command to specify the account used to call API operations of cloud services. Run the `$ aliyun configure list` command to view the current user configuration, as shown in the following table. The asterisk (*) next to the Profile field value indicates the default user configuration.**

| Profile | Credential | Valid | Region | Language |
|---|---|---|---|---|
| default * | AK:***f9b | Valid | cn-beijing | zh |
| aaa | AK:****** | Invalid | | |
| test | AK:***456 | Valid | | en |
| ecs | EcsRamRole: EcsTest | Valid | cn-beijing | en |

**To specify the authentication method, you can add the --mode <authenticationMethod > parameter to the end of the configure command in the CLI. Supported authentica tion methods are listed as follows:**

| Authentication method | Description |
| --- | --- |
| AK | Use the AccessKey ID and AccessKey Secret to authorize access permissions. |
| StsToken | Use the Security Token Service (STS) token to authorize access permissions. |
| RamRoleArn | Use AssumeRole of the RAM user to authorize access permissions. |
| EcsRamRole | Use EcsRamRole on an ECS instance to access the instance with no password. |

**Report monitoring data**

**Call the `PutCustomMetric` API operation to report the monitoring data.**

```
aliyun   cms   PutCustomM  etric   -- MetricList . 1 . MetricName
  cpu_total  -- MetricList . 1 . Dimensions  '{" sampleName  1 ":"
value1 "," sampleName  2 ":" value2 "}' -- MetricList . 1 . Time
1555390981  421  -- MetricList . 1 . Type   0  -- MetricList . 1 .
Period   60  -- MetricList . 1 . Values  '{" value ": 10 . 5 }' --
MetricList . 1 . GroupId  " 0 "
```

**The system returns Status code 200 to indicate successful reporting.**

```
{
  " Message ": " success ",
  " RequestId ": " F69F5623 - DDD6 - 42AE - AE59 - 87A2B84162  0B ",
  " Code ": " 200 "
}
```

**Error codes**

| Error code | Description |
| --- | --- |
| 200 | The error message returned because you reported data successfully. |

| Error code | Description |
| --- | --- |
| 206 | The error message returned because some data failed to be reported.<br>If the system returns "reach Max time series num", you have used up the time series quota. We recommend that you purchase a higher quota or remove unnecessary time series.<br>If the system returns "not allowed original value, please upgrade service", you have used a free edition. You cannot use this edition to report raw data.<br>If the system returns "type is invalid", the value of the type parameter was invalid. Make sure that the value of this parameter is 0 or 1. |
| 400 | The error message returned because a syntax error has occurred in the client request. |
| 403 | The error message returned because the verification failed, the target rate reached the upper limit, or the target permission was not authorized. |
| 500 | The error message returned because an internal server error has occurred. |

Create a RAM user

To use the AccessKey pair of the corresponding RAM user to report event data, you must grant CloudMonitor permissions to the RAM user. If you have not granted the permissions, the system returns the error "cannot upload data, please use ram to auth " 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. On the Users page that appears, select the RAM user that you use to report data, and click Add Permissions in the Actions column next to the target RAM user.

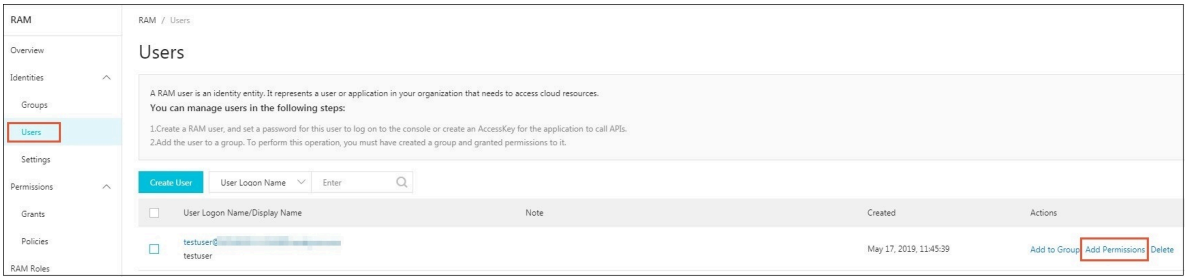


4. On the Add Permissions page, select AliyunCloudMonitorFullAccess and click OK.

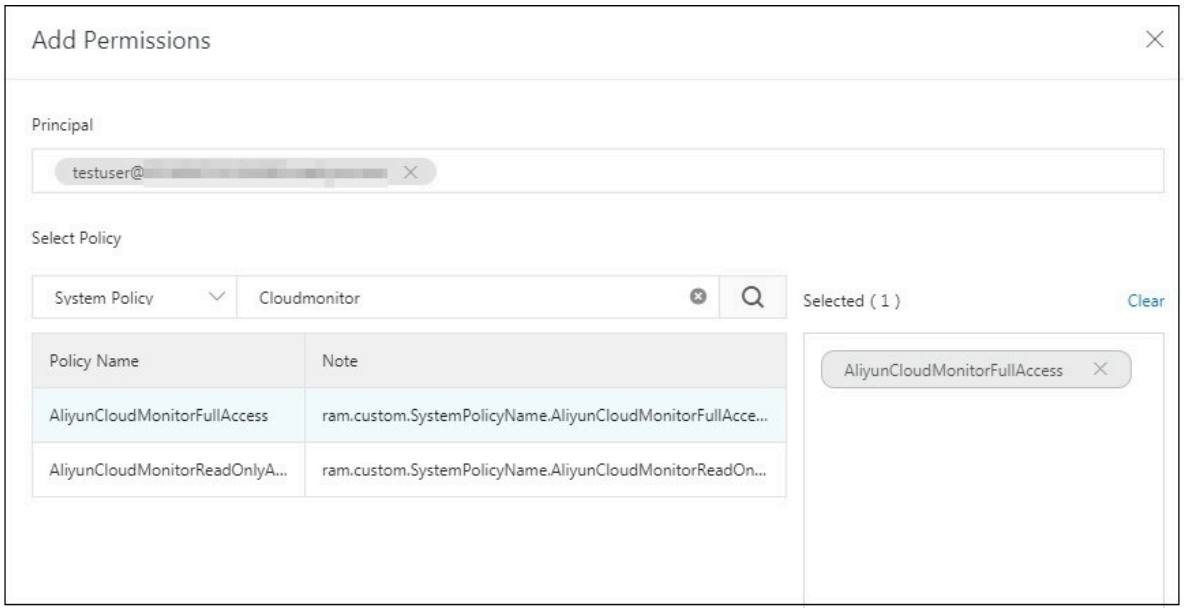


## 10.3 View custom monitoring charts

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.

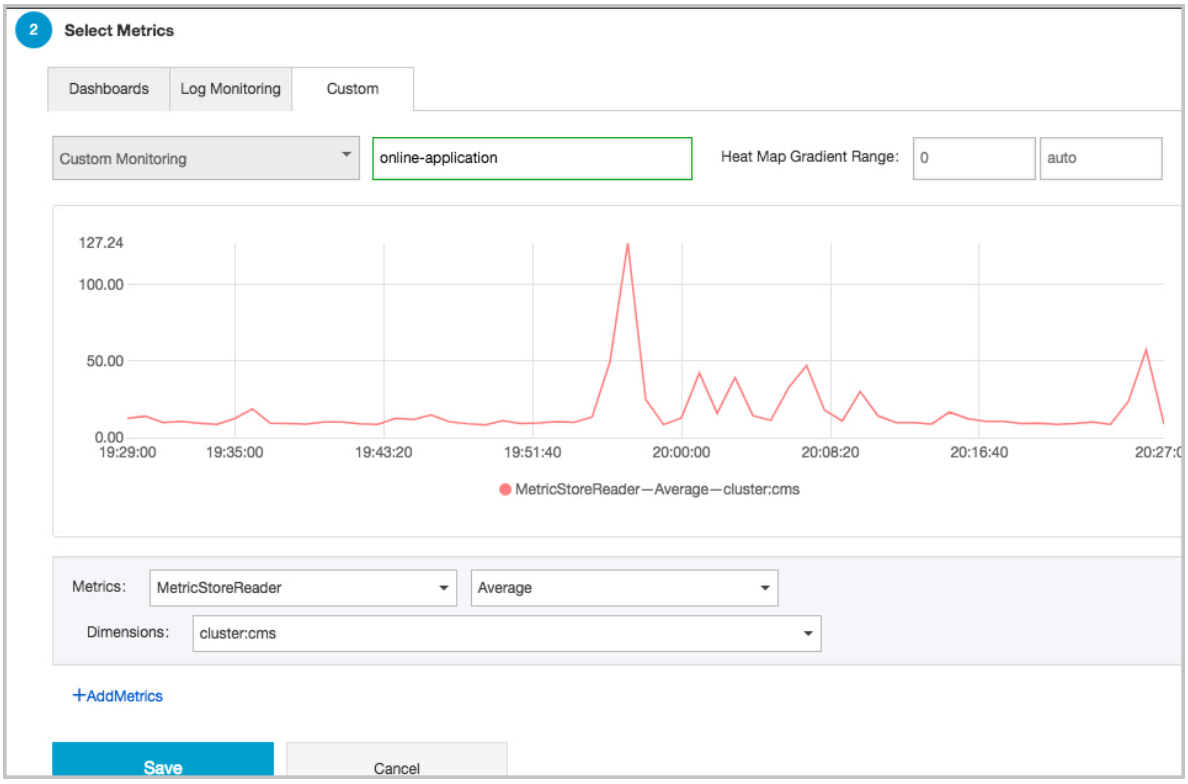


4. On the Custom tab that appears, enter a chart name. Select the metrics, statistical methods, and dimensions to be displayed.

5. Click Save. After saving these configurations, you can view the custom monitoring chart.