ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Enterprise

Apsara Big Data Manager

User Guide

Product Version: v3.16.2

Document Version: 20220915

Alibaba Cloud

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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, 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 this document 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 directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	**Danger:** Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	**Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	**Notice:** If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	**Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.What is Apsara Big Data Manager?

Apsara Big Data Manager (ABM) is an operations and maintenance (O&M) platform tailored for big data services.

ABM supports the following services:

- MaxCompute
- DataWorks
- Realtime Compute
- Quick BI
- DataHub

ABM supports O&M on big data services from the perspectives of business, services, clusters, and hosts. ABM also allows you to update big data services, customize alert configurations, and view the O&M history.

Onsite Apsara Stack engineers can use ABM to easily manage big data services. For example, they can view metrics, check and handle alerts, and modify configurations.

2.Common operations

The data tables and legends in the ABM console facilitate operations. This topic uses MaxCompute and DataHub as examples to describe the common operations.

Search for a project

You can perform a quick search for a project by project name.

1. On the **MaxCompute** page, click **O&M** in the upper-right corner, and then click the **Business** tab. The **Project List** page under **Projects** appears.
2. In the **Project** field, enter a keyword of the project name. Auto-suggestion is supported. Select the target project from the drop-down list, or select the project by using the up and down arrow keys, and then press **Enter**.

> **Note** When a project is matched, the region of the project appears before the project name.

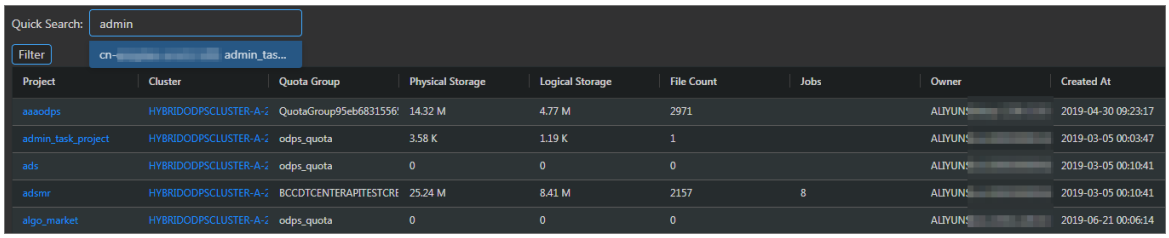

The following figure shows the search result.

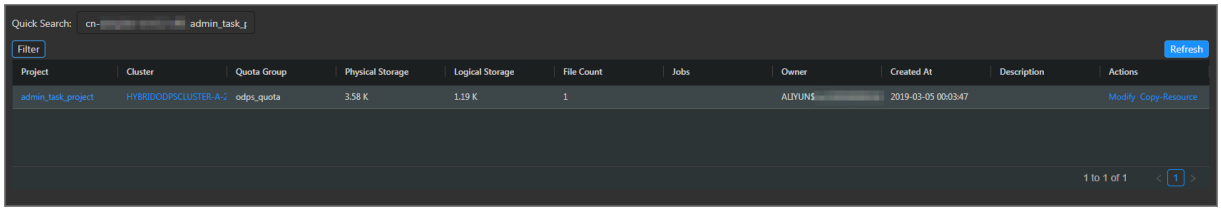

Filter projects

You can set filter conditions for multiple columns at the same time to filter projects and find the target projects.

1. On the **MaxCompute** page, click **O&M** in the upper-right corner, and then click the **Business** tab. The **Project List** page under **Projects** appears.
2. On the **Project List** page, click **Filter** in the upper-left corner of the list. A field for setting filter conditions appears for each column.
3. Click the icon next to each field for setting filter conditions and select the filtering method. The default method is **Contains**.

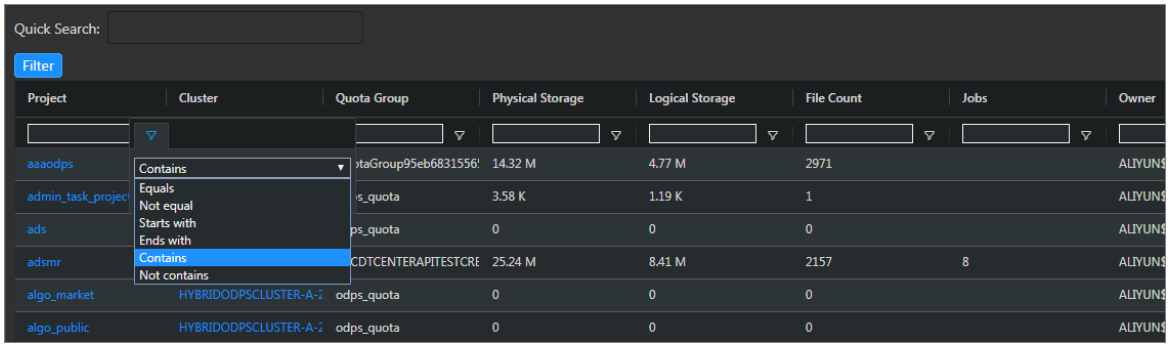

You can select one of the following filtering methods:

- **Equals**
- **Not equal**
- **Starts with**
- **Ends with**
- **Contains**
- **Not contains**

4. After you select the filtering method, enter the filter condition. The projects that meet the filter condition appear.

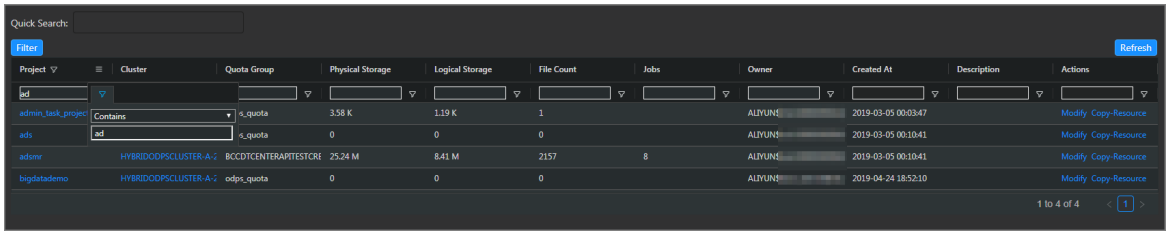

5. If the filtering result is not accurate, you can continue performing this operation on other columns.

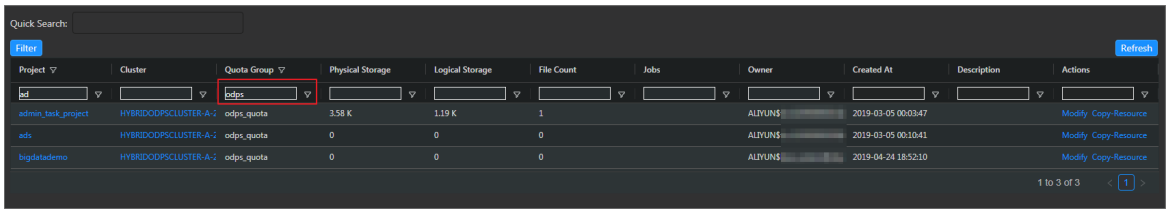

After you set the filter conditions for the projects, the **Filter** button is highlighted. If you need to cancel filtering, click the highlighted **Filter** button.

Search for an item

You can search for an item in a table by column, which is similar to filtering projects. For example, you can perform the following steps to search for a checker:

1. On the **MaxCompute** page, click **O&M** in the upper-right corner, and then click the **Clusters** tab. On the Clusters page, click the **Health Status** tab.
2. In the checker list, click the **Filter** icon in a column and enter a keyword in the search box.

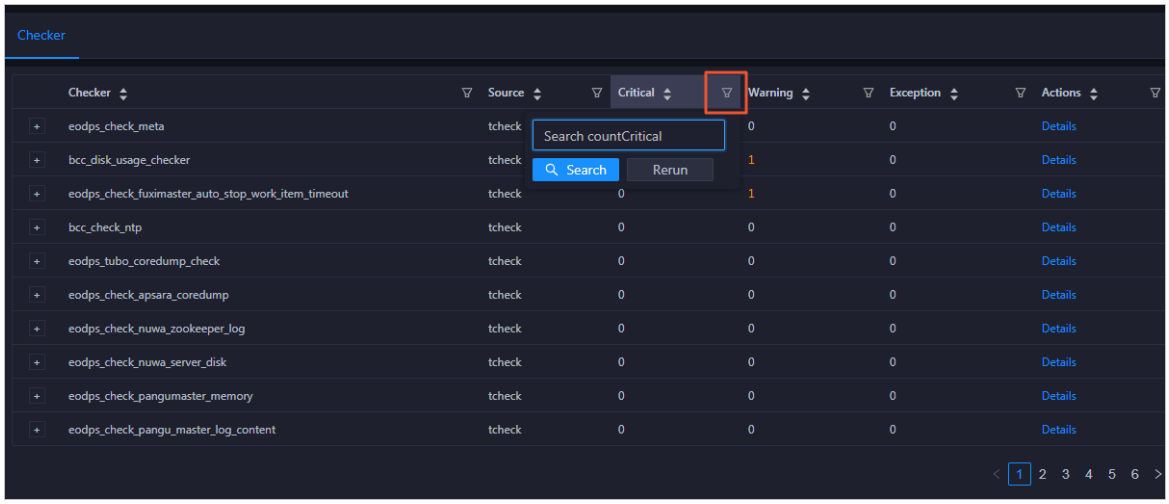

3. Click **Search**. The checkers that meet the requirements appear.
4. If the search result is not accurate, you can continue performing this operation on other columns.

Customize a column

You can customize columns in the list. For example, you can set the column position or column width, and determine whether to display a column. You can also set filter conditions for columns.

On the **Project List** page, you can drag a column to change its position.

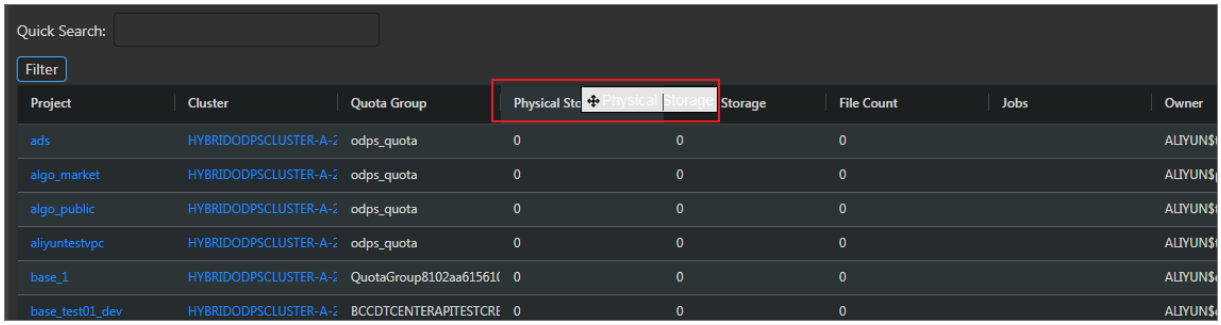

You can click ☰ in a column heading to customize the column.

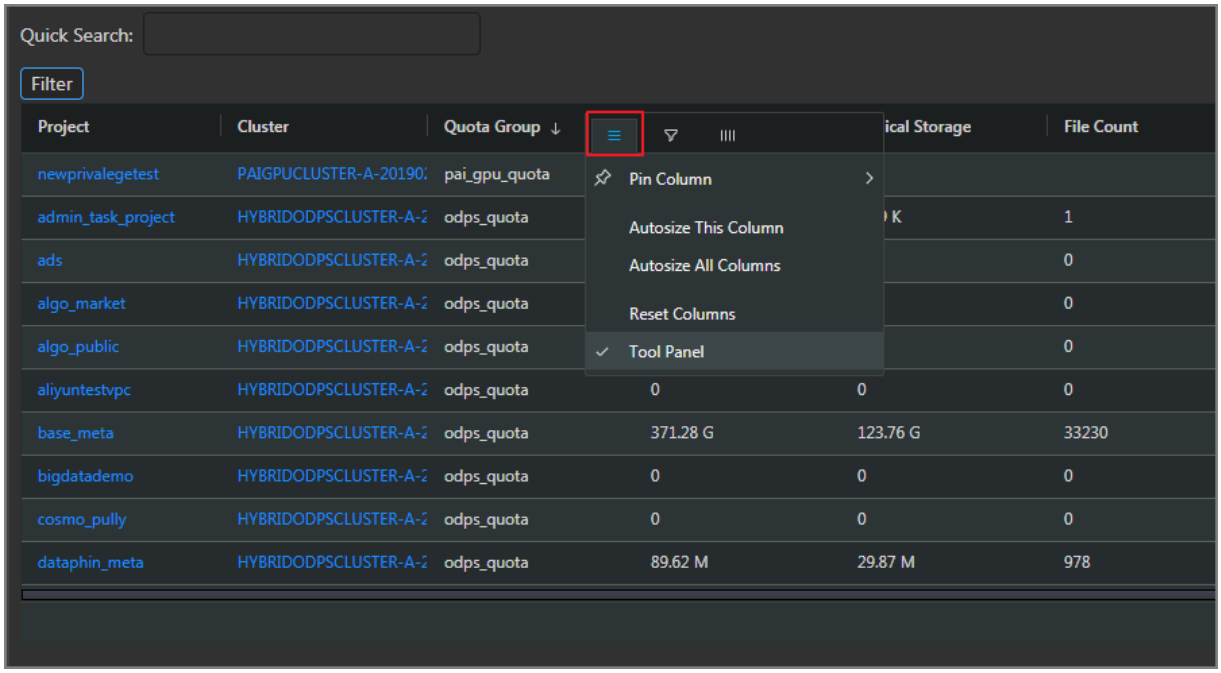

- **Pin Column**: allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.
- **Autosize This Column**: allows you to adjust the width of a column automatically.
- **Autosize All Columns**: allows you to adjust the width of all columns automatically.
- **Reset Columns**: allows you to reset a column to its initial status.
- **Tool Panel**:

Click in a column heading and set a filter condition to filter projects based on the column.

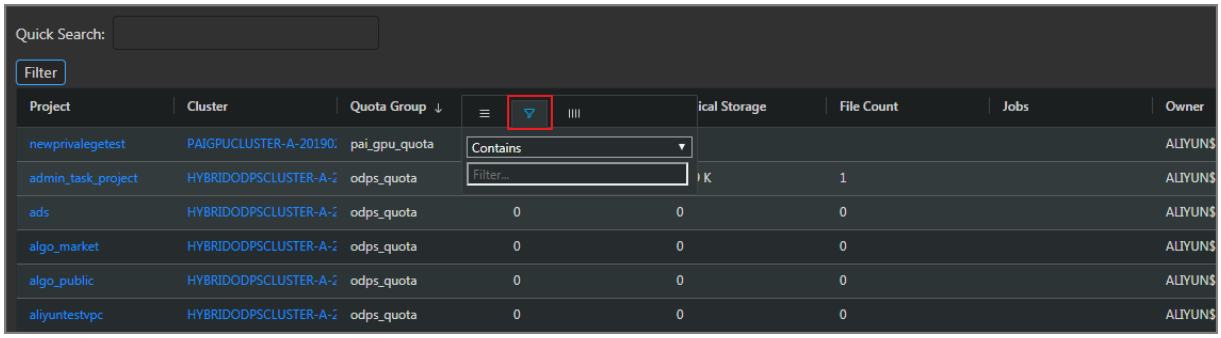

Click in a column heading and select the columns to display.

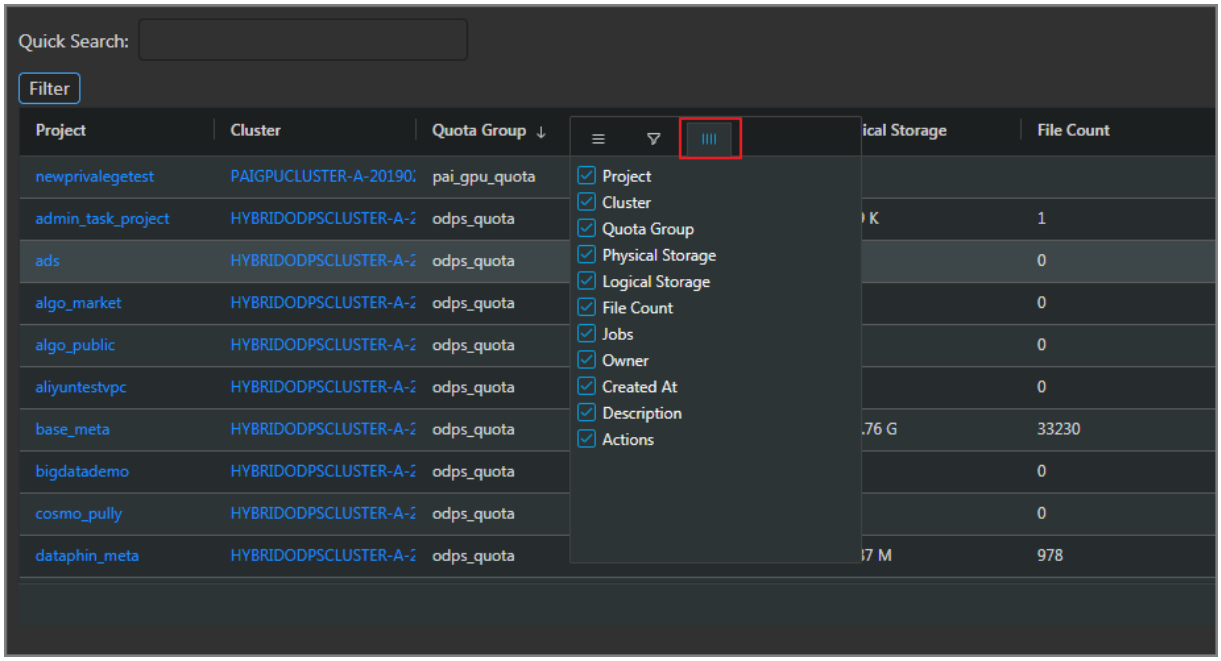

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

After the tool panel appears, it is attached to the right of the list so that you can set the columns to display.

On the **Project List** page, click ≡ in a column heading and select **Tool Panel**. The tool panel is then attached to the right of the list.

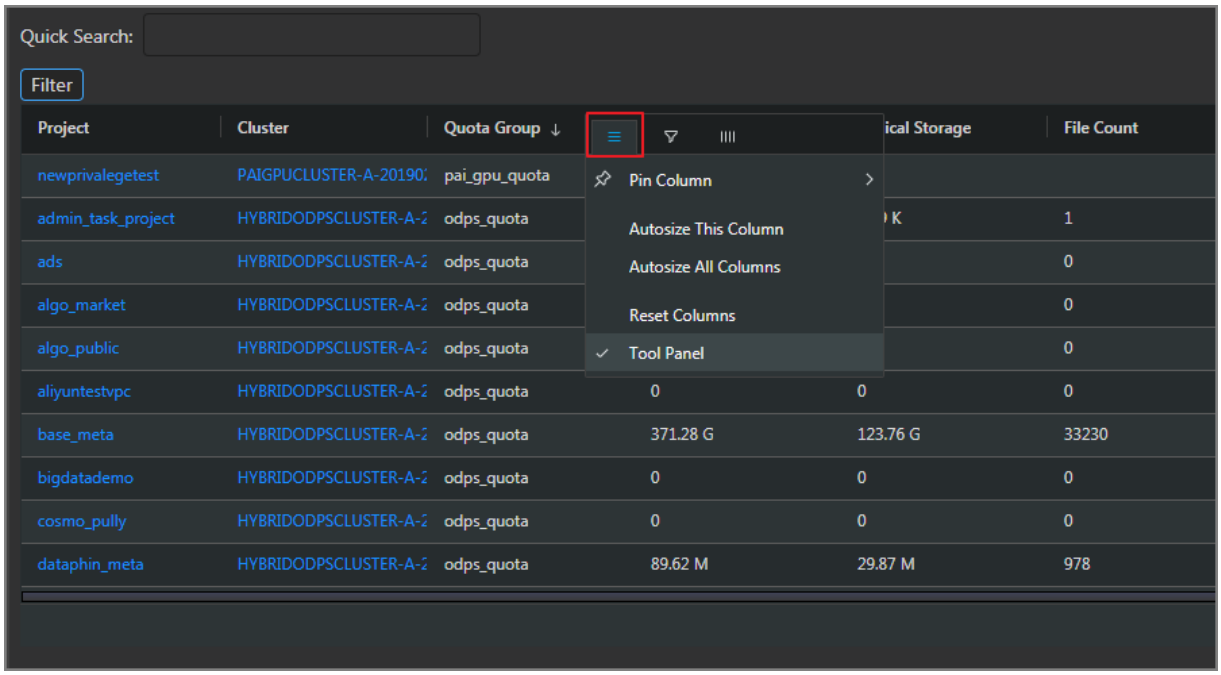

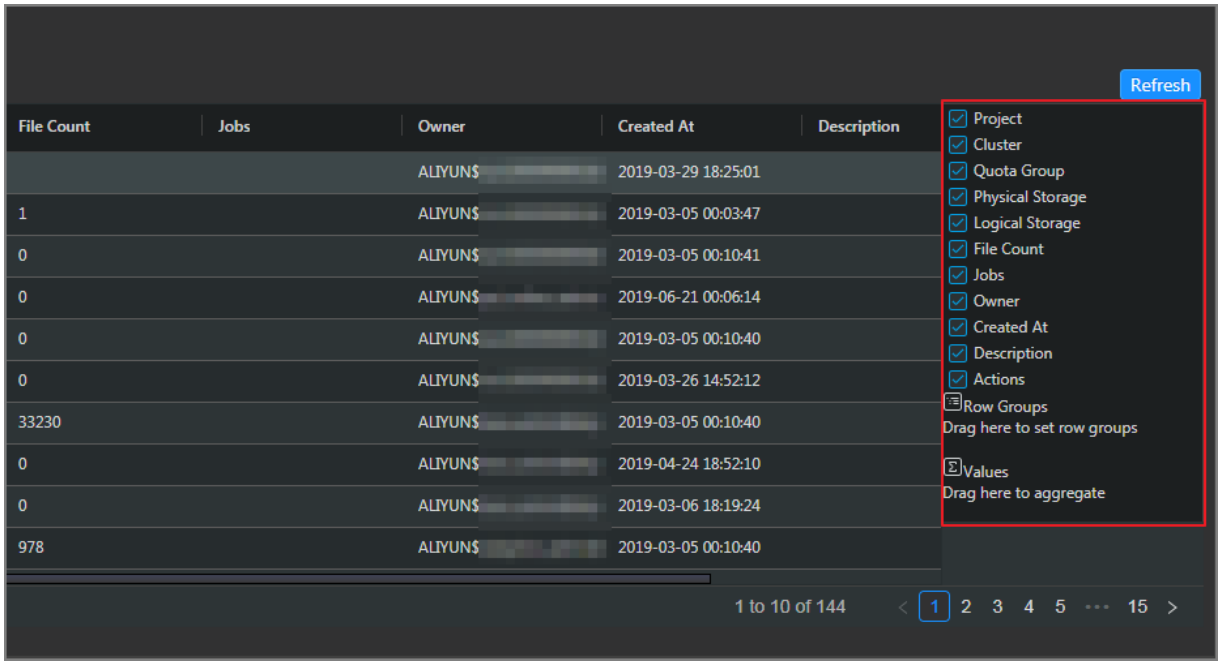

Sort projects based on a column

You can sort projects based on a column in ascending or descending order.

On the **Project List** page, click a column heading in the list. When you click the column heading for the first time, the projects are sorted based on the column in ascending order. When you click the column heading for the second time, the projects are sorted in descending order. When you click the column heading for the third time, the default sorting is restored.

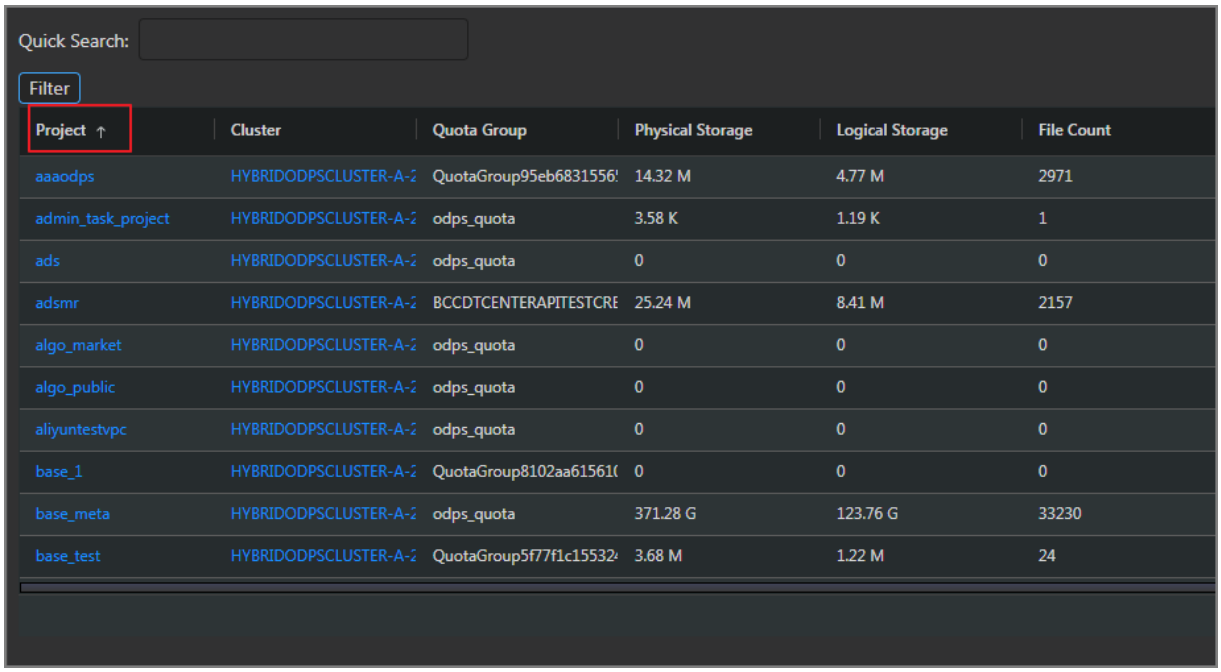

Sort items based on a column

You can sort items based on a column in ascending or descending order. The procedure and display method are different from those described in Sort projects based on a column.

1. On the **MaxCompute** page, click **O&M** in the upper-right corner, and then click the **Clusters** tab.

On the Clusters page, click the **Health Status** tab.

2. In the checker list, click a column heading or the Sort icon in the column heading to sort checkers in ascending order or descending order.

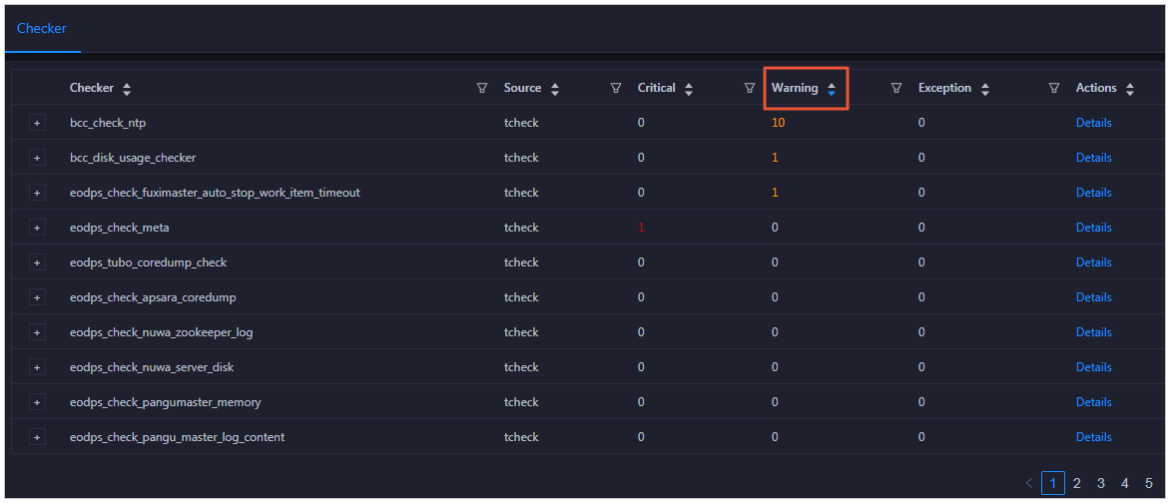

The highlighted up arrow indicates that the checkers are sorted in ascending order. The highlighted down arrow indicates that the checkers are sorted in descending order.

View the trend charts for a MaxCompute cluster

On the **MaxCompute** page, click **O&M** in the upper-right corner, and then click the **Clusters** tab. On the Clusters page, you can view relevant metrics, such as CPU and memory usage, of the selected cluster.

Take CPU usage as an example. The trend chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the specified cluster over time in different colors.

Click in the upper-right corner of the chart to zoom in the chart.

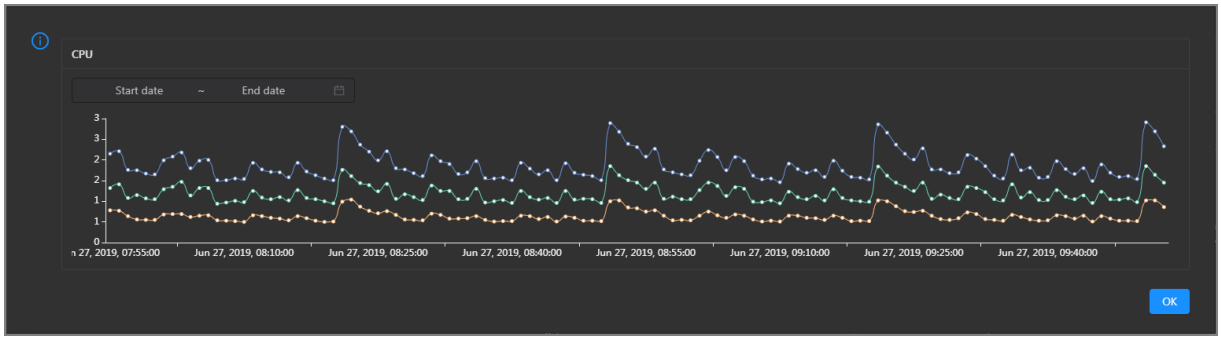

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

View the trend charts for a DataHub cluster

1. On the **DataHub** page, click **O&M** in the upper-right corner, and then click the **Services** tab. In the left-side navigation pane of the Services tab, click **Manage Service**.
2. On the **Overview** page, you can view the trend charts of resource usage for the specified cluster.

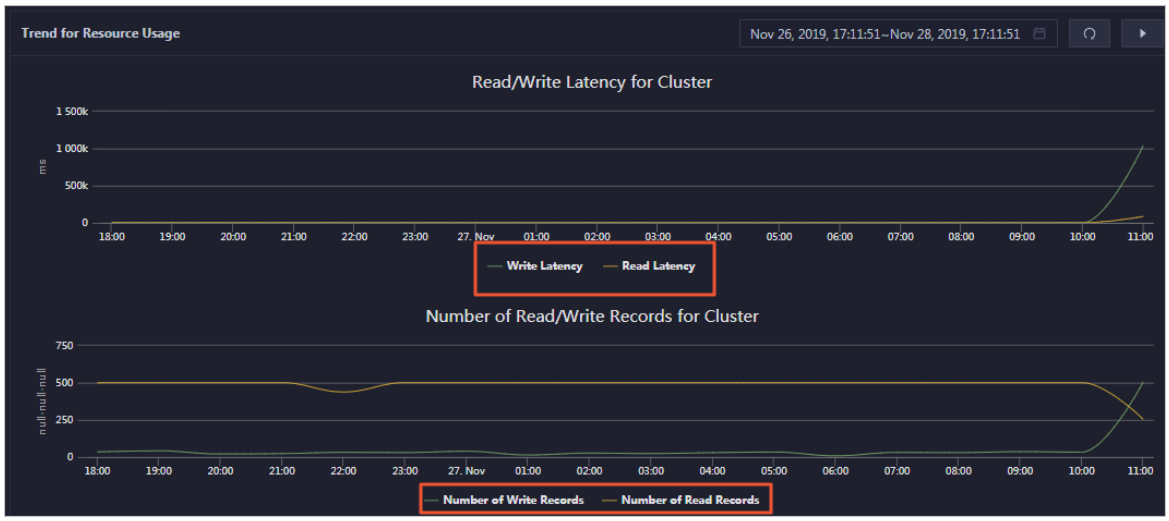

The trend charts, such as the trend charts of the read/write latency and the number of read/write records, appear in the Trend for Resource Usage section. Each chart displays the trend lines of the metrics over time in different colors. You can customize the metrics to display. You can click the name of a metric under the chart to determine whether to display the corresponding trend line in the chart. A highlighted metric name indicates that the corresponding trend line is visible, whereas a dimmed metric name indicates that the corresponding trend line is hidden.

3.Quick start

3.1. Log on to the ABM console

This topic describes how to log on to the Apsara Big Data Manager (ABM) console.

Prerequisites

- The endpoint of the Apsara Uni-manager Operations Console and the username and password used to log on to the console are obtained from the deployment personnel or an administrator.

  The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id*.ops.console.*intranet-domain-id*.

- A browser is available. We recommend that you use Google Chrome.

Procedure

1. Open your Chrome browser.
2. In the address bar, enter the endpoint of the Apsara Uni-manager Operations Console. Press the Enter key.

   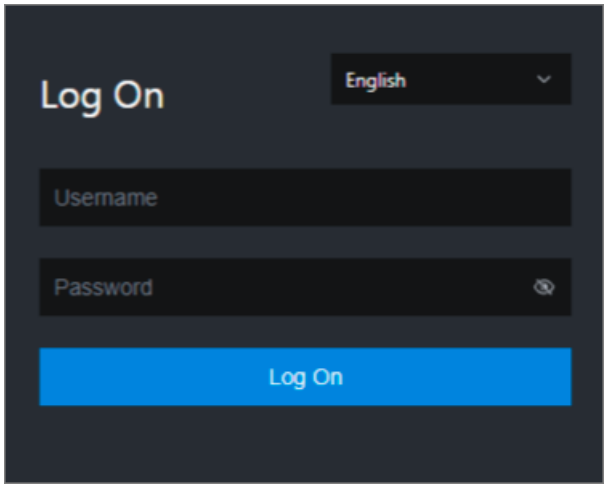

   > **Note** You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

   > **Note** Obtain the username and password used to log on to the Apsara Uni-manager Operations Console from the deployment personnel or an administrator.

   When you log on to the Apsara Uni-manager Operations Console for the first time, you must change the password of your username.

   For security reasons, your password must meet the following requirements:

   - The password contains uppercase and lowercase letters.
   - The password contains digits.
   - The password contains the following special characters: ! @ # $ %

- The password must be 10 to 20 characters in length.

4. Click **Log On**.
5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **Products**.
6. In the **Big Data** section, click **Apsara Bigdata Manager** to go to the homepage of the ABM console.

3.2. Set the theme of the console

You can set the theme of the Apsara Big Data Manager (ABM) console to dark or bright based on your preferences. By default, the dark theme is used.

Prerequisites

An ABM account and the corresponding password are obtained.

Procedure

1. Log on to the ABM console.
2. Set the theme of the ABM console to dark or bright based on your preferences.

Theme	Description
Bright	If the dark theme is used, you can move the pointer over the username in the upper-right corner and turn off the switch to change to the bright theme.
Dark	If the bright theme is used, you can move the pointer over the username in the upper-right corner and turn on the switch to change to the dark theme.

3.3. View the trace dashboards

The dashboard is used to display the key running metrics of MaxCompute, DataWorks, RealtimeCompute, and DataHub products, as well as alarms of all big data products. This allows you to understand the running status of big data products as a whole.

Prerequisites

Your ABM account is granted the required permissions on services on which you want to perform O&M.

Background information

The dashboard is a feature of the ABM console. As the homepage of the ABM console, the dashboard allows you to view the overall running information about all big data services.

Procedure

1. Log on to the ABM console.

   After logging in to the Apsara Big Data Manager, the default display **Dashboard** page. If you are currently on another page, you can click the icon and select **ABM** products to enter **Dashboard** page.

2. View and clear service alerts.

In the alert list, view the number of alerts for all big data products. **Critical** and **Warning** type alarms must be fixed in a timely manner.

i. In the **Dashboard On the** page, click the **Critical** or **Warning** quantity, into the product **Cluster O&M > Health** page.

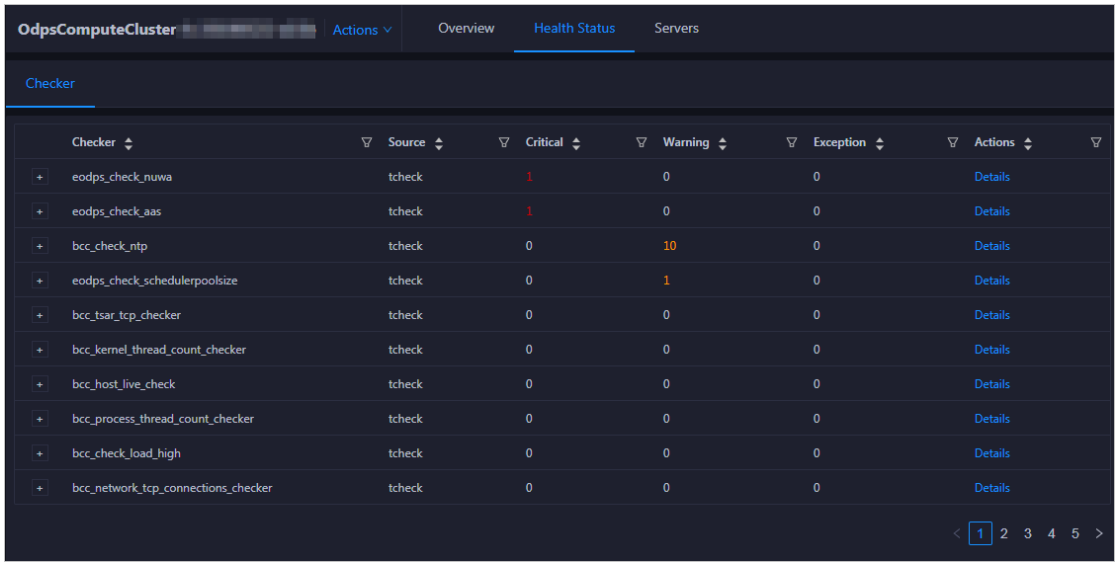

In the **Health** On the page that appears, you can view all check items of the product.

ii. Click the **Details** to view the details of the check item and the alert solution of the check item, and press **Solution The steps in** to handle alerts.

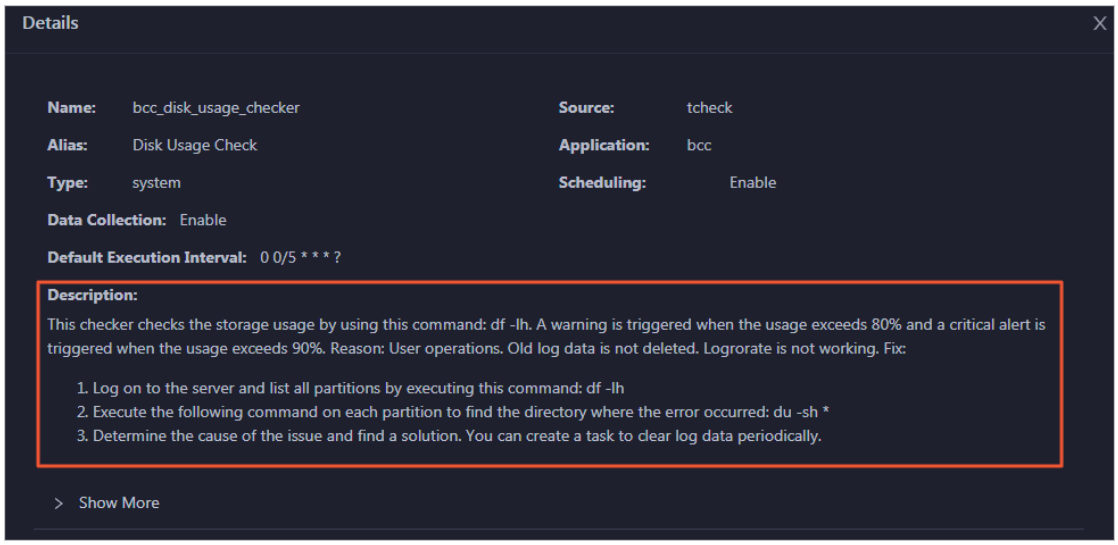

iii. Log on to the hosts on which the alerts are detected to handle the alerts.

Click in front of the check item that has an alarm. **Fold** icon, and then click the **Logon** icon.

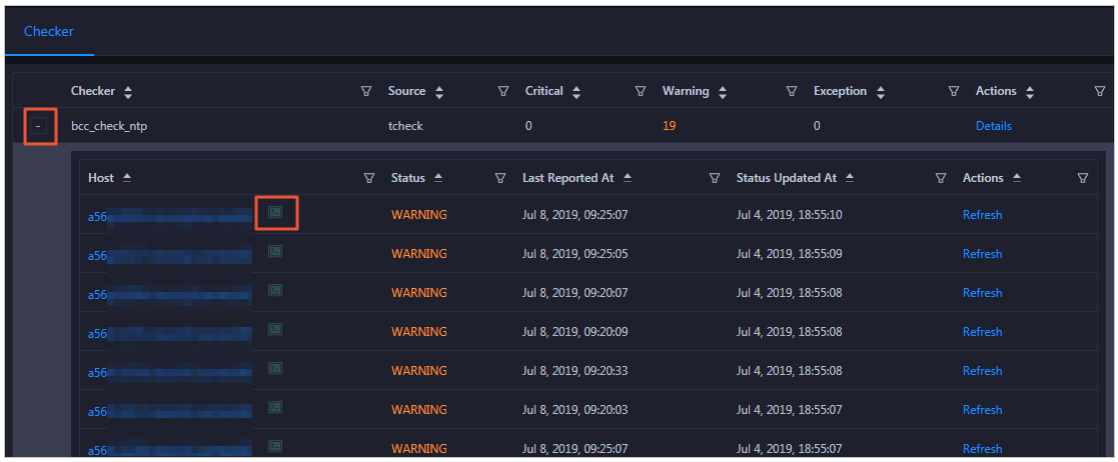

iv. In the newly opened **TerminalService** On the page that appears, select a host on the left to log on.

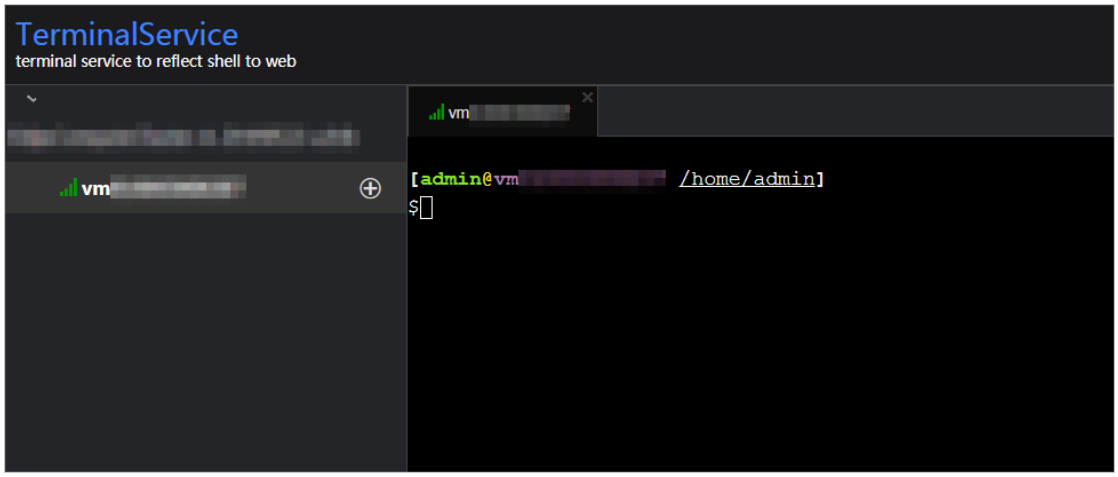

3. In the **Dashboard** On the page that appears, click the **MaxCompute** , view **MaxCompute** .

MaxCompute The section displays the job running overview, control system saturation, data import traffic, computing resource usage, storage resource usage, and the logical and physical CPU usage trend charts of the MaxCompute cluster.

4. In the **Dashboard** On the page that appears, click the **DataWorks** , view **DataWorks** .

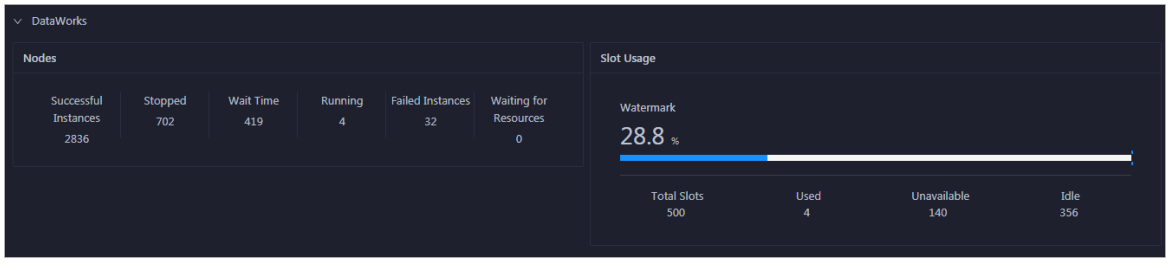

DataWorks The section displays the node scheduling overview, slot resource overview, and the cumulative trend chart of task completion in the DataWorks cluster.

5. In the **Dashboard** On the page that appears, click the **RealtimeCompute** , view **RealtimeCompute** .

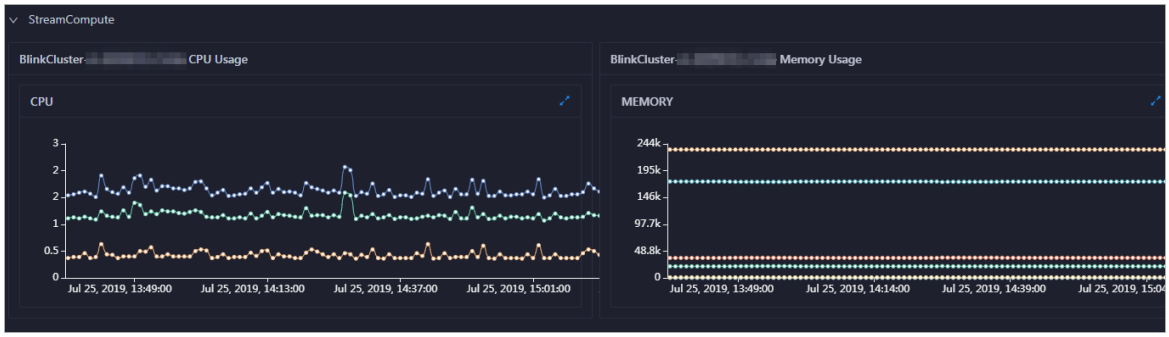

RealtimeCompute The section displays the trend charts of TPS and FAILOVER for RealtimeCompute cluster jobs, and the trend charts of CPU and memory usage.

6. In the **Dashboard** On the page that appears, click the **DataHub** , view **DataHub** .

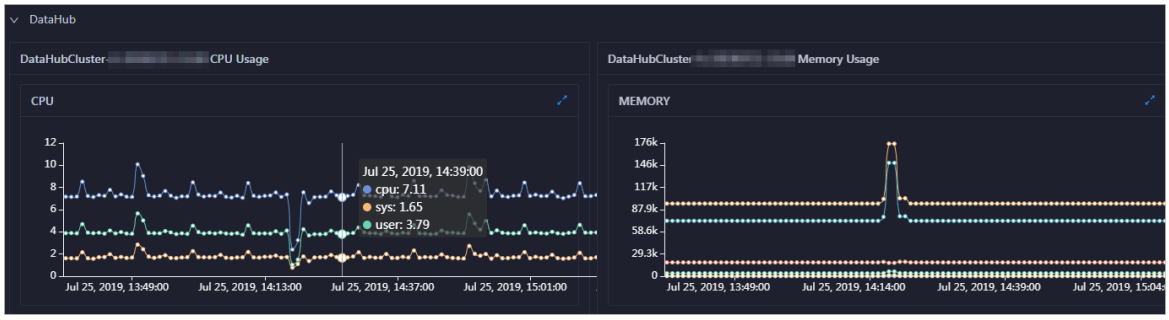

DataHub section displays the trend charts of read /write latency, number of read /write records, read /write QPS, read /write byte traffic, CPU level, and memory level of the DataHub cluster.

3.4. View the cluster running status

Apsara Big Data Manager (ABM) provides you with several operation metrics of clusters, such as CPU usage, memory usage, load, storage, and health check result. This helps you understand the running status of clusters at any time. Based on relevant metrics, you can evaluate whether the selected cluster has operation risks.

Prerequisites

Your ABM account is granted the required permissions to perform O&M operations on the corresponding service.

Context

In the ABM console, the procedures of viewing the cluster running status for different services are the same. This topic uses one of the services as an example.

Procedure

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner and then click a service.
3. On the page that appears, click **O&M** in the upper-right corner, and then click the **Clusters** tab.
4. On the **Clusters** page, select a cluster in the left-side navigation pane. The **Overview** page for the cluster appears.

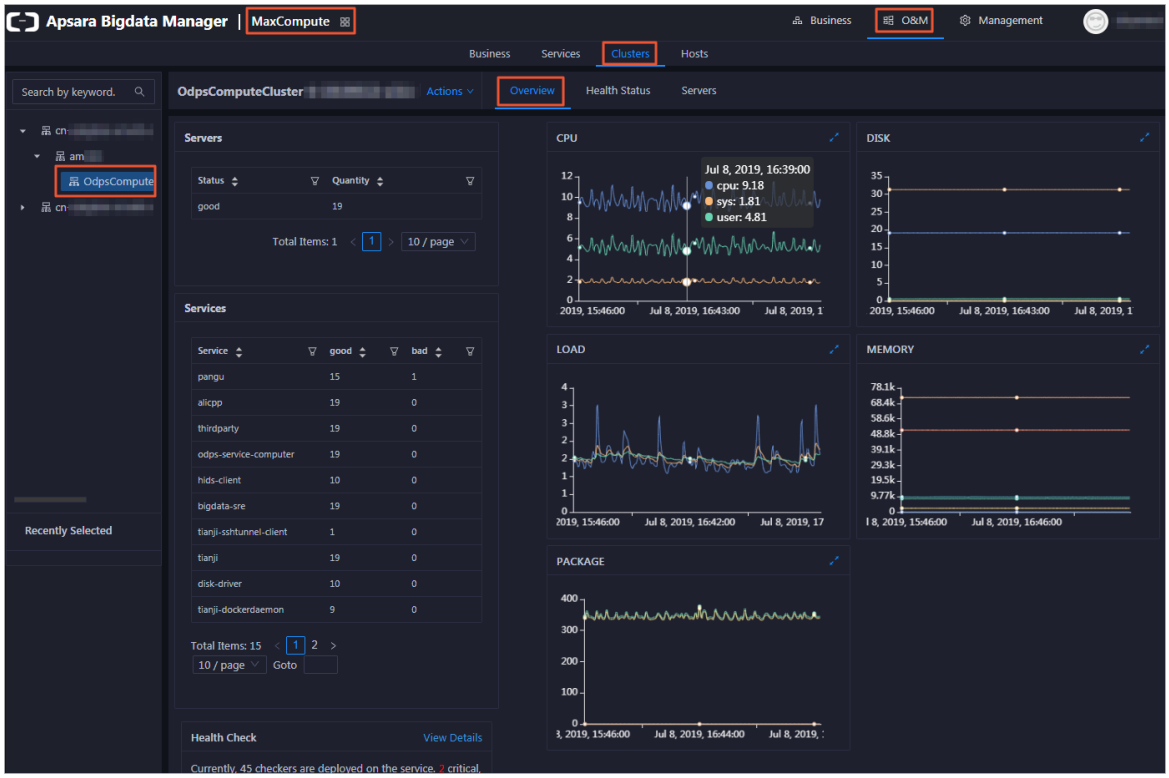

On the **Overview** page, you can view the host status, service status, health check result, and health check history of the selected cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the cluster.

What's next

You can evaluate the operation risks of a cluster based on the metrics such as the service status, CPU usage, disk usage, memory usage, and load.

If the cluster has any Critical, Warning, or Exception alerts, you need to check and clear them in a timely manner. You need to pay special attention to the Critical and Warning alerts. For more information, see View and clear cluster alerts.

3.5. View and clear cluster alerts

If you find alerts on the cluster overview page, go to the cluster health status page to view and clear the alerts. This topic uses one Apsara Big Data Manager (ABM) service as an example to describe how to view and clear alerts.

Prerequisites

Your ABM account is granted the required permissions to perform O&M operations on the corresponding service.

Context

In the ABM console, the procedures of viewing and clearing alerts for different services are the same. If a service has alerts, especially the Critical and Warning alerts, pay attention to them and clear them in a timely manner to make sure that the cluster can run properly.

Procedure

1. Log on to the ABM console.
2. Click in the upper-left corner and then click a service.
3. On the **Clusters** page, select a cluster in the left-side navigation pane, and then click the **Health Status** tab. The Health Status page for the cluster appears.

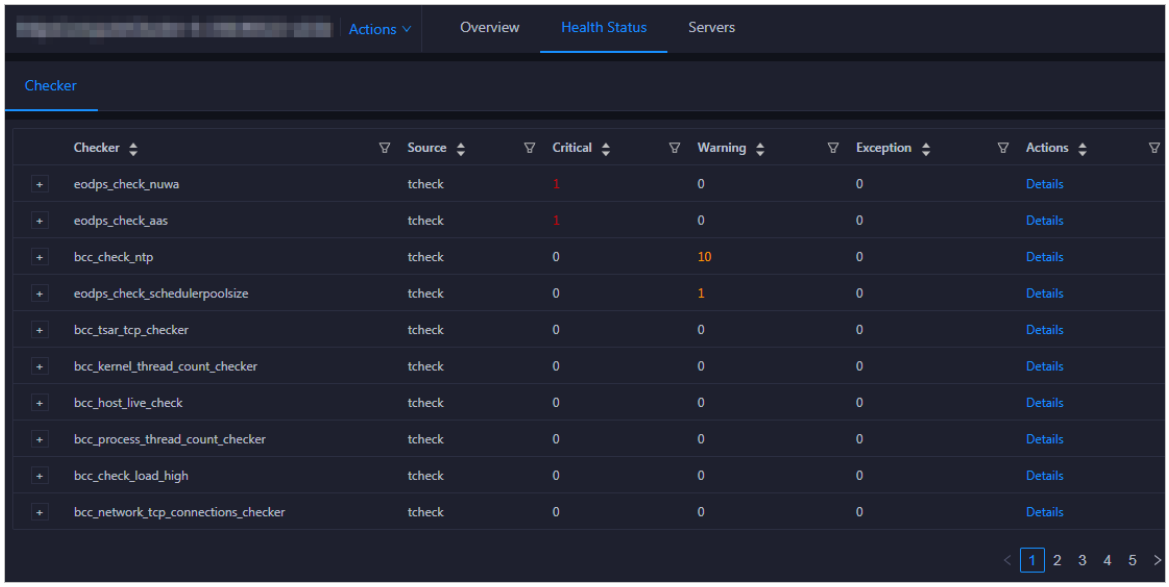

4. On the **Health Status** page, click **+** to expand a checker with alerts. You can view all hosts where the checker is run.

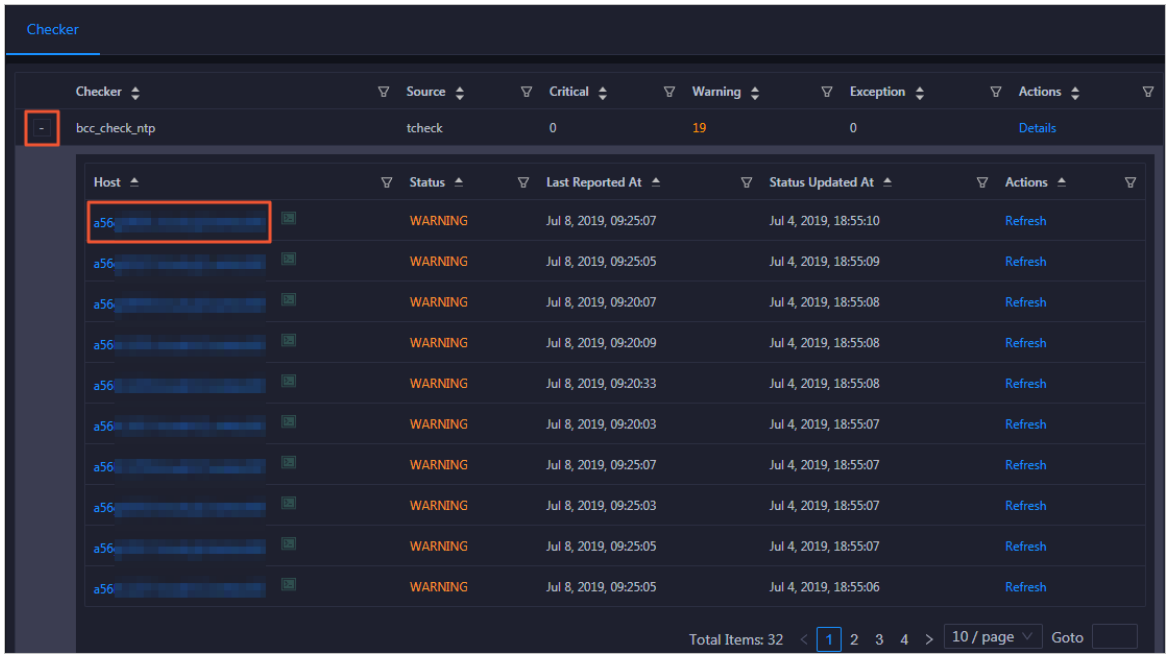

5. Click a hostname. In the dialog box that appears, click **Details** in the Actions column of a check result to view the alert causes.

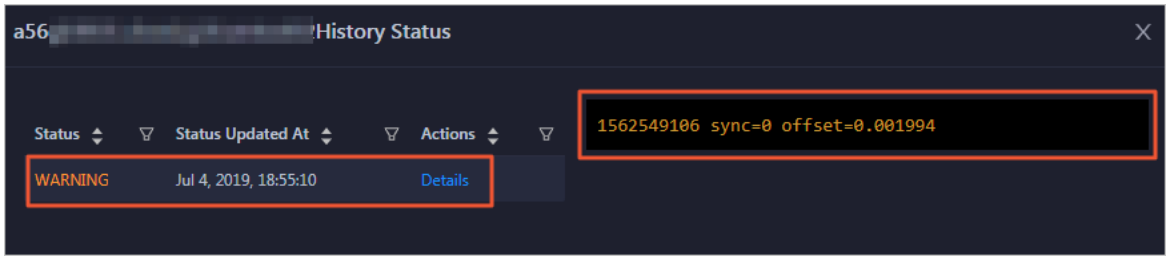

6. On the **Health Status** page, click **Details** in the Actions column of the checker to view the schemes to clear the alerts.

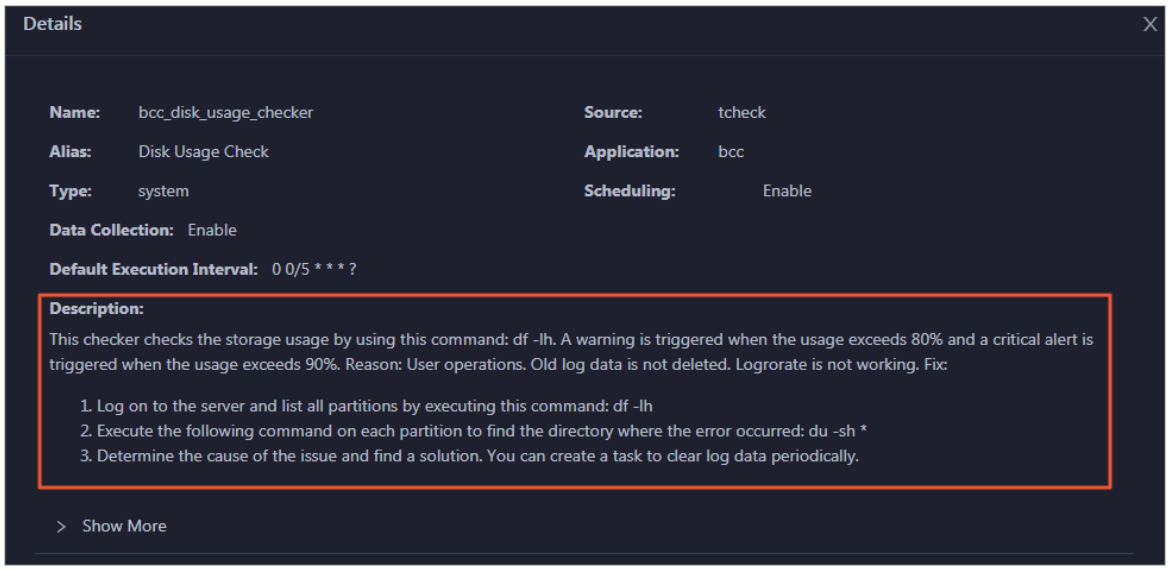

7. Clear the alerts according to the schemes.

To log on to a host with alerts for related operations, click the **Log On** icon next to the name of the host. On the **TerminalService** page that appears, click the hostname on the left to log on to the host.

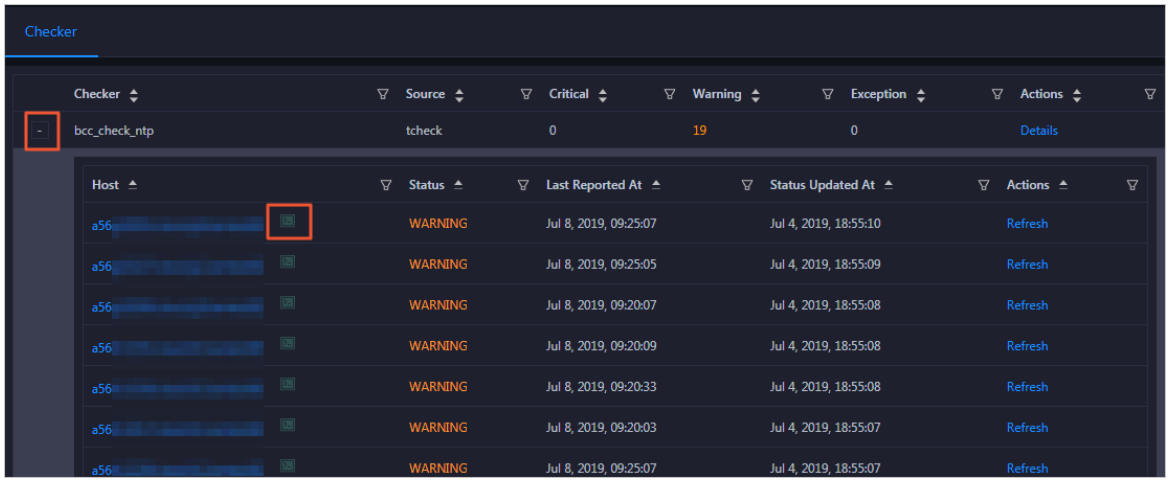

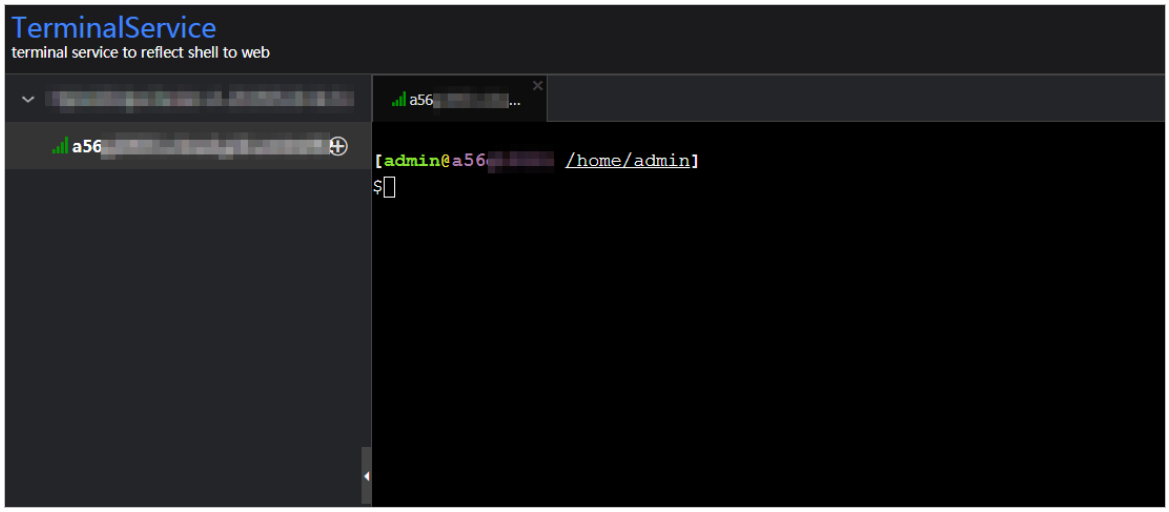

8. After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

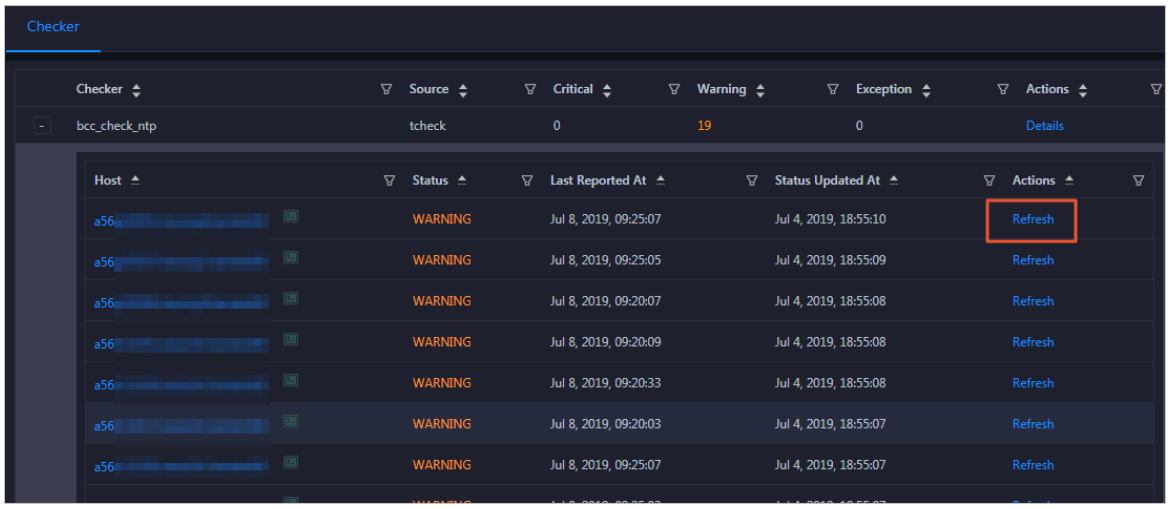

4.ABM

4.1. ABM dashboard

The Apsara Bigdata Manager (ABM) dashboard shows the key indicators of MaxCompute, DataWorks, Realtime Compute for Apache Flink, and DataHub. The dashboard also provides information about the alerts for all big data services and helps you understand the overall status of these services. The dashboard supports auto-refresh and full-screen display.

Go to the Dashboard tab

After you log on to the ABM console, the **Dashboard** tab appears by default. To return to the **Dashboard** tab, click the icon in the upper-left corner and click **ABM**.

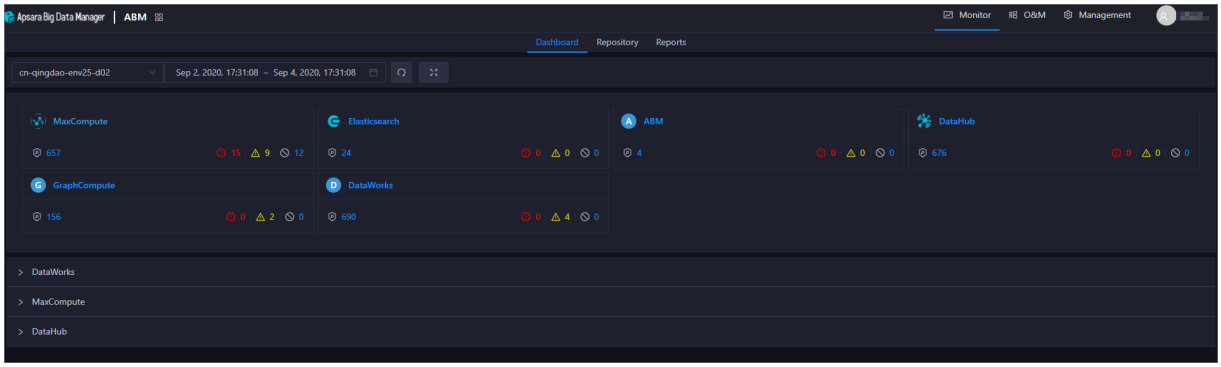

In the upper-left corner of the **Dashboard** tab, you can select a region from the drop-down list to view the cluster status of each big data service in the region.

View and handle the alerts of various services

In the Overview section, you can view the numbers of **Critical**, **Warning**, and **Exception** alerts that are reported for each big data service. If a service has alerts, especially **Critical** or **Warning** alerts, handle these alerts on time.

1. On the **Dashboard** tab, find the check item of a service that you want to query, and click the number in the **Critical** or **Warning** column of the service. The **Health Status** page for the service appears on the **Clusters** tab.

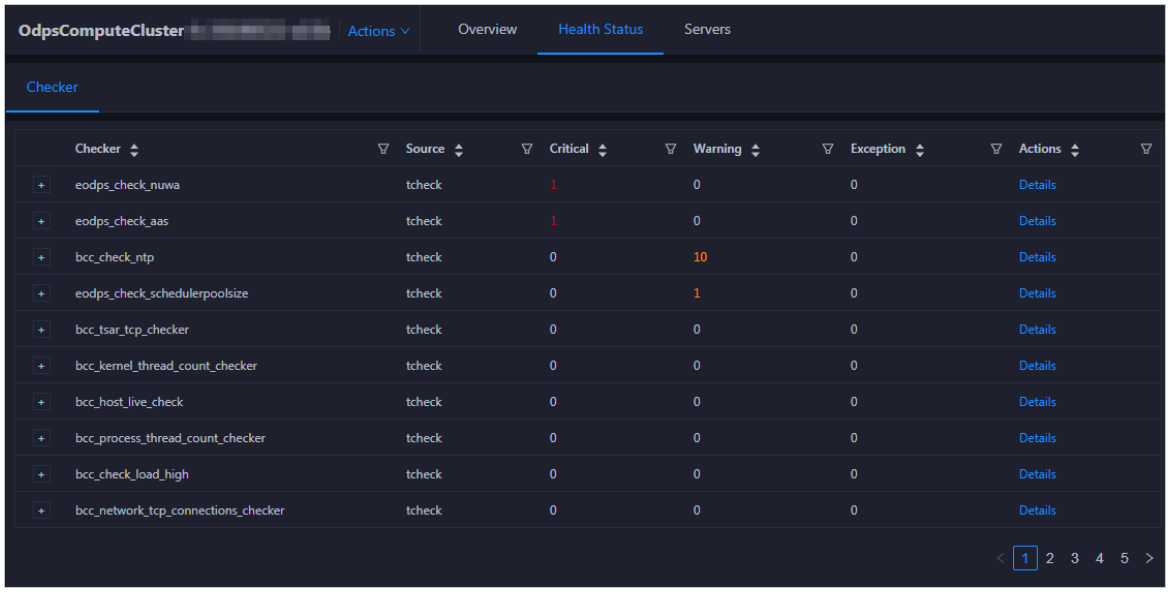

On the **Health Status** page, you can view all the check items of the service.

2. Click **Details** in the Actions column of a check item for which alerts are reported. In the Details dialog box, view the details of the check item and the descriptions to handle the alerts. Perform the steps provided in the **Description** section to handle the alerts.

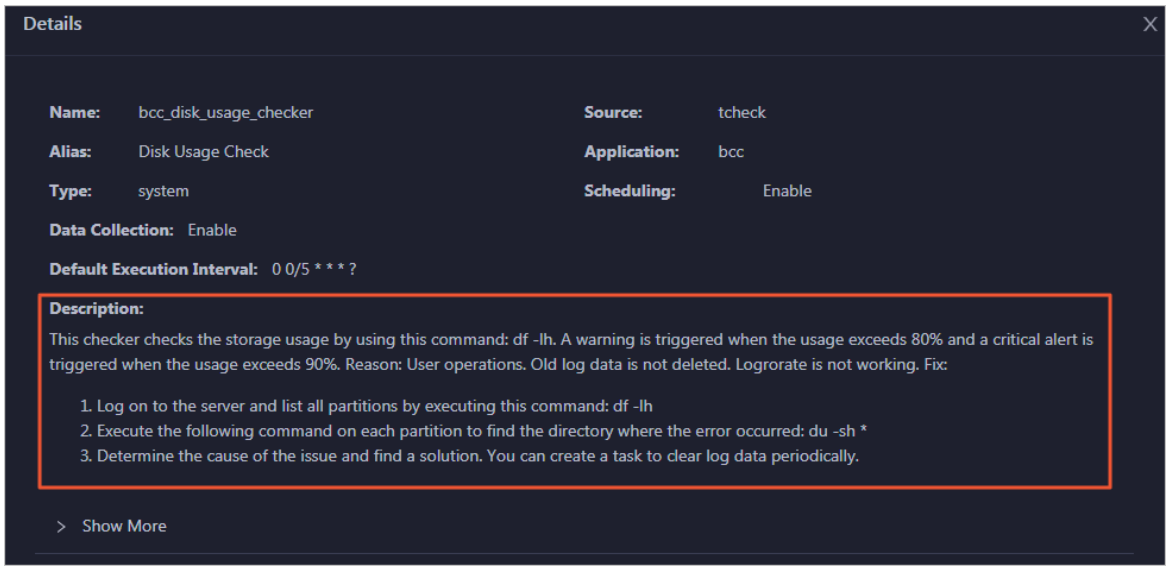

3. Log on to the hosts on which the alerts are detected to handle the alerts.

   Click the plus sign (+) to expand a check item with alerts, and click the **Log On** icon next to the name of a host with alerts. On the TerminalService page that appears, click the hostname on the left to log on to the host.

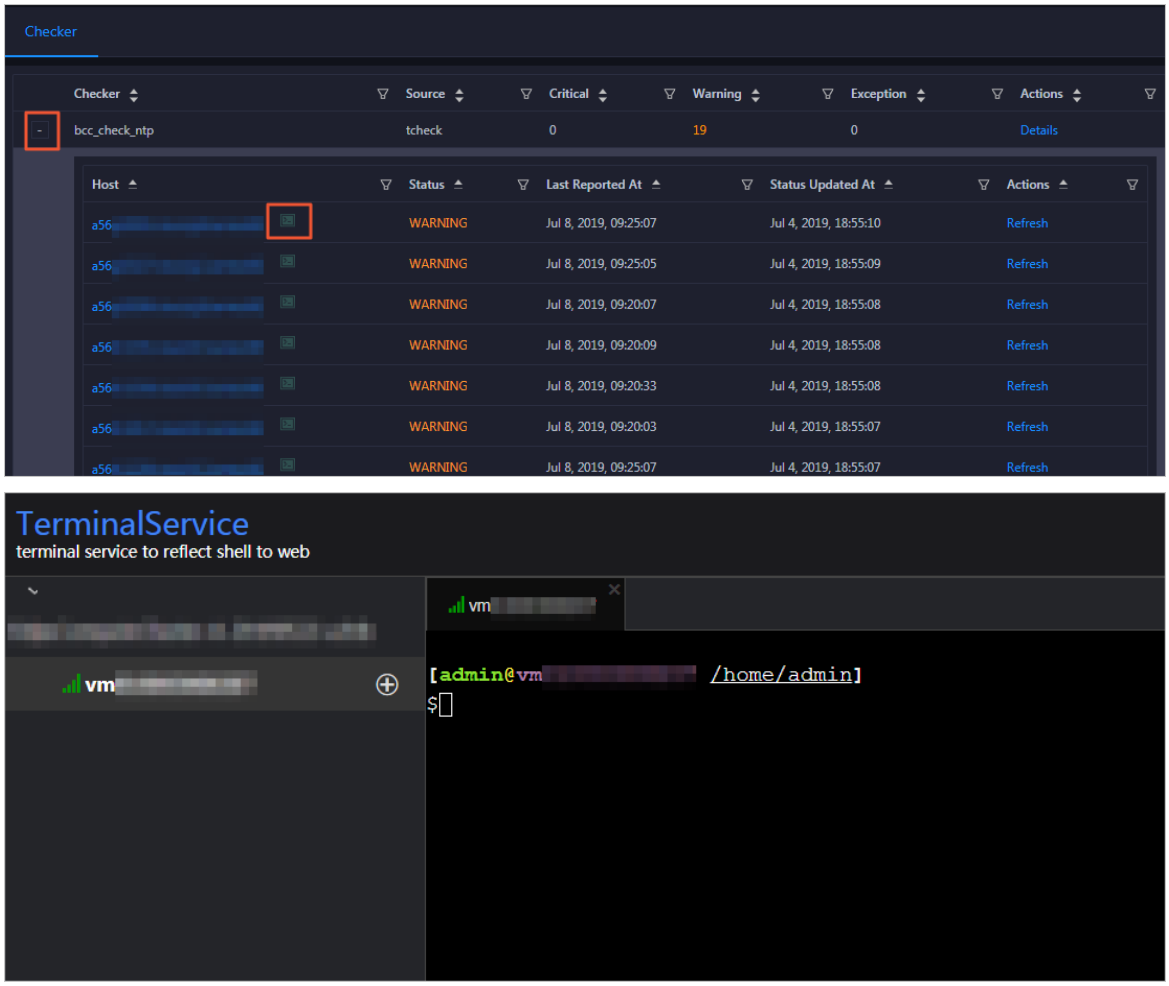

View key indicators of MaxCompute

The ABM dashboard shows the key indicators of MaxCompute. On the **Dashboard** tab, click **MaxCompute** to view the information.

In the **MaxCompute** section, you can view the job status, the real-time capacity for the control system, computing resource usage, and storage resource usage. You can also view the trend charts of imported data traffic, logical CPU utilization, and physical CPU utilization.

View key indicators of DataWorks

The ABM dashboard shows key indicators of DataWorks. On the **Dashboard** tab, click **DataWorks** in the **Monitoring** column to view the information.

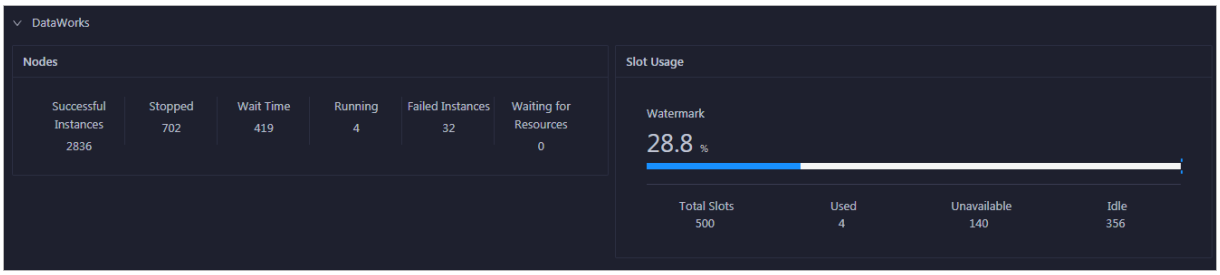

In the **DataWorks** section, you can view the node scheduling and slot usage of a DataWorks cluster. You can also view the trend chart of the total number of daily finished tasks.

View key indicators of Realtime Compute for Apache Flink

The ABM dashboard shows key indicators of Realtime Compute for Apache Flink. On the **Dashboard** tab, click **Realtime Compute** in the **Monitoring** column to view the information.

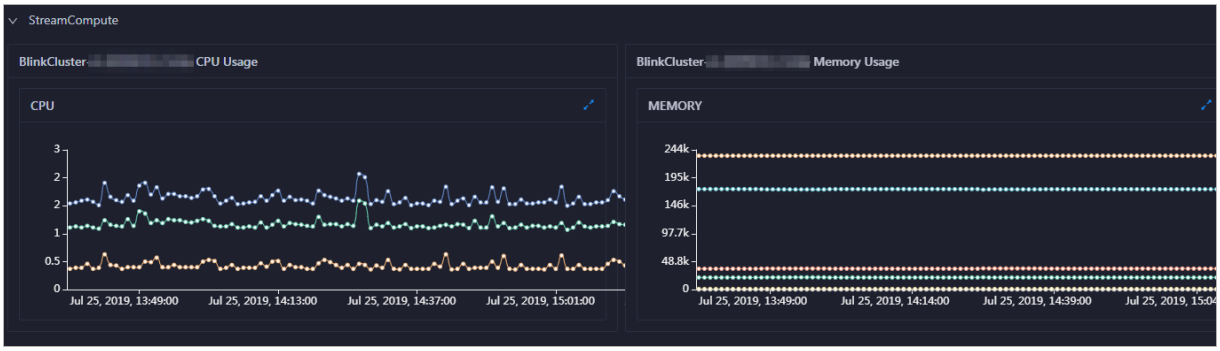

In the **Realtime Compute** section, you can view the trend charts of the transactions per second (TPS), failover rate, CPU utilization, and memory usage for a Realtime Compute for Apache Flink cluster.

View key indicators of DataHub

The ABM dashboard shows key indicators of DataHub. On the **Dashboard** tab, click **DataHub** in the **Monitoring** column to view the information.

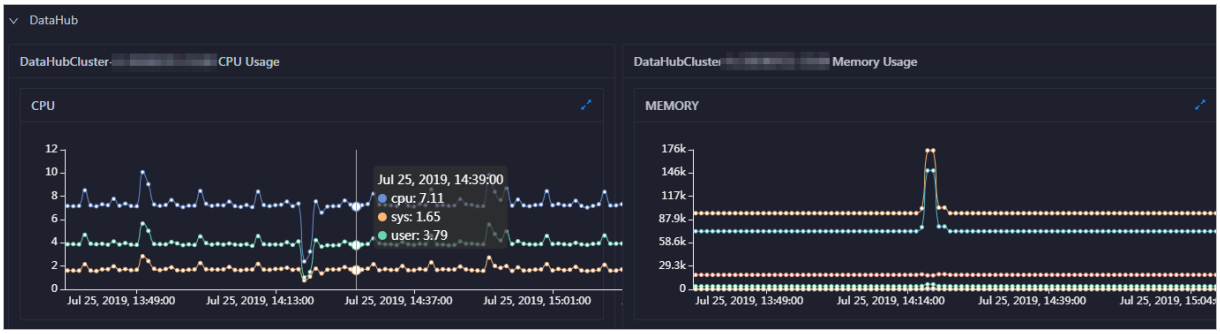

In the **DataHub** section, you can view the trend charts of the read/write latency, read/write records, read/write queries per second (QPS), and read/write throughput. You can also view the trend charts of CPU utilization and memory usage of a DataHub cluster.

Enable and disable auto-refresh

By default, auto-refresh is disabled on the Dashboard tab, and the statistics of cluster metrics from the last two days are displayed on this page. You can specify a time range to view the metric statistics. If you enable auto-refresh, the system automatically updates the metric data of clusters based on the specified interval.

1. At the top of the **Dashboard** tab, click the icon.
2. In the dialog box that appears, configure the **Refreshing every** and **Refreshing range** parameters.

   The **Refreshing range** parameter specifies a time period for the trend charts, such as those for the CPU utilization and memory usage of each cluster.
3. After you configure these parameters, click **OK** to enable auto-refresh.

   If auto-refresh is enabled, the icon is replaced with the icon. The system automatically updates all data on the dashboard based on the specified time interval.

   If you want to disable auto-refresh, click the icon.

Display the dashboard in full-screen

The dashboard supports full-screen display. This feature allows you to view the status of big data services.

At the top of the **Dashboard** tab, click the icon to display the **Dashboard** tab in full-screen mode.

4.2. ABM repository

The Repository page in the Apsara Big Data Manager (ABM) console displays the resource usage in MaxCompute, DataWorks, and DataHub. This topic describes the features of the ABM repository and how to access the Repository page.

Entry

1. Log on to the ABM console.

   > **Note**
   >
   > By default, the **Dashboard** page appears. To return to the **Dashboard** page from any other page, click in the upper-left corner and then click **ABM**.

2. On the **Dashboard** page, click the **Repository** tab. The **Repository** page appears.

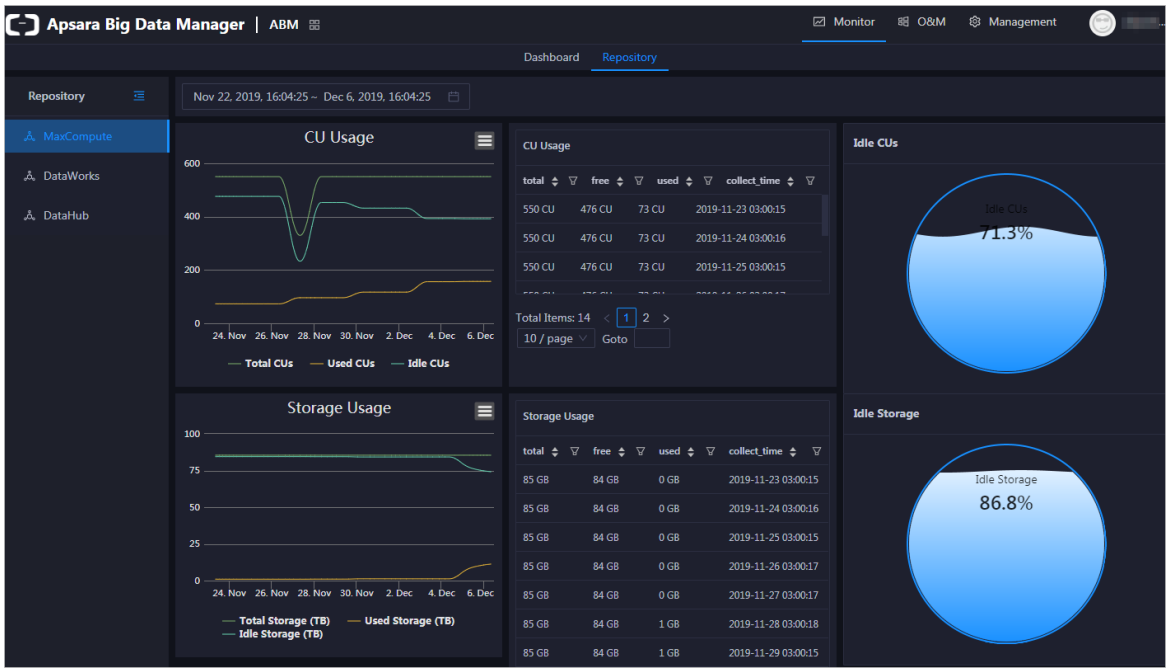

View the resource usage in MaxCompute

In the left-side navigation pane of the **Repository** page, click **MaxCompute**. On the page that appears, you can view the resource usage in MaxCompute.

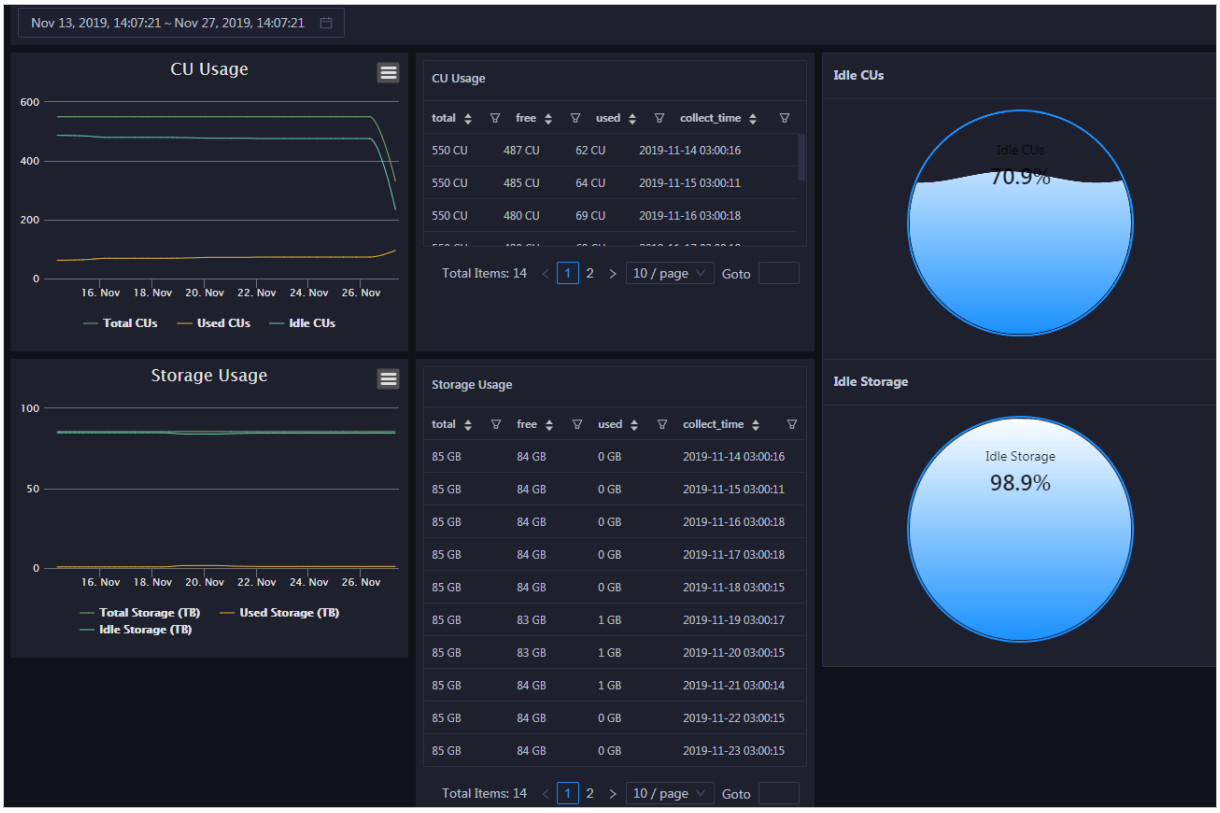

For MaxCompute, the Repository page displays the trend charts of CU and storage usage, records of CU and storage usage, and proportions of idle CUs and storage.

View the resource usage in DataWorks

In the left-side navigation pane of the **Repository** page, click **DataWorks**. On the page that appears, you can view the resource usage in DataWorks.

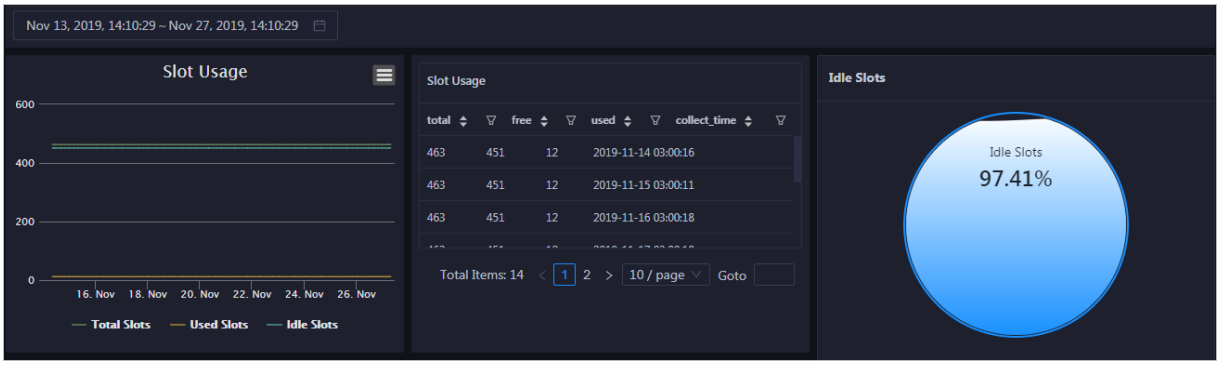

For DataWorks, the Repository page displays the trend chart of slot usage, records of slot usage, and proportion of idle slots.

View the resource usage in DataHub

In the left-side navigation pane of the **Repository** page, click **DataHub**. On the page that appears, you can view the resource usage in DataHub.

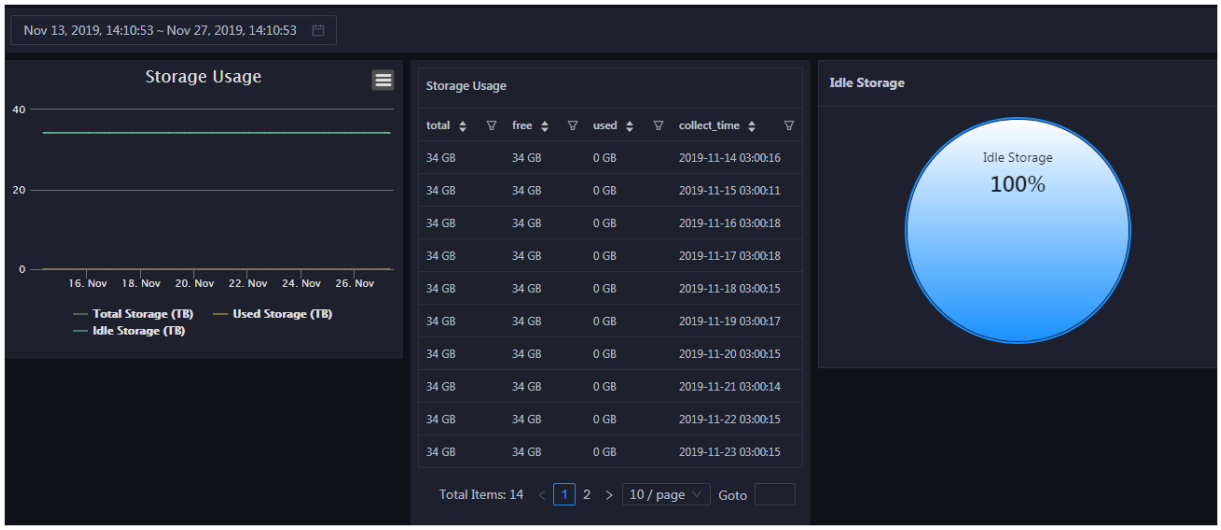

For DataHub, the Repository page displays the trend chart of storage usage, records of storage usage, and proportion of idle storage.

Other operations

You can filter or sort records of CU, storage, and slot usage based on a column to facilitate information retrieval. For more information, see Common operations.

4.3. O&M overview

This topic describes the O&M modules of Apsara Big Data Manager (ABM) and how to go to the O&M page of ABM.

O&M modules

The O&M page of ABM provides the following modules: Business, Clusters, and Hosts. The following table describes these modules.

Module	Submodule or feature	Description
Environment O&M	Single-Host Basic Components	Shows the status of single-host basic components within a cluster.
Clusters	Overview	Shows the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Health Status	Shows all check items of a cluster, including the check item details, check results for the hosts in the cluster, and schemes to clear alerts. You can also log on to a host and manually check the host.

Module	Submodule or feature	Description
Hosts	Overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.

Go to the O&M page

1. Log on to the ABM console
2. In the upper-left corner of the ABM console, click the icon and click **ABM**.
3. In the upper-right corner of the ABM page, click **O&M**. On the page that appears, click the **Business** tab to go to the Business page.

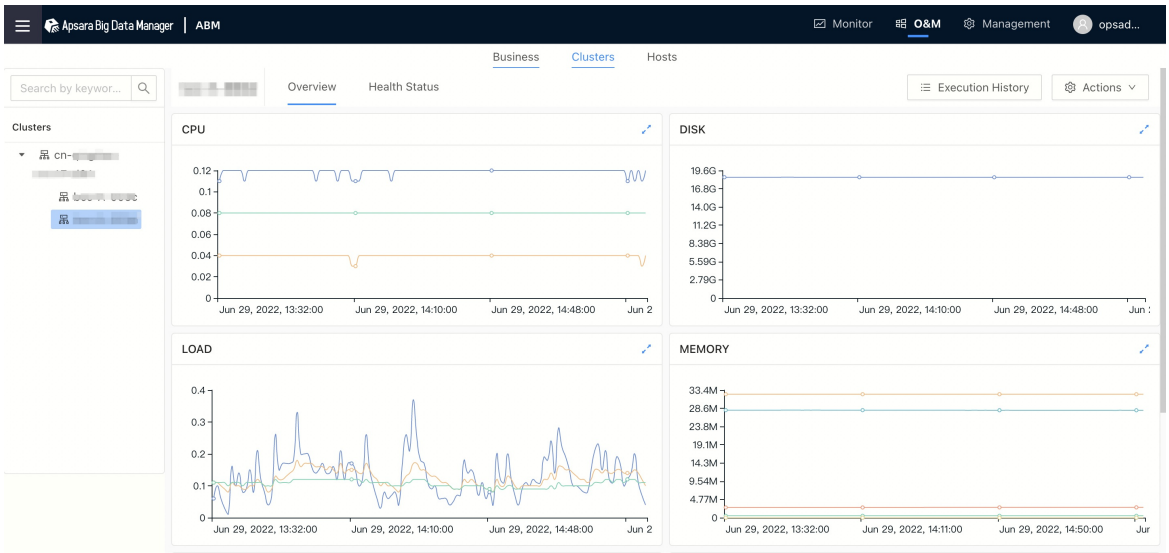

The **O&M** page includes the following modules: **Business**, **Clusters**, and **Hosts**.

4.4. Service O&M

4.4.1. Service overview

The Overview page lists all Apsara Big Data Manager (ABM) services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

On the **Services** page, select a cluster above the left-side service list, select a service in the service list, and then click the **Overview** tab. The **Overview** page for the service appears.

On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click in the upper-right corner of the chart to zoom in the chart.

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.

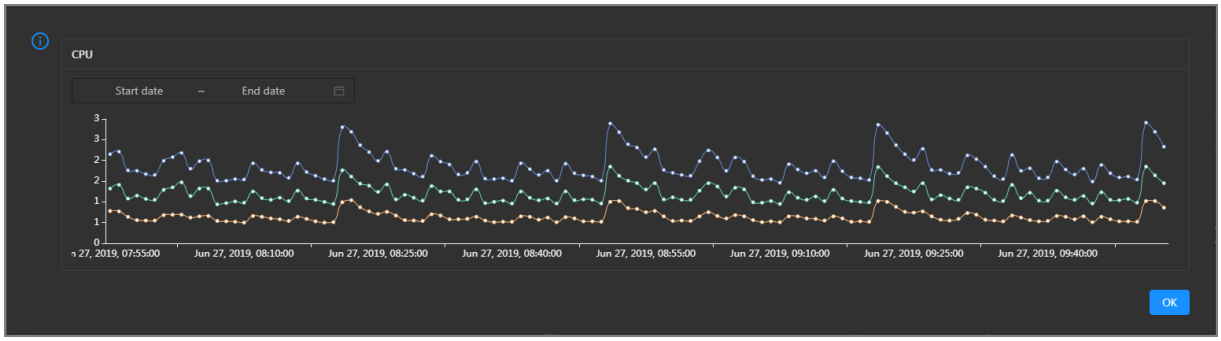

DISK

This chart displays the trend lines of the storage space usage on the */*, */boot*, */home/admin*, and */home* directories for the selected service over time in different colors.

Click in the upper-right corner of the chart to zoom in the chart.

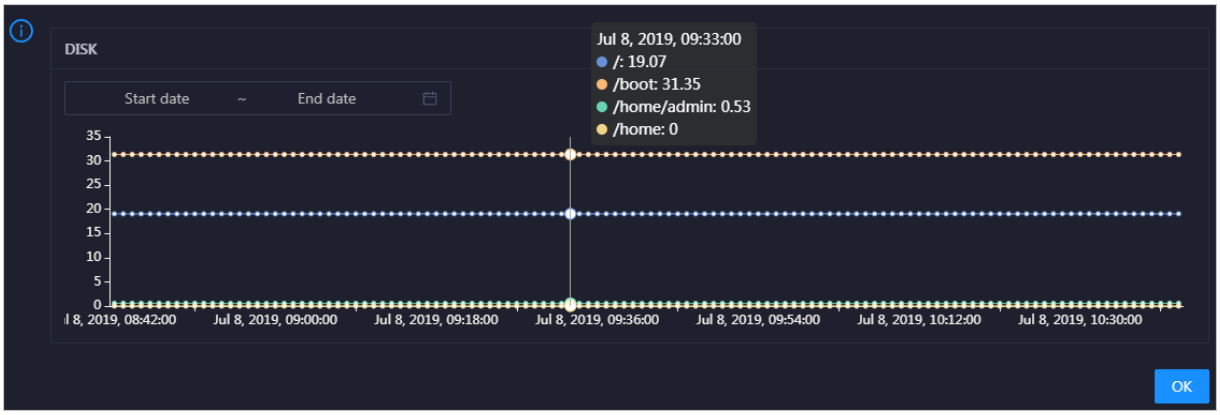

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click in the upper-right corner of the chart to zoom in the chart.

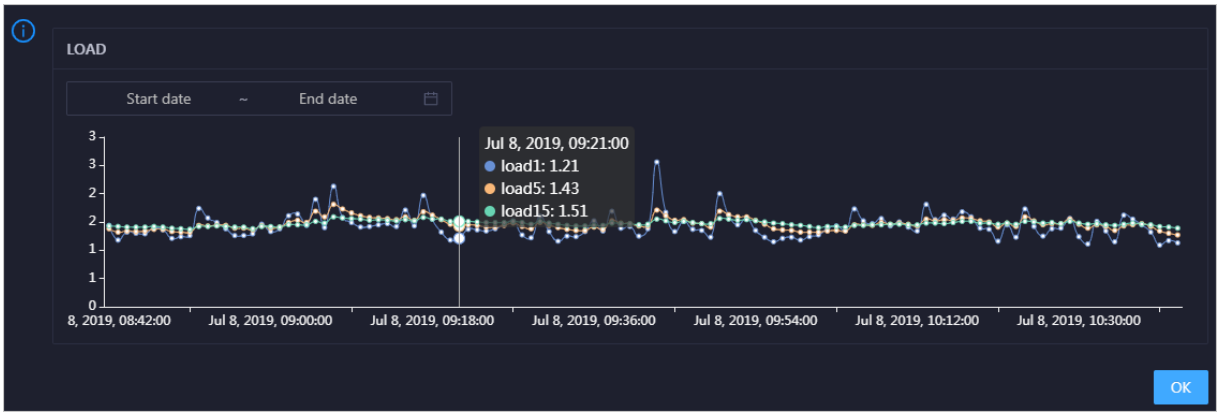

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click in the upper-right corner of the chart to zoom in the chart.

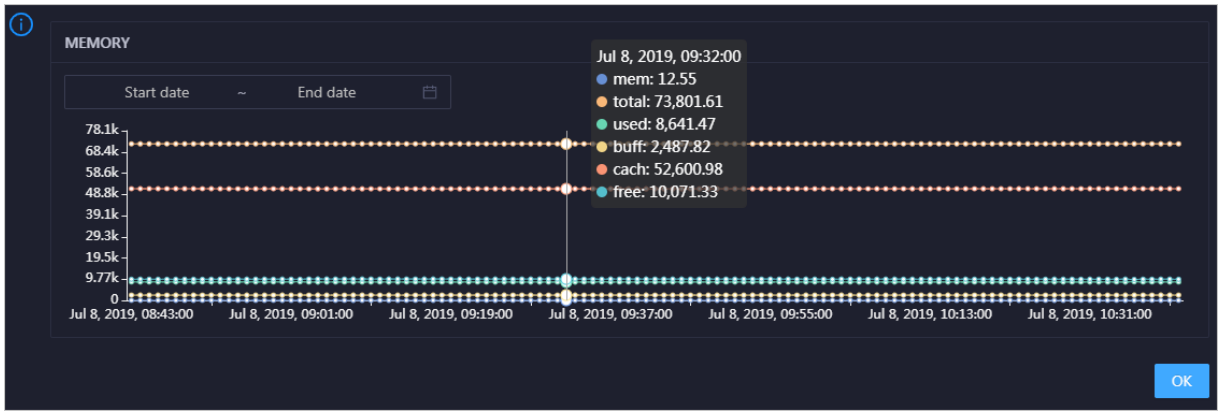

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click in the upper-right corner of the chart to zoom in the chart.

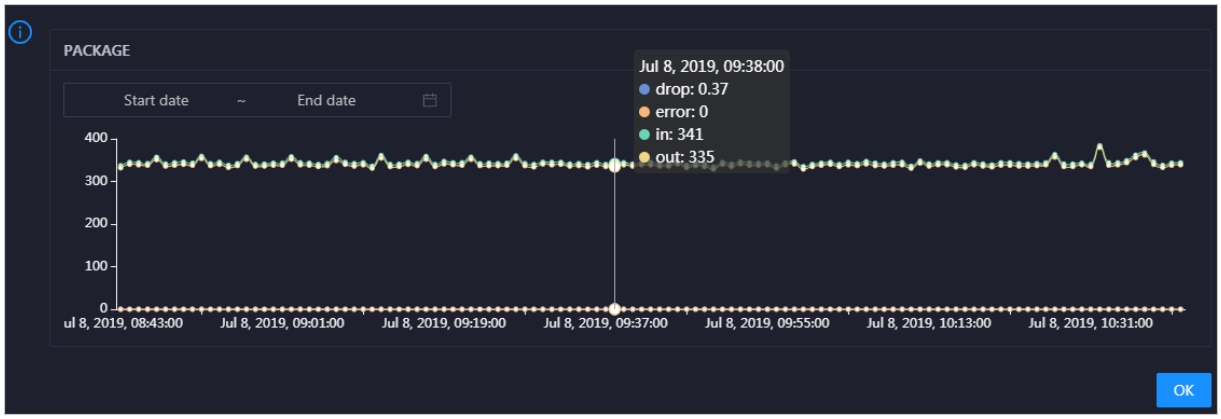

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click in the upper-right corner of the chart to zoom in the chart.

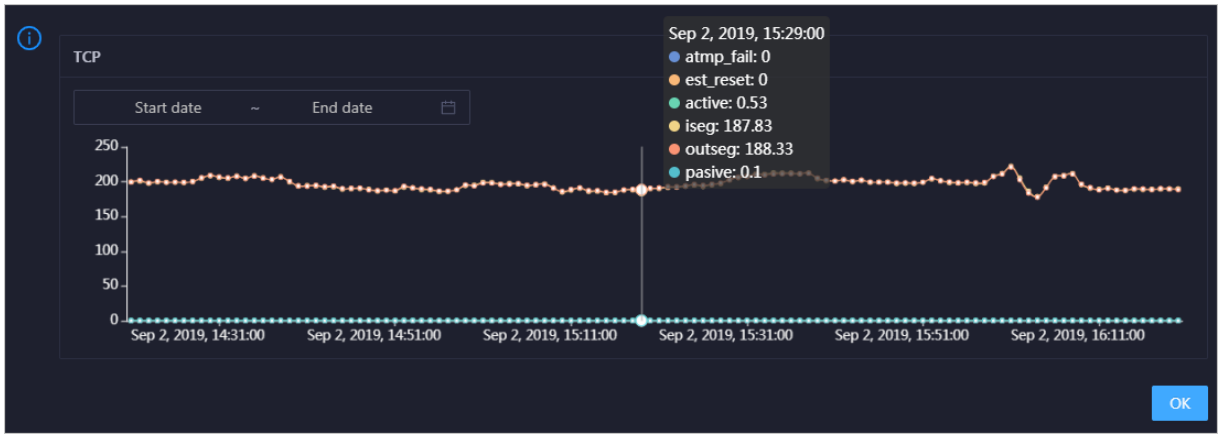

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click in the upper-right corner of the chart to zoom in the chart.

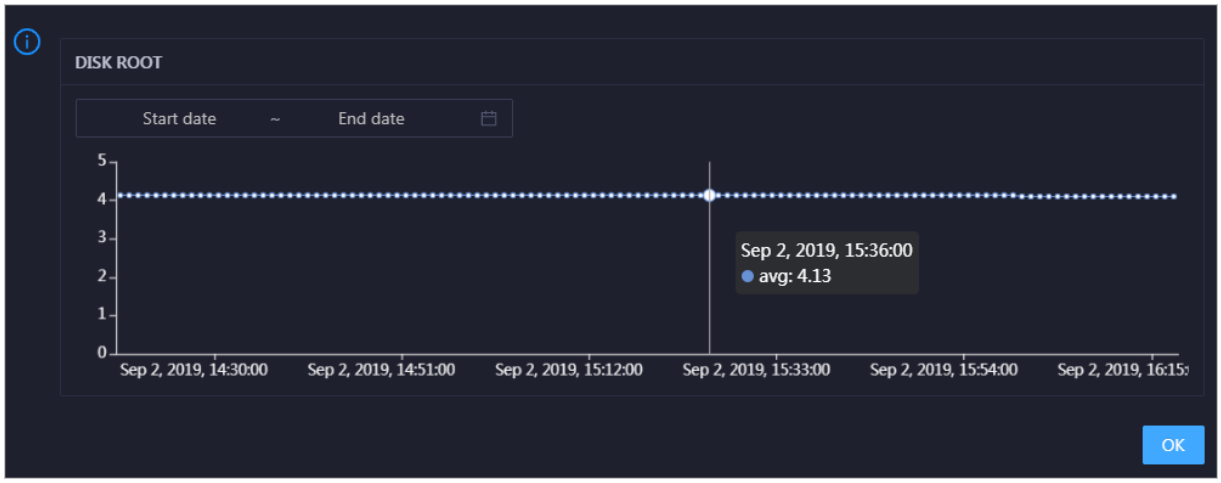

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

4.4.2. Service hosts

Apsara Big Data Manager (ABM) allows you to view the host list of each ABM service so that you can understand the service deployment on hosts.

On the **Services** page, select a cluster above the left-side service list, select a service in the service list, and then click the **Server** tab. The **Server** page for the service appears.

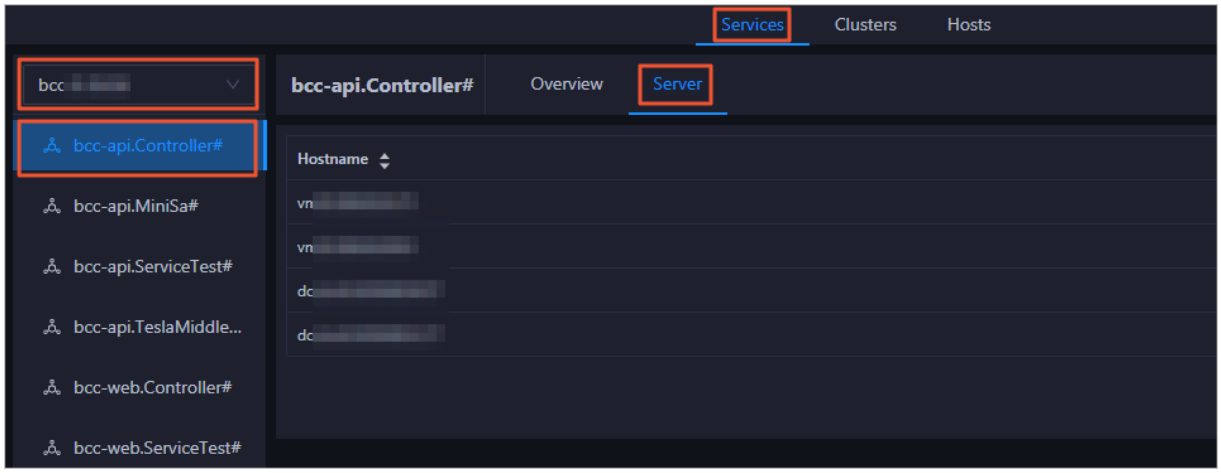

On the **Server** page, you can view the hosts where the selected service is run.

4.5. Cluster O&M

4.5.1. Cluster overview

The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.

Entry

On the **Clusters** page, select a cluster in the left-side navigation pane, and then click the **Overview** tab. The Overview page for the cluster appears.

The cluster overview page displays the trend charts of CPU usage, memory usage, and load for a cluster. The trend charts are described as follows:

CPU

This chart shows the trend lines of the total CPU utilization (cpu), CPU utilization for executing code in kernel space (sys), and CPU utilization for executing code in user space (user) for the cluster in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the cluster in the specified period.

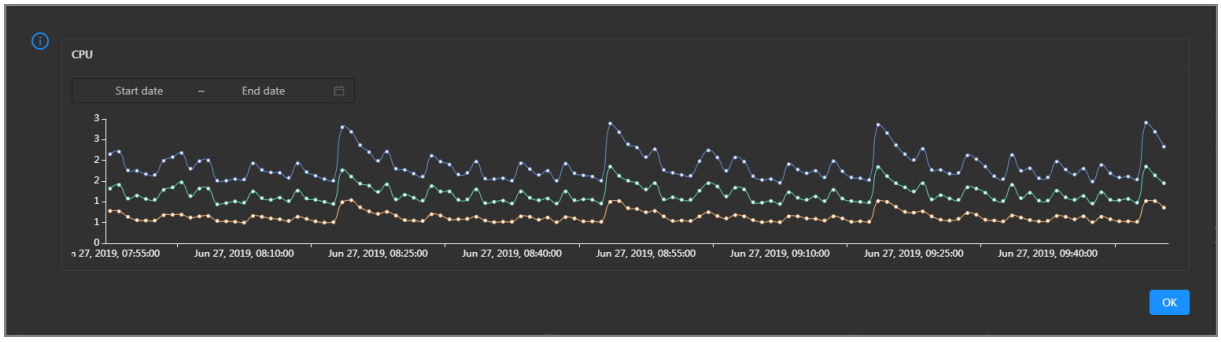

MEMORY

This chart shows the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

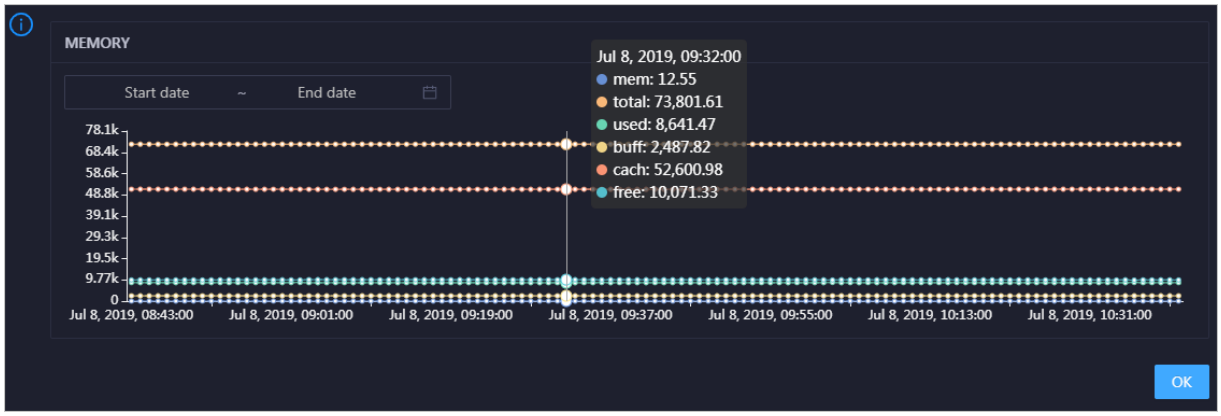

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

LOAD

This chart shows the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

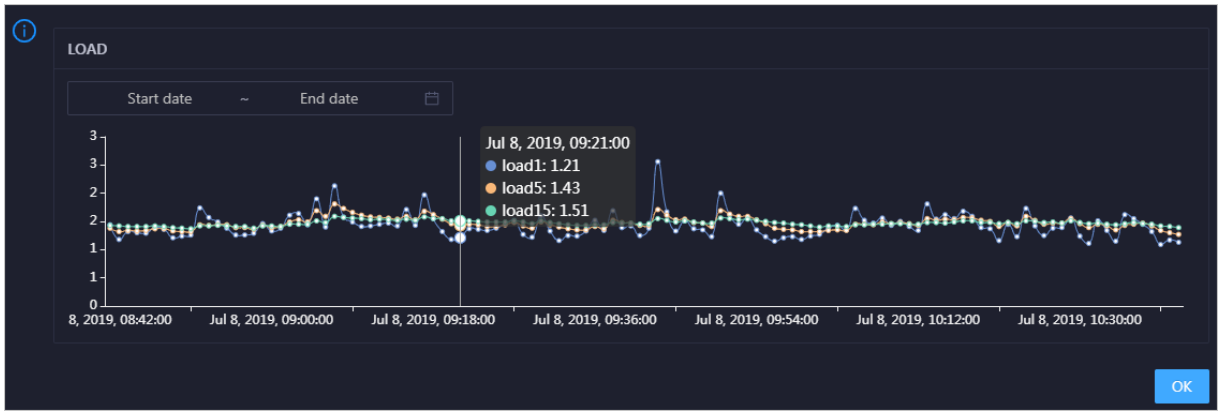

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

4.5.2. Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

On the **Clusters** page, select a cluster in the left-side navigation pane, and then click the **Health Status** tab. The Health Status page for the cluster appears.

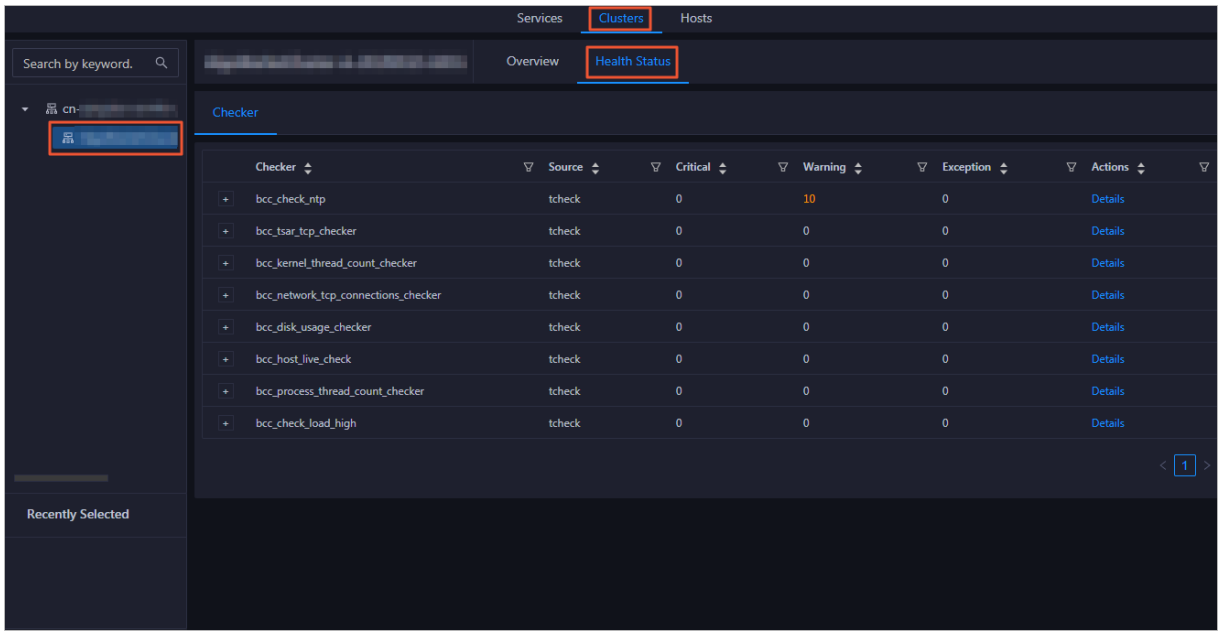

On the **Health Status** tab, you can view all checkers for the cluster and the check results for the hosts in the cluster. The following alerts may be reported on a host: **CRITICAL**, **WARNING**, and **EXCEPTION**. The alerts are represented in different colors. You must handle the alerts in a timely manner, especially the **CRITICAL** and **WARNING** alerts.

View checker details

1. On the Health Status tab, click **Details** in the Actions column of a checker. On the Details page, view checker details.

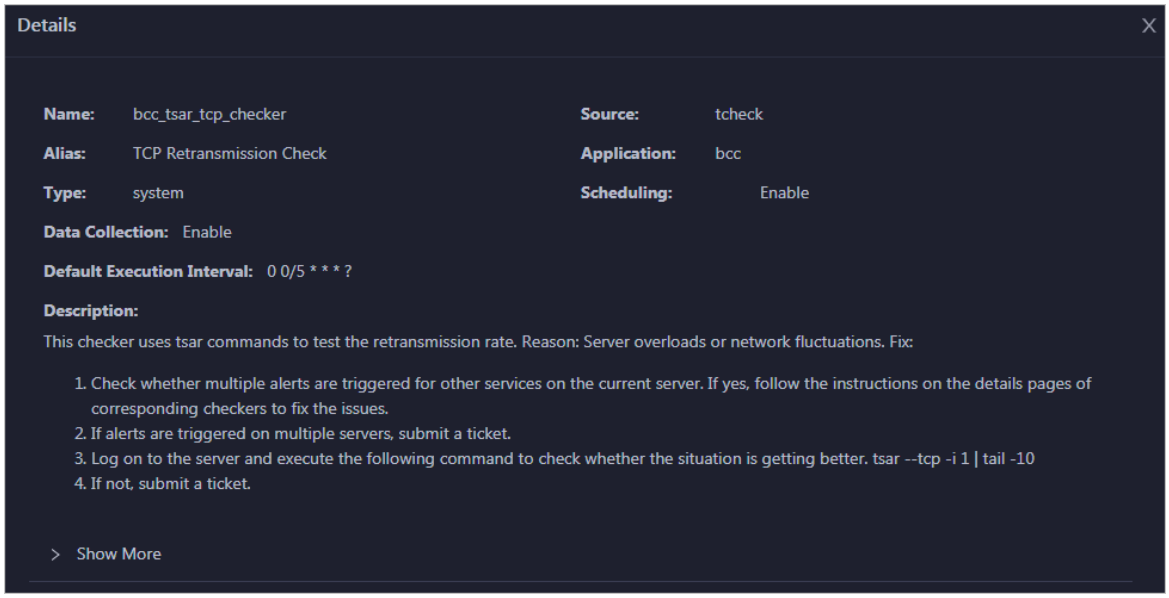

The checker details include **Name**, **Source**, **Alias**, **Application**, **Type**, **Scheduling**, **Data Collection**, **Default Execution Interval**, and **Description**. The schemes to clear alerts are provided in the description.

2. Click **Show More** to view more information about the checker.

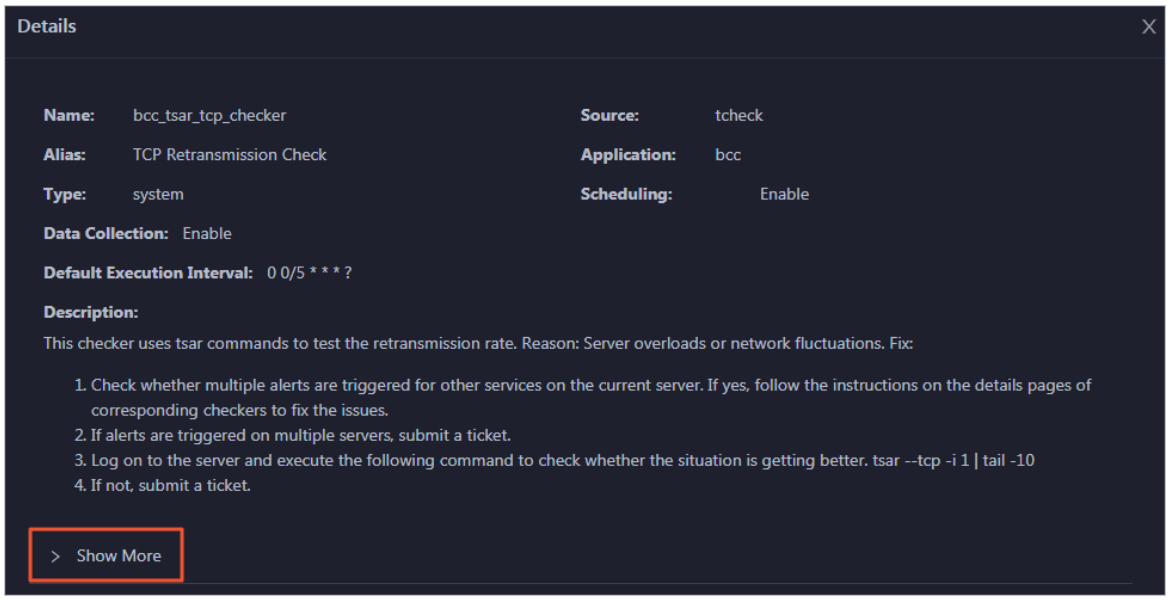

You can view information about **Script**, **Target (TianJi)**, **Default Threshold**, and **Mount Point**.

View the hosts for which alerts are reported and causes for the alerts

You can view the check history and check results of a checker on a host.

1. On the Health Status tab, click **+** to expand a checker for which alerts are reported. You can view all hosts where the checker is run.

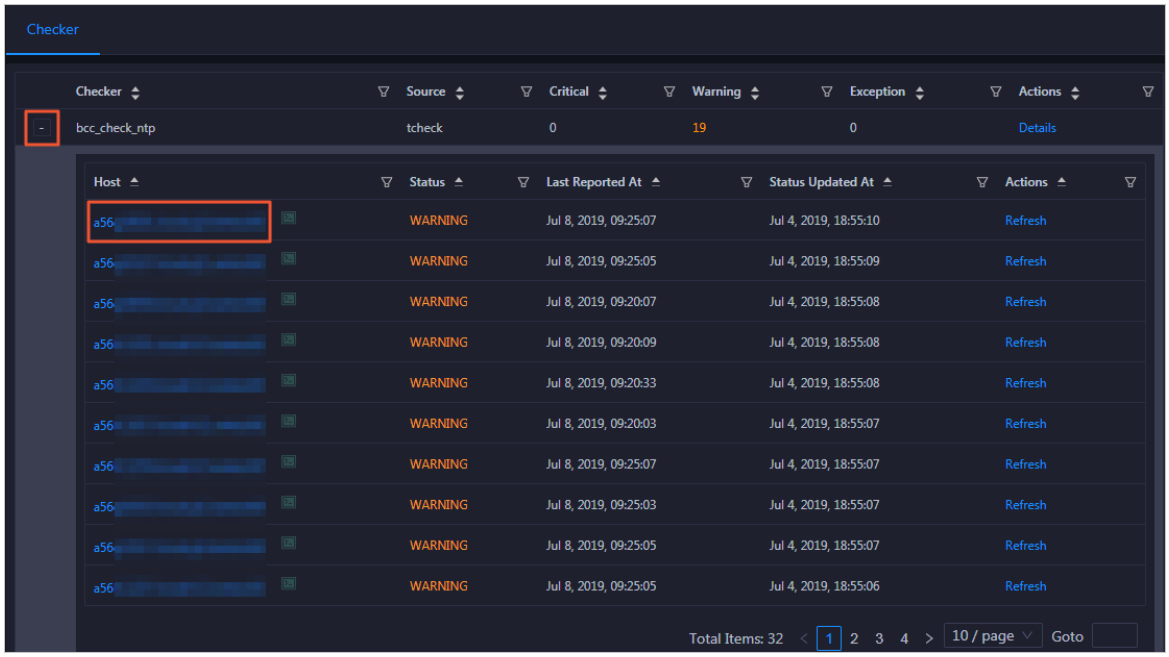

2. Click a hostname. In the panel that appears, click **Details** in the Actions column of a check result to view the cause of the alert.

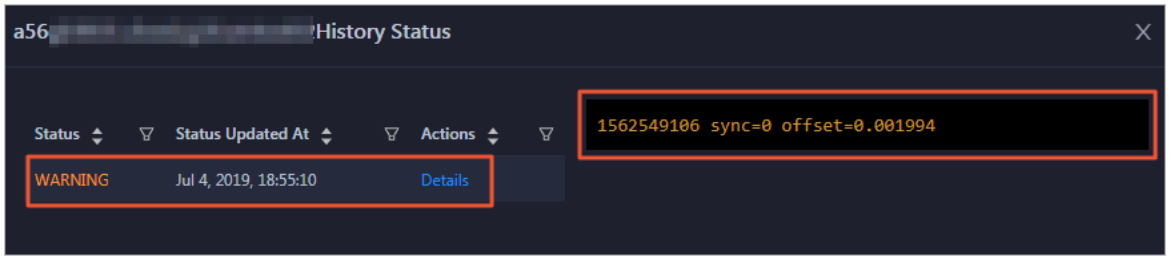

Clear alerts

On the Health Status tab, click **Details** in the Actions column of a checker for which alerts are reported. On the Details page, view the schemes to clear alerts.

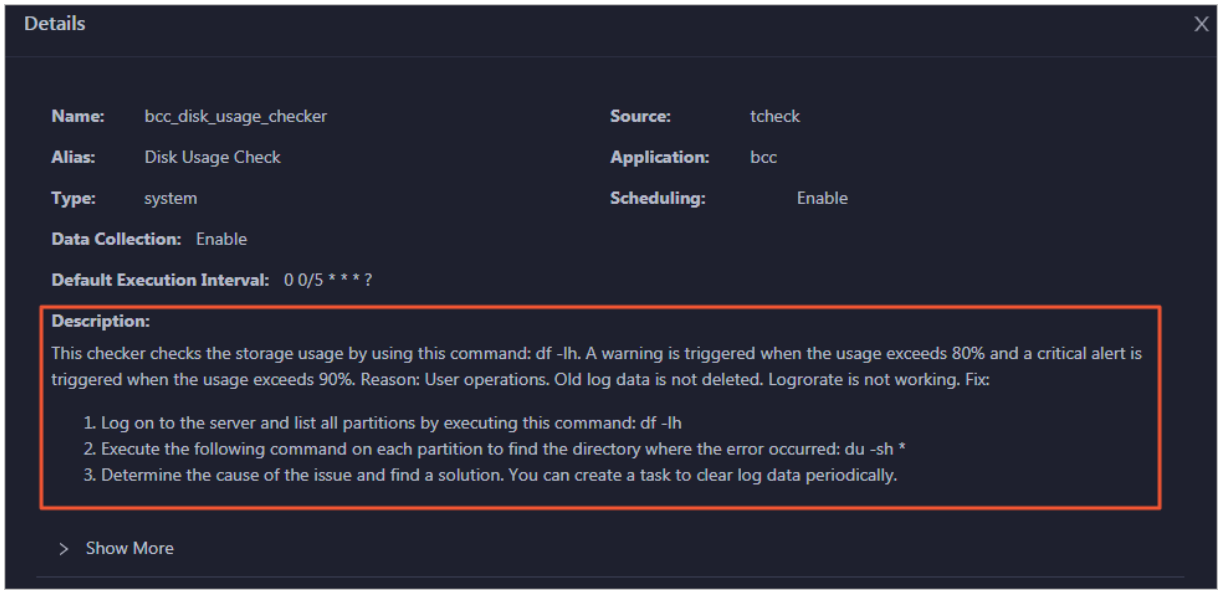

Log on to a host

You may need to log on to a host to handle alerts or other issues that occurred on the host.

1. On the Health Status tab, click + to expand a checker for which alerts are reported.

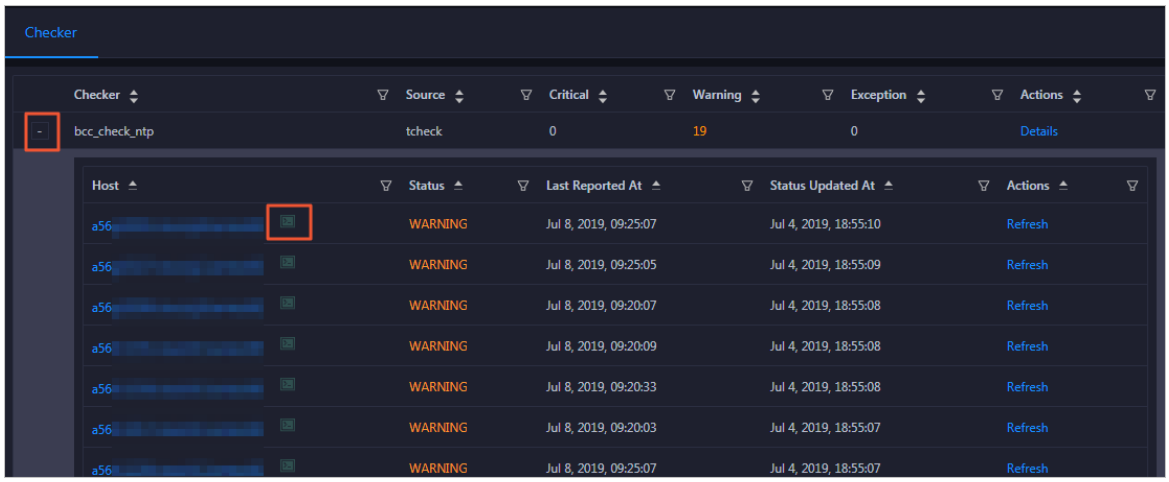

2. Click the **Login in** icon of a host. The **TerminalService** page appears.

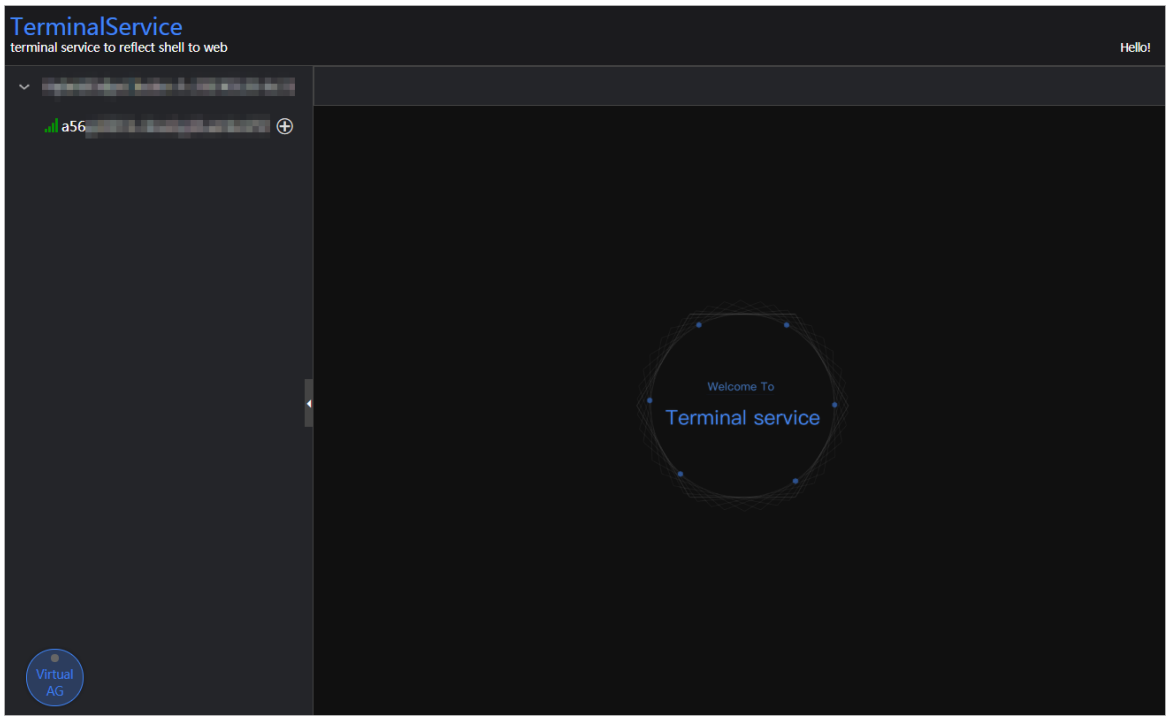

3. On the **TerminalService** page, click the hostname in the left-side navigation pane to log on to the host.

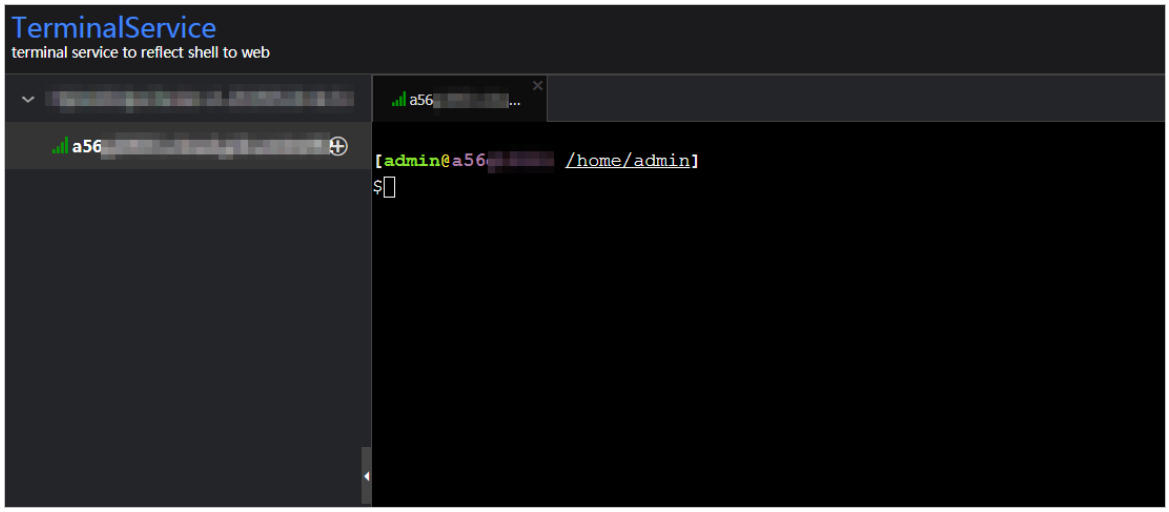

Run a checker again

After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. This way, you can check whether the alert is cleared.

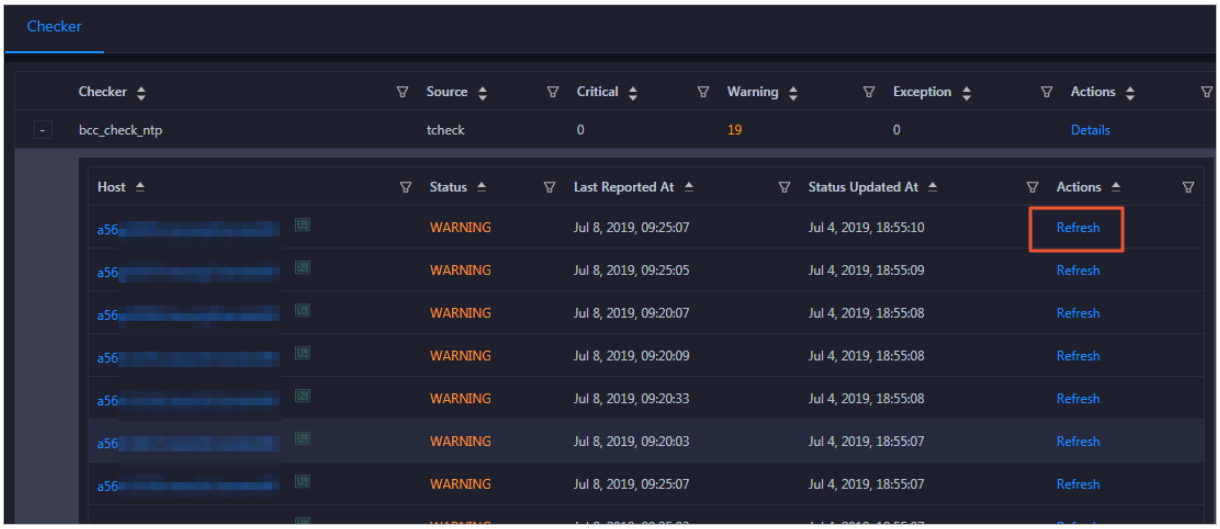

4.5.3. Restore environment settings

If a host in the cluster encounters RPMDB errors, Apsara Big Data Manager (ABM) allows you to restore environment settings.

Prerequisites

bigdata-sre is installed on the machine that you want to manage. If the machine is a Docker container, make sure that the staragent process runs in the container.

Restore environment settings

1. Log on to the ABM console.
2. In the upper-left corner, click the icon and then click **ABM**.
3. In the top navigation bar of the ABM page, click **O&M**. Then, click the **Clusters** tab.
4. In the left-side navigation pane of the **Clusters** tab, select a cluster. Then, click the **Health**

Status tab. The **Health Status** tab appears.

5. In the upper-right corner of the tab, click **Actions** and select **Restore Environment Settings**. In the **Restore Environment Settings** pane, enter a hostname. If you enter multiple hostnames, separate them with commas (,).

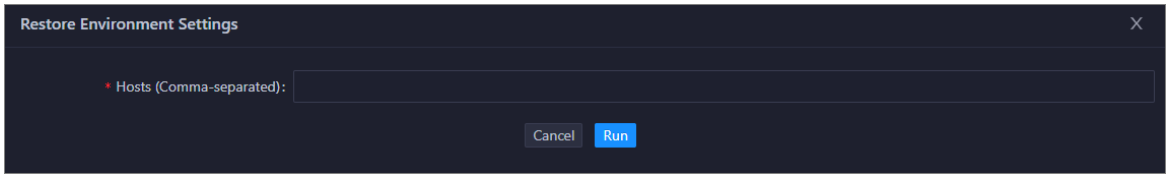

6. Click **Run**.
7. Check the execution status.

   Click **Actions** and select **Execution History** next to **Restore Environment Settings** to view the execution history.

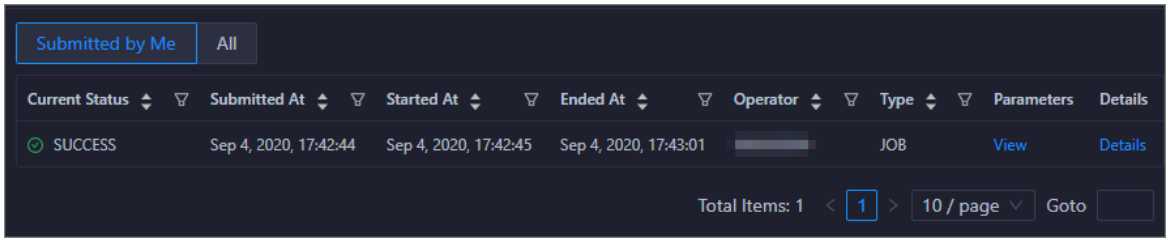

   It requires a long time to restore environment settings. **RUNNING** indicates that the execution is in progress. **SUCCESS** indicates that the execution succeeded. **FAILED** indicates that the execution failed.

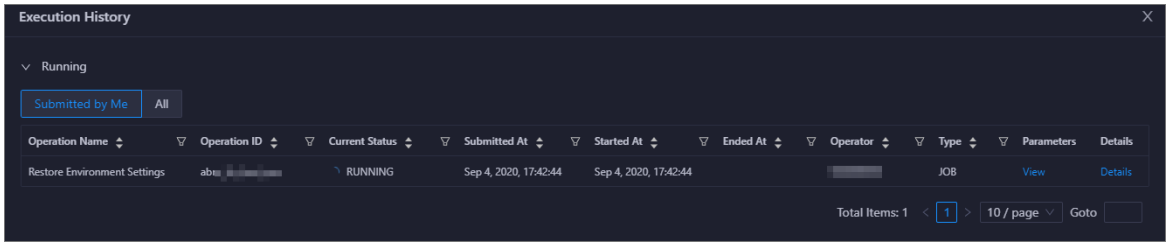

8. If the status is RUNNING, click **Details** in the Details column to view the steps and progress of restoration.

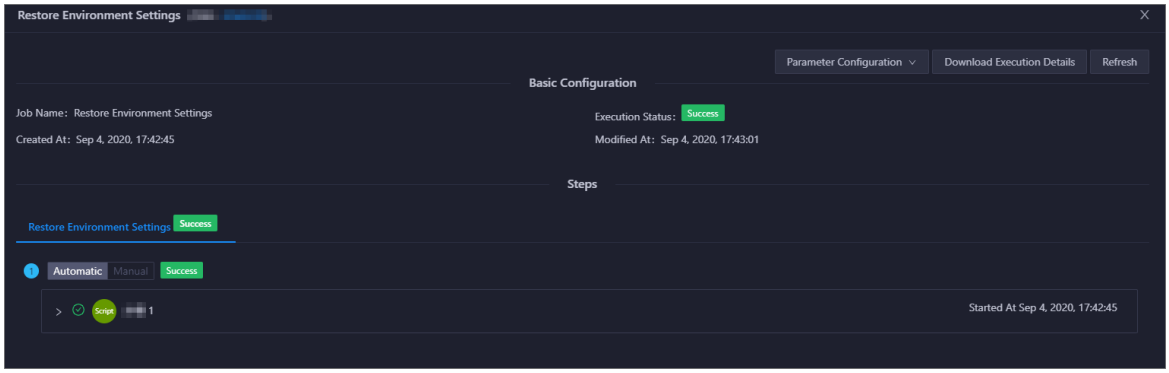

9. If the status is **FAILED**, click **Details** in the Details column to identify the cause of the failure. For more information, see Identify the cause of the failure to restore environment settings.

Identify the cause of the failure to restore environment settings

This section describes how to identify the cause of the failure to restore environment settings.

1. In the upper-right corner of the **Clusters** tab, click **Actions** and select **Execution History** next to **Restore Environment Settings** to view the execution history.
2. Click **Details** in the Details column of a failed record to identify the cause of the failure.

   You can also view information about parameter settings, host details, script, and runtime parameters to identify the cause of the failure.

4.6. Host O&M

4.6.1. Host overview

The host overview page displays the overall running information about a host in an Apsara Big Data Manager (ABM) cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the host.

Entry

On the **Hosts** page, select a host in the left-side navigation pane. The **Overview** page for the host appears.

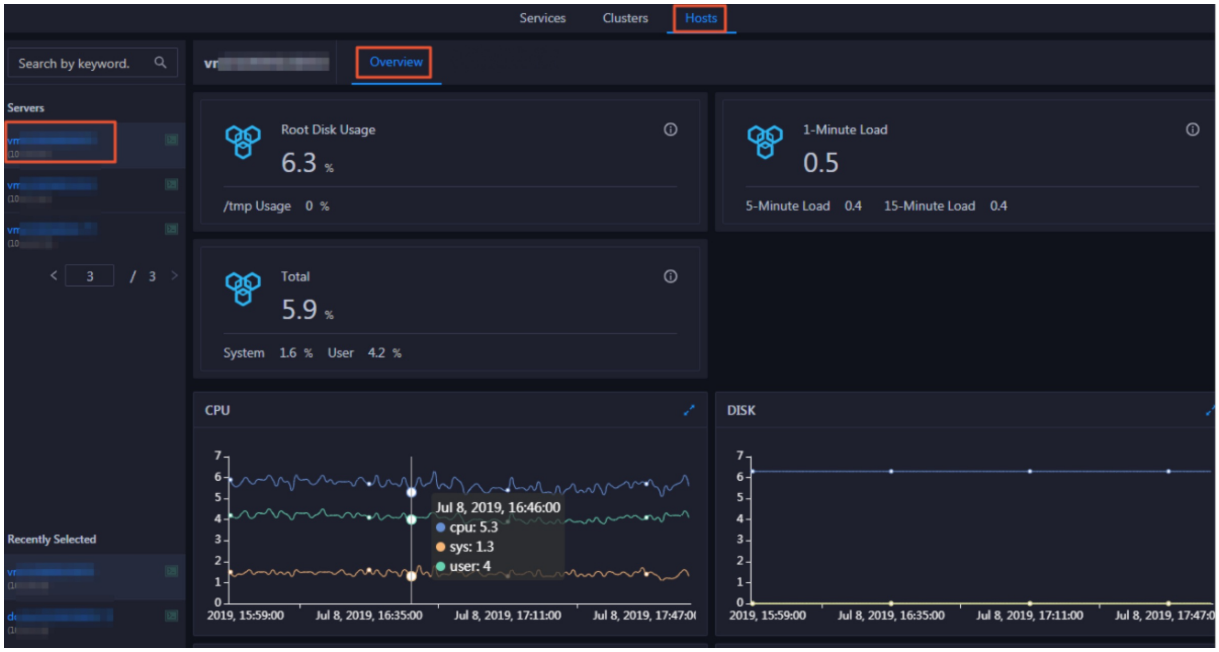

Root Disk Usage, Total, and 1-Minute Load

These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the */tmp* directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.

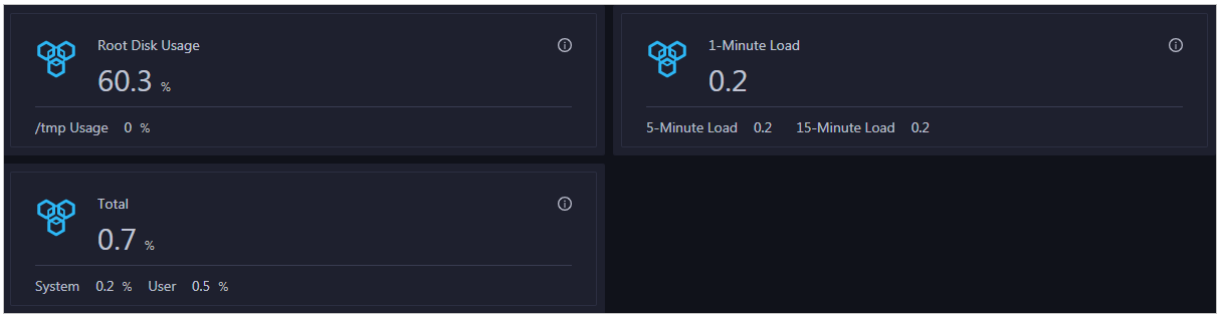

CPU

The CPU chart shows the trend lines of the total CPU utilization (cpu), CPU utilization for executing code in kernel space (sys), and CPU utilization for executing code in user space (user) of the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

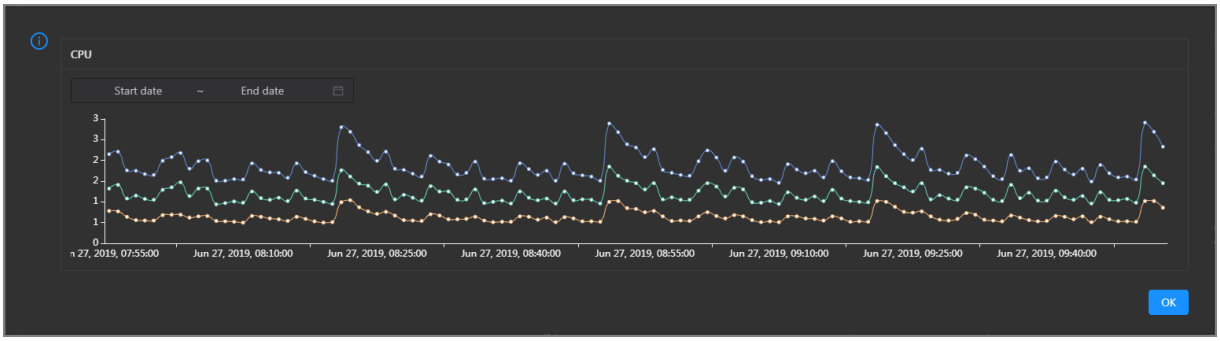

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU utilization of the host in the specified period.

DISK

The DISK chart shows the trend lines of the storage usage in the */*, */boot*, */home/admin*, and */home* directories for the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

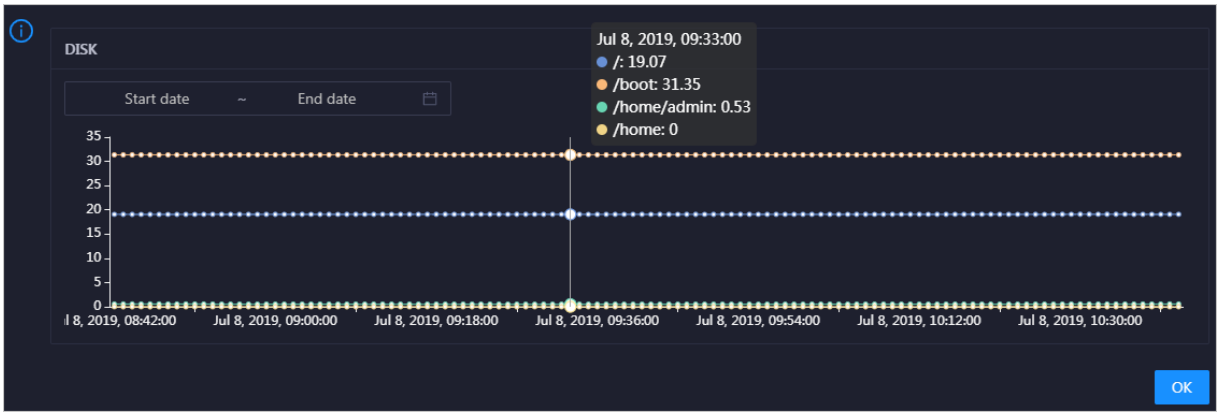

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

The MEMORY chart shows the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

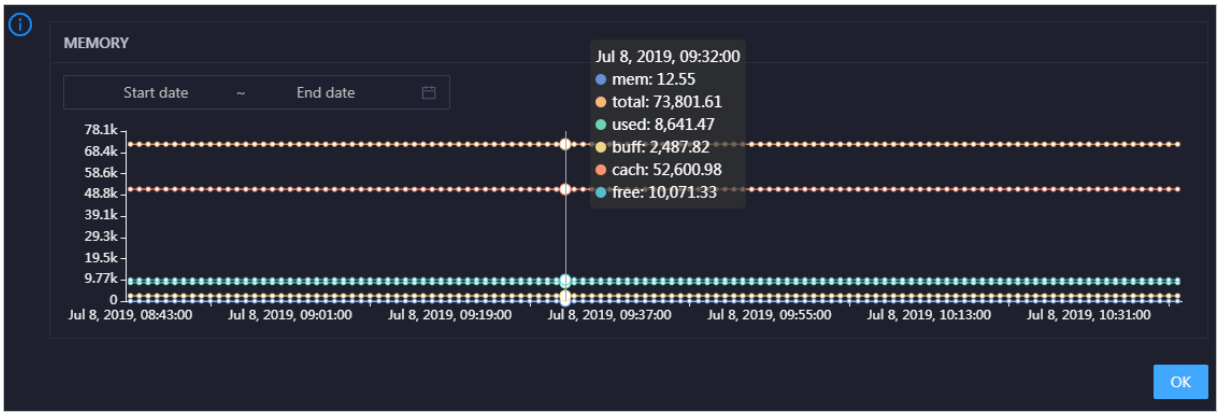

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

The LOAD chart shows the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

In the upper-right corner of the chart, click the icon to zoom in the chart.

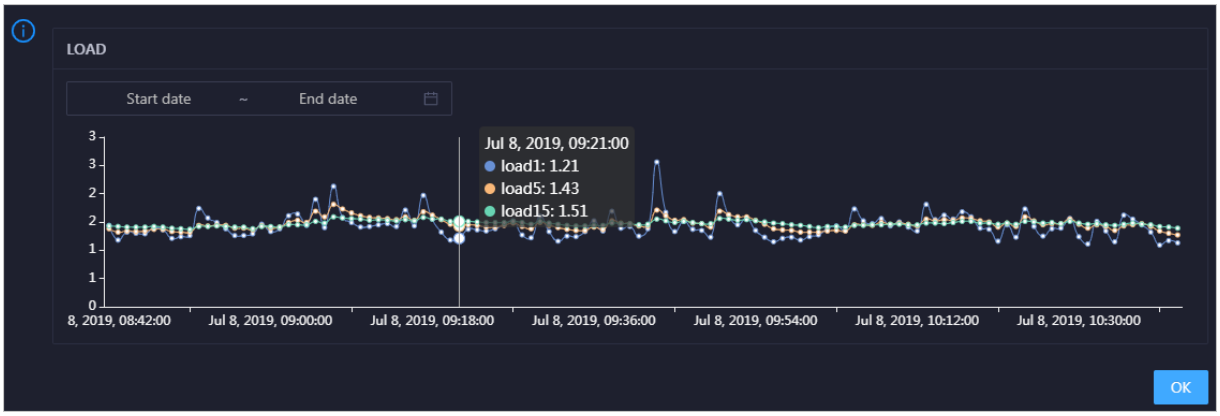

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

The PACKAGE chart shows the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

In the upper-right corner of the chart, click the icon to zoom in the chart.

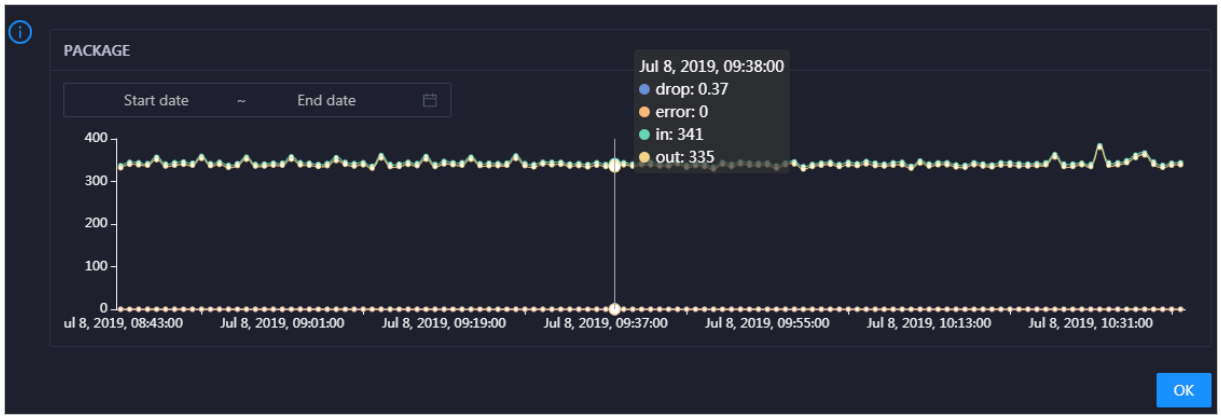

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click in the upper-right corner of the chart to zoom in the chart.

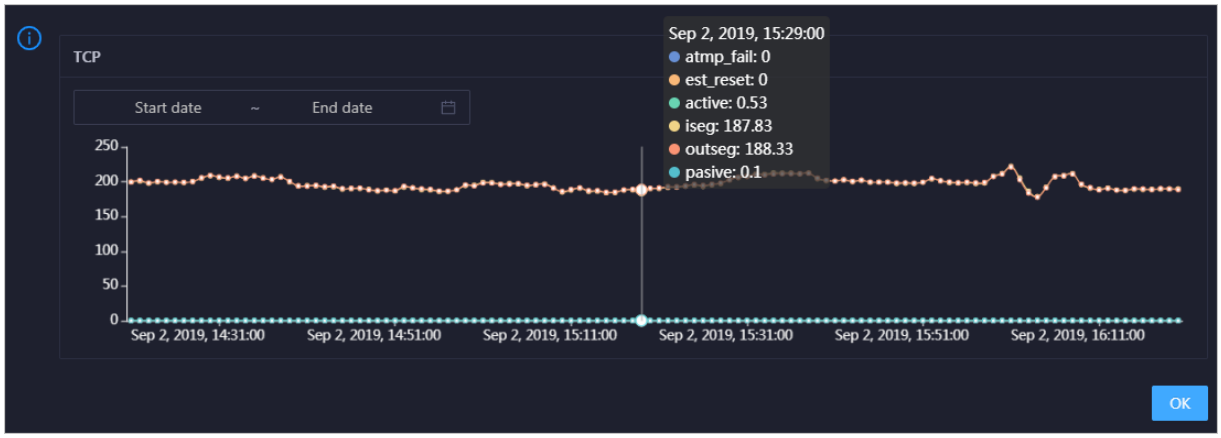

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This chart displays the trend line of the average usage of the root disk (/) for the host over time.

Click in the upper-right corner of the chart to zoom in the chart.

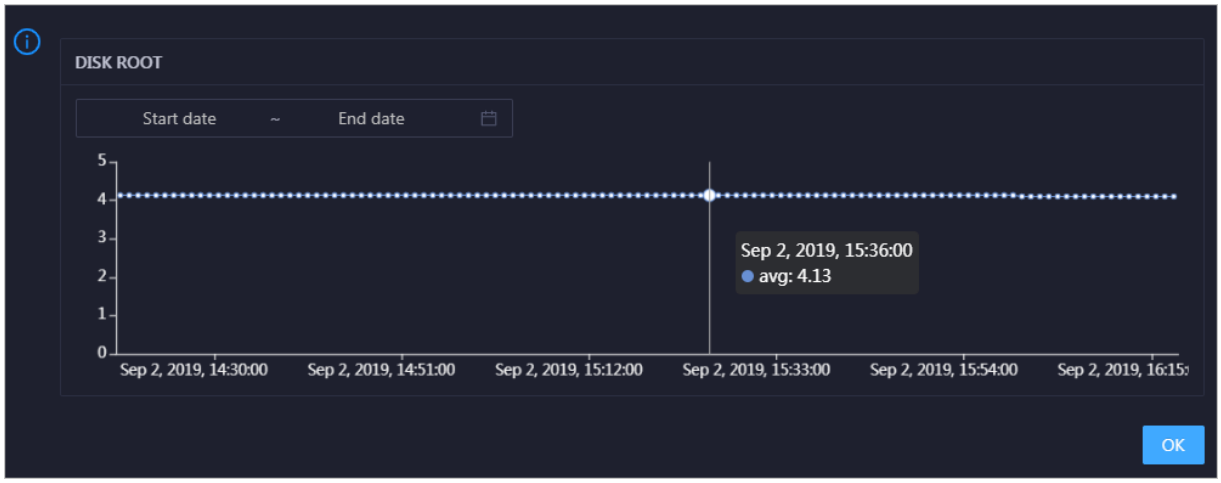

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

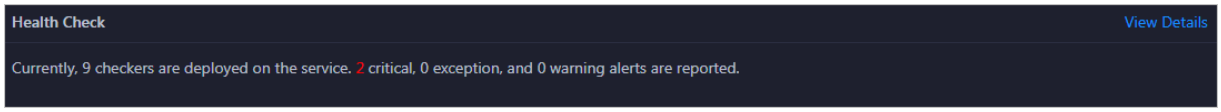

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

Health Check History

This section displays a record of the health checks performed on the host.

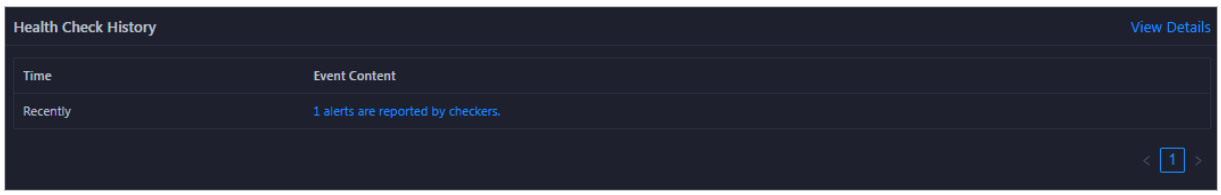

Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see Host health.

You can click the event content of a check to view the exception items.

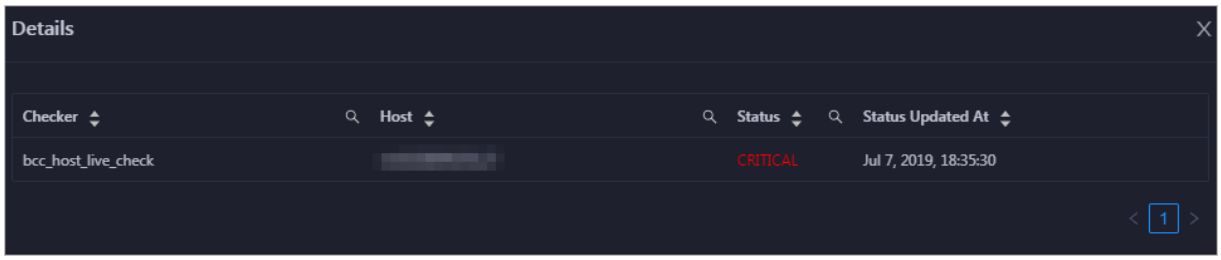

5.Management

5.1. Overview

The management module is the configuration and software management center of Apsara Big Data Manager (ABM). It is an important functional module that supports and customizes O&M items for services.

The management module supports the following features:

- Job execution and management: You can generate jobs based on the scheme library to perform O&M operations on services.
- Patch management: You can deploy upgrade patches for various services.
- Hot upgrade: You can perform hot upgrades on the monitoring configuration and monitoring items of ABM so that services are not interrupted during the upgrade process.
- Health management: You can create health checkers and apply them to service hosts.
- Operation audit: You can view the records of job execution and other service O&M operations in ABM.

5.2. Jobs

5.2.1. Overview

This topic describes the UIs for job management and terms related to jobs in Apsara Big Data Manager (ABM).

ABM allows you to run jobs to perform O&M operations on big data services. Jobs in ABM are run to perform O&M operations on physical devices in a cluster. The Jobs page contains the **Job Execution** page.

Terms

Terms related to jobs include:

- Ordinary job: A job that can only be manually run. No timer is set.
- Cron job: A job that is automatically run based on timer settings.

Job Execution page

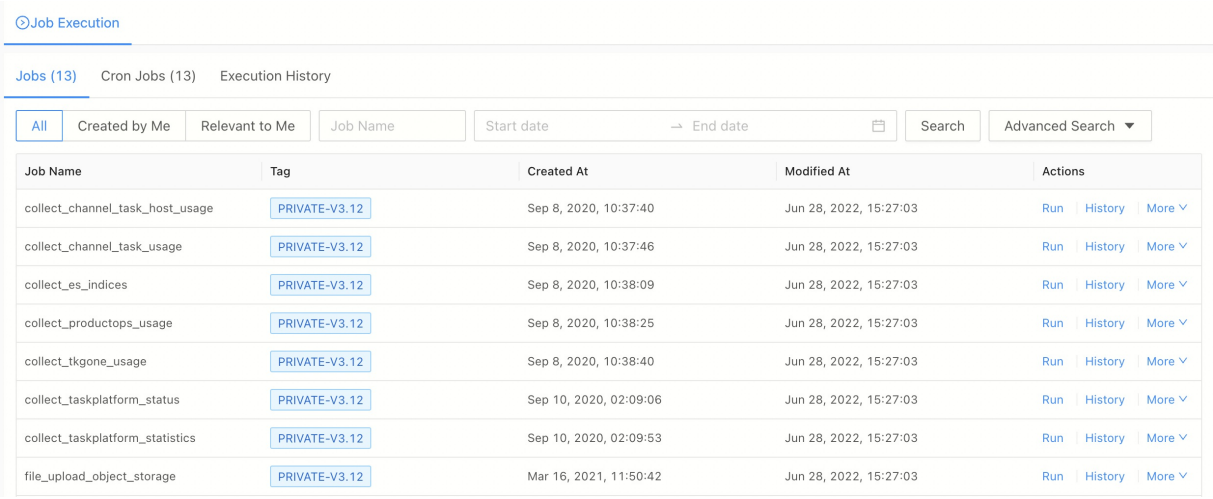

The **Job Execution** page contains the following tabs:

- Jobs

  You can view and run ordinary jobs, and view their execution history.

  You can search for a specific ordinary job.
- Cron Jobs

  You can enable, disable, view, or run cron jobs, and view their execution history.

  You can search for a specific cron job.
- Execution History

  You can view the execution history of ordinary and cron jobs.

  You can specify multiple filter conditions to search for a specific job and view the execution history of the specific job.

5.2.2. Jobs

5.2.2.1. Enable or disable a cron job

When a cron job is generated from a scheme, the job is disabled by default. You must manually enable it. If you do not need the cron job to run during a specified time period, you can manually disable it.

Prerequisites

You must have an ABM administrator account.

Procedure

1. Log on to the ABM console.
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. On the **Job Execution** page, click **Cron Jobs**.

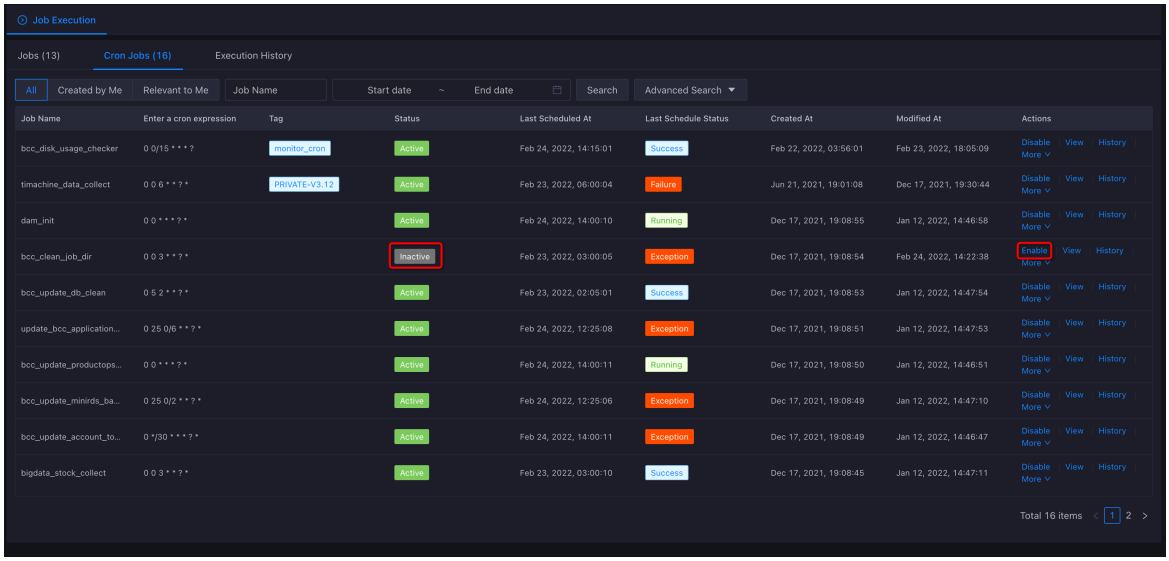

4. On the **Cron Jobs** page, you can enable or disable a cron job.
   - To enable a cron job in the inactive status, click **Enable** in the Actions column of the cron job.

     After a cron job is enabled, its **status** changes to **Active**. The **Enable** button is replaced by **Disable**.
   - To disable a cron job in the active status, click **Disable** in the Actions column of the cron job.

     After a cron job is disabled, its **status** changes to **Inactive**. The **Disable** button is replaced by **Enable**.

5.2.2.2. Manually run a job

After you have created an ordinary job, you must manually run the job in order to perform O&M operations on the product. You can also manually run a cron job.

Prerequisites

You must have an ABM administrator account.

Procedure

1. Log on to the ABM console.
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. Click **Ordinary Jobs** on the **Job Execution** page.

   If you need to manually run a cron job, click **Cron Jobs**. The procedure to manually run a cron job is the same as that of an ordinary job. This topic takes ordinary jobs as an example.

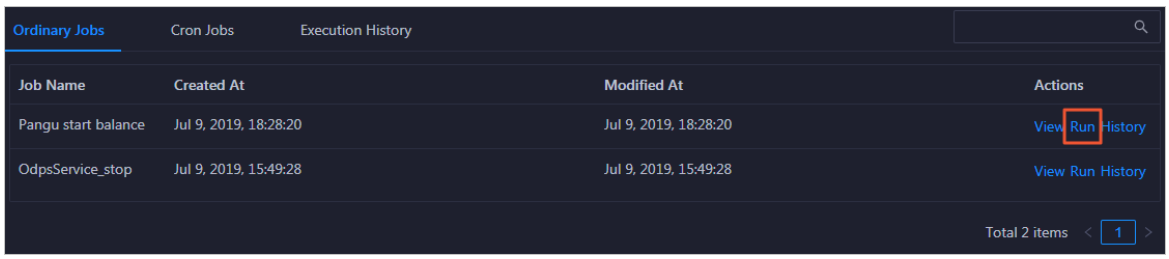

4. In the **Ordinary Jobs** list, click **Run** in the Actions column of a job.
5. Confirm the job risks in the dialog box that appears, and click **Confirm**.

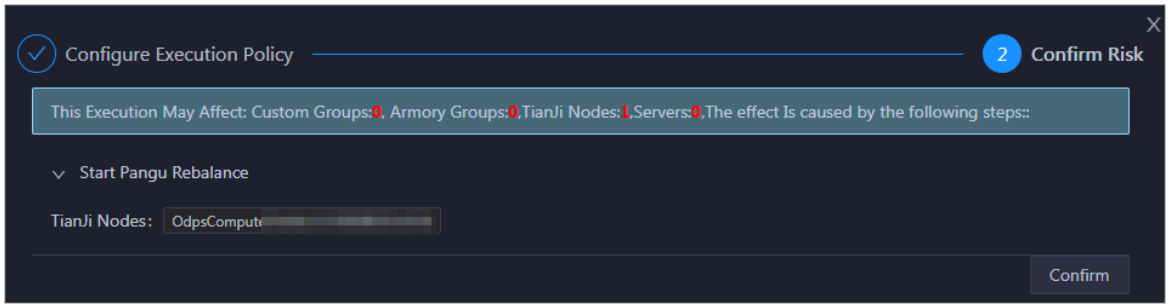

After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see View the execution history.

6. On the job execution page, click **Start** at the top to start the execution.

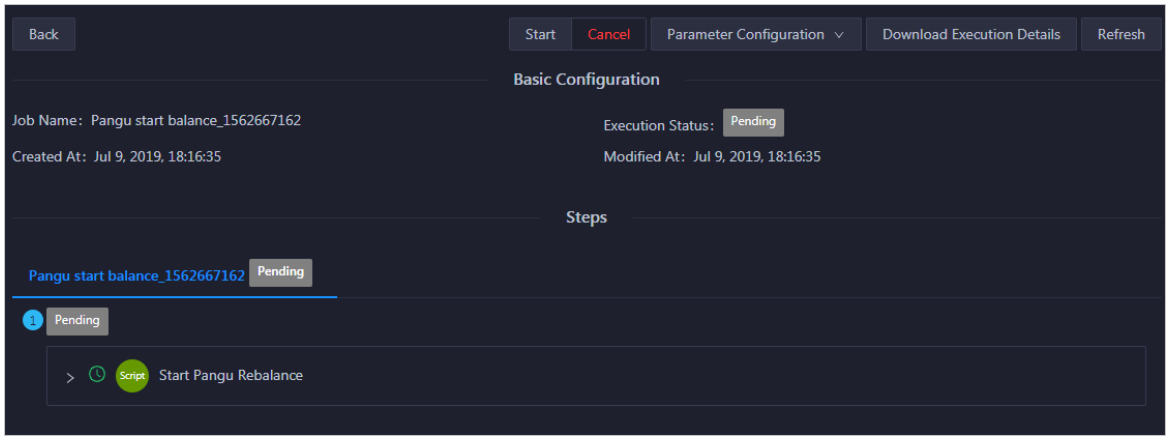

You can find the record about a job on the **Execution History** page, and click **View** to go to the detailed execution page.

5.2.2.3. View the execution history of a job

Apsara Big Data Manager (ABM) allows you to view the execution history of a specific job to learn the execution status of it.

Prerequisites

An ABM administrator account is obtained.

Context

After you confirm to run a job, ABM generates logs for the job execution. You can learn the execution status by using the log data.

The **Execution History** page provides the following features:

- Provides information such as the trigger mode, current status, start time, and end time of each job.
- Provides job execution details and parameter setting information, and allows you to download execution details.
- Allows you to perform certain operations depending on the job status. For example, you can run a job that is in the **Pending** state or retry the execution of a job that is in the **Exception** state.

This topic describes how to view the execution history of an ordinary job. You can follow a similar procedure to view the execution history of a cron job.

Procedure

1. Log on to the ABM console.
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. Click the **Ordinary Jobs** tab on the **Job Execution** page.
4. On the **Ordinary Jobs** page, click **History** in the Actions column of an ordinary job. The **Execution History** page appears.

   You can view the execution history of this job on the **Execution History** page. For more information, see View the execution history.

5.2.3. View the execution history

Apsara Big Data Manager (ABM) allows you to view the execution history of jobs and schemes so that you can learn about their execution details.

Prerequisites

An ABM administrator account is obtained.

Context

After you have confirmed the execution of a job, a record is automatically generated on the Execution History page.

The **Execution History** page provides the following features:

- Provides information such as the trigger mode, current status, start time, and end time of each job.
- Provides job execution details and parameter setting information, and allows you to download execution details.
- Allows you to perform certain operations depending on the job status. For example, you can run a job that is in the **Pending** state or retry the execution of a job that is in the **Exception** state.

Procedure

1. Log on to the ABM console.
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. Click the **Execution History** tab on the **Job Execution** page.

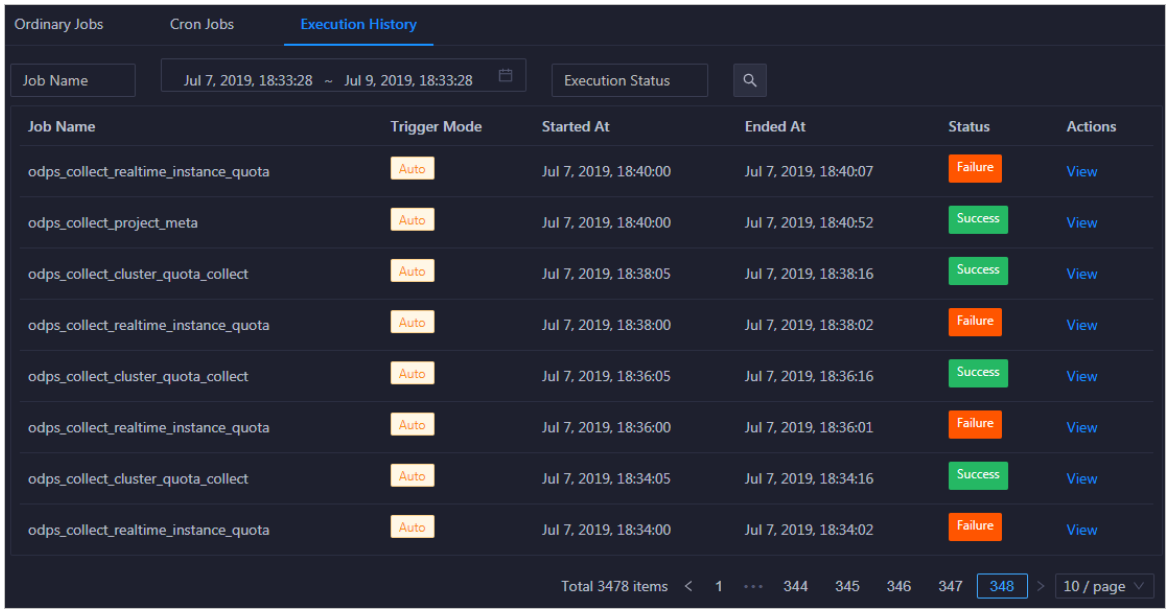

4. If there are too many execution records, filter them by a combination of one or more of the following filter conditions: job name, creator, execution status, and time range. Then, click to search for required records.
5. Click **View** in the Actions column of a record to view the execution details.

   The following table lists the operations that you can perform on records in different states.

Execution status	Feature	Operation
All statuses	View the parameter configuration	Click **Parameter Configuration** at the top, and select **Context Parameters** or **Global Parameters** to view the context parameters or global parameters of the task.
	Download execution details	Click **Download Execution Details** at the top to download the job execution details to the local device. Save it into a TXT file. The execution details record the JSON and raw data of job execution.
	View the execution details of steps	○ On the **Servers** page of a step, click **View Details** in the Actions column of a certain server. The execution details of the step on the server, including the execution output, appear in the Execution Details section. ○ If the step includes a script, the **Script Content** and **Execution Parameters** pages will appear, where you can view the script content and the script execution parameters. ○ If the step includes a command, the **Commands** and **Execution Parameters** pages will appear, where you can view the command content and the command execution parameters.

Execution status	Feature	Operation
	Refresh the page	If the task is in progress, you can click **Refresh** at the top to view the latest execution status.
Pending	Start the execution	Click **Start** at the top to start the execution.
	Cancel the execution	Click **Cancel** at the top to cancel the execution.
Unconfirmed	Complete the manual operation	At the manual step to be operated, follow the instructions and click **OK** to go to the next step.
	Roll back to the complete status of the previous step	At the manual step to be operated, click **Rollback** to roll back to the complete status of the previous step.
	Cancel the execution	Click **Cancel** to cancel the execution.
Exception	Retry the step with exceptions	At the step with exceptions, click **Retry** to execute the step again.
	Skip the step with exceptions	At the step with exceptions, click **Skip** to skip this step and execute the subsequent steps.
	Roll back to the complete status of the previous step	At the step with exceptions, click **Rollback** to roll back to the complete status of the previous step.
	Reset the step with exceptions to the Pending state	At the step with exceptions, click **Reset** to reset the step to the **Pending** state. When the step with exceptions is reset to the Not Started state, the execution status becomes **Paused**. You can click **Continue** at the top to execute the step again.

Execution status	Feature	Operation
	View the execution details of steps with exceptions	◦ On the **Servers** page of a step, click **View Details** in the Actions column of a certain server. The execution details of the step on the server, including the execution output and error message, appear in the Execution Details section. After you have viewed the details of the server with exceptions during the execution, you can click **Skip** to skip this server. Alternatively, you can click **Retry** to execute the step again on the server. ◦ If the step includes a script, the **Script Content** and **Execution Parameters** pages will appear, where you can view the script content and the script execution parameters. ◦ If the step includes a command, the **Commands** and **Execution Parameters** pages will appear, where you can view the command content and the command execution parameters.
Failure	Retry the failed step	At the failed step, click **Retry** to execute the step again.
	Skip the failed step	At the failed step, click **Skip** to skip this step and execute the subsequent steps.
	Roll back to the complete status of the previous step	At the failed step, click **Rollback** to roll back to the complete status of the previous step.
	Reset the failed step to the Pending state	At the failed step, click **Reset** to reset the step to the **Pending** state. When the failed step is reset to the Not Started state, the execution status becomes **Paused**. You can click **Continue** at the top to execute the step again.

Execution status	Feature	Operation
	View the execution details of failed steps	◦ On the **Servers** page of a step, click **View Details** in the Actions column of a certain server. The execution details of the step on the server, including the execution output and error message, appear in the Execution Details section. After you have viewed the details of the server with exceptions during the execution, you can click **Skip** to skip this server. Alternatively, you can click **Retry** to execute the step again on the server. ◦ If the step includes a script, the **Script Content** and **Execution Parameters** pages will appear, where you can view the script content and the script execution parameters. ◦ If the step includes a command, the **Commands** and **Execution Parameters** pages will appear, where you can view the command content and the command execution parameters.
	Cancel the execution	Click **Cancel** at the top to cancel the execution.

5.3. Health management

Apsara Big Data Manager (ABM) provides a wide range of built-in scheduling items and monitoring items for each service. These items check service faults and send alerts when necessary, enabling you to detect and fix service faults in time.

Prerequisites

- Your ABM account is granted the required permissions to perform O&M operations on the corresponding service.
- The alert sources and checkers of the monitoring items are obtained.

Background

Different services have different scheduling and monitoring items, but their configuration and operations are the same. This topic uses MaxCompute as an example.

Scheduling: You can run checkers on all hosts of a specified Apsara Infrastructure Management Framework role as scheduled to generate raw alert data. The raw alert data includes the checker, host, alert severity, and alert information. ABM stores the raw alert data in its database.

Monitoring: You can mount checkers to service pages in ABM. When mounting a checker to a service page, you can set a filter policy to display only required alerts.

Both the scheduling items and monitoring items are built-in and cannot be added. However, you can modify some parameters of the items, such as whether to enable an item, running parameters, and description. In addition, you can configure mount points of the monitoring items or delete monitoring items.

View details and mount points of scheduling items

The mount points of scheduling items are built-in and cannot be added, modified, or deleted. The mount points of the scheduling items correspond to the list of all hosts corresponding to the Apsara Infrastructure Management Framework role that runs the scheduling script.

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page. The **Scheduling** page appears.

   The **Scheduling** page displays all scheduling items of the current service.
4. On the **Scheduling** page, click **View** in the Actions column of a scheduling item to view the details.

   The details of a scheduling item include the name, alias, description, alert cause, and alert solution.
5. Click + to expand a scheduling item, and then view the mount points of the scheduling item.

Modify a scheduling item

You can set the scheduling interval and running parameters of a scheduling item, and set whether to enable the scheduling item.

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page. The **Scheduling** page appears.
4. On the **Scheduling** page, click **Edit** in the Actions column of a scheduling item. In the dialog box that appears, set relevant parameters.

   Type: The value **System Default** indicates that parameters such as **Execution Interval** and **Parameters** use the default settings. The value **Custom** indicates that the parameters can be customized.

   > **Note** Set the **Execution Interval** parameter based on the **crontab** command.
5. Click **OK**. The system prompts that the configuration has been modified.

View faulty hosts

You can view all the faulty hosts in the current cluster.

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page. The **Scheduling** page appears.
4. Click **Faulty Servers** in the upper-right corner to view the faulty hosts in the cluster.

The faulty host list displays all faulty hosts in the current cluster and the Apsara Infrastructure Management Framework role of each host.

Modify a monitoring item

You can modify the name and description of a monitoring item and determine whether to enable it. The alert sources and checkers of monitoring items are built-in. Do not modify them.

1. Log on to the ABM console.
2. Click in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page.
4. On the **Health Management** page, click the **Monitoring** tab. The **Monitoring** page appears.

   The **Monitoring** page displays all monitoring items of the current service.
5. On the **Monitoring** page, click **Modify** in the Actions column of a monitoring item to modify its configuration.
6. Click **OK**. The system prompts that the configuration has been modified.

Add a mount point for a monitoring item

After a mount point is added for a monitoring item, the monitoring item mounts the raw alert data to the O&M page of each service in the ABM console.

1. Log on to the ABM console.
2. Click in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page.
4. On the **Health Management** page, click the **Monitoring** tab. The **Monitoring** page appears.
5. On the **Monitoring** page, click + to expand a monitoring item, and then view the mount points of the monitoring item.
6. Click **Add Mount Point** under the mount point list. In the dialog box that appears, set relevant parameters.

   The following table describes some key parameters.

Parameter	Description
Mount Point	The mount point to which the required inspection result of this monitoring item is to be mounted. For example, the value **odps/host** indicates that the result is mounted to the host O&M page of MaxCompute.
Filter Policy	Valid values: ○ **None**: Display all alerts generated by the monitoring item. ○ **Custom**: Display the alerts generated by the monitoring item in accordance with the filter configured for the service tree node. ○ **Node Name**: Display the alerts whose node name is the same as the name of the current node.

Parameter	Description
Enabled	Specifies whether the mount point takes effect.

7. Click **OK**. The system prompts that the configuration has been modified.

Delete a mount point for a monitoring item

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page.
4. On the **Health Management** page, click the **Monitoring** tab. The **Monitoring** page appears.
5. On the **Monitoring** page, click + to expand a monitoring item, and then view the mount points of the monitoring item.
6. Click **Delete** in the Mount Point column of the mount point to be deleted. In the dialog box that appears, click **OK**. The system prompts that the deletion is successful.

Delete a monitoring item

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the **Management** page.
4. On the **Health Management** page, click the **Monitoring** tab. The **Monitoring** page appears.
5. Click **Delete** in the Actions column of the monitoring item to be deleted. In the dialog box that appears, click **OK**. The system prompts that the deletion is successful.

5.4. Operation auditing

This feature allows you to view the O&M operations of the current service of Apsara Big Data Manager (ABM). The details of each operation are provided for retrieval and fault locating.

Prerequisites

Your ABM account is granted the required permissions to perform O&M operations on the corresponding service.

Background

You can view operation logs by service. For example, to view the operation logs of MaxCompute, you must go to the MaxCompute page first. The following describes how to view the operation logs of MaxCompute.

> **Note** This page displays only the O&M operations of a service. Note that the O&M operations of job services are not included.

Procedure

1. Log on to the ABM console.
2. Click ≡ in the upper-left corner, and then click **MaxCompute**.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Operation Audit** in the left-side navigation pane of the **Management** page.

   The **Operation Audit** page displays the O&M operations of the current service. In this example, the information about MaxCompute O&M operations is displayed, including the operation name, operation ID, status, submission time, start time, end time, operator, and implementation method.
4. Click **Details** for an operation to view the O&M operation details.

   You can also view the causes of failed steps in detail.
5. If an O&M operation fails, view the cause of the failure.
6. When the task is in the Failure, Not Started, Pending, or Exception state, perform the operations listed in the following table based on your situation.

State	Executable operation
Not Started	○ Click **Start** to start the task. ○ Click **Parameter Configuration** to view the parameter configuration of the task. ○ Click **Cancel** to cancel the task.
Pending	○ Follow the instructions and click **OK** to go to the next step. ○ Click **Rollback** to roll back to the complete status of the previous step. ○ Click **Parameter Configuration** to view the parameter configuration of the task. ○ Click **Cancel** to cancel the task.
Exception	○ Click **Retry** to run the step again. ○ Click **Skip** to skip this step and execute the subsequent steps. ○ Click **Rollback** to roll back to the complete status of the previous step. ○ Click **Parameter Configuration** to view the parameter configuration of the task. ○ Click **Cancel** to cancel the task.

State	Executable operation
Failure	◦ Click **Retry** to run the step again. ◦ Click **Skip** to skip this step and execute the subsequent steps. ◦ Click **Rollback** to roll back to the complete status of the previous step. ◦ Click **Parameter Configuration** to view the parameter configuration of the task. ◦ Click **Cancel** to cancel the task.

7. To download the O&M operation execution logs, click **Download Execution Details** at the top to save the logs to your local device.

6.Go to other consoles

Apsara Big Data Manager (ABM) provides links to the Apsara Uni-manager Operations Console (ASO), Apsara Infrastructure Management Framework console, and Apsara Stack Security Center console to facilitate the operations and maintenance (O&M) of big data services.

Prerequisites

An ABM account that works properly and its password are obtained.

Procedure

1. Log on to the ABM console.
2. On the homepage of ABM, click the icon in the upper-left corner. In the **Site Navigation** section, click **ASO**, **TIANJI**, or **YUNDUN** to go to the console that you want to access.

Result

After you click **ASO** or **TIANJI**, you can log on to the Apsara Uni-manager Operations Console or the Apsara Infrastructure Management Framework console without the need to enter your username and password.

After you click **YUNDUN**, you must enter the username and password to log on to the Apsara Stack Security Center console.